Understanding Drug Resistance through Computational Models of the Genotype–Phenotype Mapping

Daniel Nichol

Supervisors: Prof. Peter Jeavons, Dr. Alexander RA Anderson

Department of Computer Science
University of Oxford

Thesis submitted for the degree of Doctor of Philosophy

University College November 2018

For Sophie

Declaration

I declare that all work presented in thesis was produced by myself, except where explicitly stated in the text. No portion of this thesis has been submitted for any other degree or qualification.

Daniel Nichol
November 2018

Acknowledgements

The submission of this thesis concludes eight years as a student in Oxford. I am honoured to have spent all eight of those years under the supervision of Peter Jeavons, who believed in me when I didn't, taught me almost everything I know, and shaped my life indescribably. I am immensely grateful for the guidance you have provided and look forward to continuing our work as colleagues and friends.

This thesis was made possible by the support and intellectual guidance of my supervisor, and now dear friend, Sandy Anderson. I am especially grateful for the chance to spend a year in the department of Integrated Mathematical Oncology at the Moffitt Cancer Center, where I learned the clinical significance of mathematical modelling. I would like to thank all members of the IMO, who made Tampa feel like home.

I must acknowledge the contribution of Jacob Scott, who enthusiastically introduced me to the field of mathematical oncology when I could barely solve a differential equation, and whose boundless optimism served as the perfect antidote to my pessimistic disposition. The modelling presented in this thesis benefited greatly from discussions with Mark Robertson–Tessi, Alexander Fletcher, Jill Gallaher and Philip Gerlee. I would also like to thank the lab of Dr Robert Bonomo, in particular Chris Bryant and Joseph Rutter, for their experimental efforts.

For their support, I thank my parents, Elaine and Zain Nichol, who must have thought I'd never finish school, but never seemed to mind. Finally, and most importantly, I thank Sophie Hibbin for four years of inexhaustible patience, support and tea.

Abstract

The emergence of drug resistance is ultimately driven by Darwinian evolution. These evolutionary dynamics are inherently two-tiered, with mutational processes at the genetic scale inducing variation in cellular phenotypes that are subject to natural selection. If we are to predict or reverse evolution, as we must to determine effective treatments for drug–resistant infections and cancers, then we must first understand the relationship between genetic and phenotypic change. This relationship, known as the genotype–phenotype (GP) mapping, is governed by a complex cascade of potentially stochastic molecular interactions that integrate genetic, epigenetic and environmental factors to determine cellular phenotype. In this thesis, we introduce abstract models of the GP–mapping to explore how its structure influences the evolution of drug resistance, and determines the efficacy of novel therapeutic strategies such as drug holidays or adaptive therapy. We begin by providing a comprehensive review of previous abstract modelling of the GP–mapping (Chapter 2). In Chapter 3, we demonstrate, by means of a fitness landscape model for the GP–mapping, that through careful selection of drug sequences evolution can be 'steered' such that a highly drug resistant population does not emerge. In Chapter 4, we introduce a novel model for the GP–mapping wherein phenotypes are determined by the stochastic simulation of bistable chemical reaction networks (CRNs). Through this model we explore the evolutionary dynamics of ecological 'bet–hedging', a common driver of drug resistance wherein phenotypes are determined stochastically, demonstrating that the structure of the GP–mapping can slow the evolutionary loss of this trait, preserving the survival mechanism when harsh environments occur very infrequently. Thus, the capacity to steer the evolutionary loss of bet–hedging through drug holidays is dependent on the structure of the GP–map. In Chapter 5, the CRN model is extended to account for epigenetic inheritance and the potential for phenotypic memory in bet–hedging. Critically, we find that genetics, epigenetics and the GP–mapping interact non–additively to determine organismal fitness, indicating that evolution likely cannot be predicted without accounting for each of these biological processes. Finally, in Chapter 6, we explore how properties of the GP–mapping are manifested at the population scale and suggest experimental approaches to identify bet–hedging and phenotypic plasticity from population measurements. The implications of our results for the treatment of drug–resistant diseases are explored throughout.

Table of contents

List of figures **xv**

List of tables **xix**

1 Introduction **1**
1.1 Drug Resistance and Evolution 1
1.2 Mechanisms of Drug Resistance 1
1.2.1 Antibiotic Resistance 1
1.2.2 Cancer Drug Resistance and Targeted Therapy 3
1.3 The Adaptive Therapy Paradigm 6
1.3.1 Predicting Evolution and the Genotype–Phenotype Mapping 7
1.4 Thesis Overview 8
1.5 Scientific Contributions and Output 9
1.5.1 Scientific Contributions 9
1.5.2 Papers, Preprints and Manuscripts in Preparation 11

2 Genotype–Phenotype Mapping **13**
2.1 Evolution and the GP–Map through the Algorithmic Lens 16
2.1.1 The Algorithmic Structure of Evolution 17
2.1.2 Fitness and Population Dynamics 18
2.2 Mathematical Models of Evolution and the GP–Map 20
2.2.1 Fisher's Geometric Model 20
2.2.2 Fitness Landscapes 23
2.2.3 Secondary Structure Prediction 29
2.2.4 Network Models 38
2.3 The Future of GP–Mapping Models 46
2.3.1 Properties of the Genotype–Phenotype Mapping 46
2.3.2 Conclusions 50

3 Steering Evolution with Sequential Therapy **53**
3.1 A Mathematical Model of Evolution on Fitness Landscapes 53
3.1.1 Encoding Evolution as a Random Walk 55
3.1.2 Model Parametrisation 58

3.2 Simple Algebraic Results for Evolution . . . 62
3.2.1 Non–Commutativity of Evolution and the Importance of Drug Ordering 62
3.2.2 Irreversibility and Cycling of Evolutionary Trajectories . . . 63
3.3 Implications for Collateral Sensitivity Analysis . . . 64
3.3.1 Collateral Response is not Necessarily Repeatable . . . 65
3.3.2 Collateral Sensitivity Likelihood and Worst Case Response . . . 69
3.4 Deriving Optimal Sequential Therapies from Empirical Landscapes . . . 71
3.4.1 Evolutionary Steering to Avoid Resistance . . . 73
3.4.2 The Dangers of Arbitrary Drug Ordering . . . 76
3.5 The Impact of Variance in Fitness Measurements . . . 76
3.6 Implementation Details and Computational Considerations . . . 80
3.6.1 Performance . . . 80
3.6.2 Testing . . . 80
3.7 Discussion . . . 83
3.7.1 The Pitfalls of Collateral Response Analysis . . . 83
3.7.2 Model–Driven Evolutionary Steering . . . 85
3.7.3 Variance in Fitness Measurements . . . 86
3.7.4 Future Directions . . . 86

4 Modelling Non–Genetic Heterogeneity 91
4.1 Introduction . . . 91
4.1.1 The Causes of Bet–Hedging . . . 92
4.2 Previous Mathematical Modelling of Bet–Hedging and Non–Genetic Heterogeneity . . . 93
4.3 A Stochastic Genotype–Phenotype Mapping Model . . . 95
4.3.1 The Model Genotype–Phenotype Mapping . . . 96
4.3.2 Examples of CRN Genotype–Phenotype Mappings . . . 98
4.3.3 The CRN Model can Approximate Arbitrary Hedges . . . 100
4.4 Robustness and Redundancy . . . 100
4.5 Implications of the GP–Map for the Evolution of Bet–Hedging . . . 103
4.5.1 Population Dynamics . . . 103
4.5.2 Invasion Dynamics . . . 105
4.5.3 The Genotype–Phenotype Map Preserves Bet–Hedging . . . 107
4.5.4 Robustness to Parameter Variation . . . 113
4.6 Implications for Sequential Therapies and Treatment Holidays . . . 114
4.7 Implementation Details and Computational Considerations . . . 121
4.7.1 Performance . . . 121
4.7.2 Testing . . . 122
4.8 Discussion . . . 122
4.8.1 Implications for the Drivers of Bet–Hedging . . . 123
4.8.2 Implications for Therapy . . . 124

4.9 Future Directions . . . 126

5 Mechanisms of Phenotypic Memory 127
5.1 Bet–Hedging, Epigenetic Inheritance and Phenotypic Memory . . . 127
5.2 Extending the CRN Model to Account for Phenotypic Memory . . . 131
5.2.1 Intra–Cellular Dynamics . . . 131
5.2.2 Simulating Population Dynamics . . . 137
5.3 Parametrisation . . . 144
5.3.1 Robustness to Parameter Variation . . . 146
5.4 The Selective Advantage of Memory . . . 147
5.4.1 Fixed Environments . . . 148
5.5 Memory in Fluctuating Environments . . . 153
5.6 Implications for Therapy . . . 155
5.6.1 Maximum Dose Density Fails for Bet–Hedging Populations . . . 156
5.6.2 Metronomic/Cyclic Therapy . . . 157
5.7 Memory Decay and Optimal Therapies . . . 163
5.7.1 Dose Skipping Adaptive Therapy . . . 168
5.8 Implementation Details and Computational Considerations . . . 172
5.8.1 Performance . . . 172
5.8.2 Testing . . . 173
5.9 Discussion . . . 174
5.9.1 The Evolution of Bet–Hedging in Constant Environment . . . 175
5.9.2 Fluctuating Environments . . . 176
5.9.3 Implications for Therapy . . . 177
5.9.4 Conclusions . . . 180

6 The Impact of Genotype–Phenotype Mappings at the Population Scale 183
6.1 Introduction . . . 183
6.2 Impact of Exponential and Stationary Growth . . . 184
6.2.1 Modelling Growth Phases . . . 185
6.2.2 Results . . . 186
6.2.3 Implications for Experimental Evolution . . . 190
6.3 Mutational Spectra and GP-Mappings . . . 190
6.3.1 A Model of Mutational Accumulation . . . 192
6.3.2 Mutational Signatures of Genetic Resistance and Bet–Hedging . . . 200
6.3.3 Conclusions . . . 203
6.4 Implementation Details and Computational Considerations . . . 204
6.4.1 Performance . . . 204
6.4.2 Testing . . . 204
6.5 Discussion . . . 204

7 Conclusions and Future Work **207**
7.1 Evolutionary Steering and Fitness Landscapes 207
7.1.1 Empirical Fitness Landscapes 208
7.1.2 Extensions to the Model 209
7.2 Models of Bet–Hedging 211
7.2.1 Phenotypic Memory 212
7.2.2 Combining Phenotypic Plasticity and Bet–Hedging 213
7.3 Identifying Model Systems 214
7.4 Summary and Final Conclusions 215

References **219**

Appendix A Additional Figures **241**
A.1 Scatterplots of Fitness Between Landscapes 241

Appendix B Complementary Analysis for Chapter 4 and Chapter 5 **247**
B.1 Parameter Robustness for the Results of Chapter 5 250
B.1.1 Fixed Environments 250
B.1.2 Fluctuating Environments 250

Appendix C Program Code Listing **261**

List of figures

2.1 Cancer and the biological hierarchy . . . 15
2.2 A graphical representation of Fisher's geometric model . . . 22
2.3 Fitness landscapes as an evolutionary system . . . 24
2.4 Forms of pairwise epistasis . . . 25
2.5 The RNA folding model . . . 30
2.6 The stochastic RNA model . . . 35
2.7 Schematic of the gene regulatory model for a GP–Mapping . . . 40
2.8 The fitness function for the GRN model . . . 41
2.9 The neural network model and phenotypic plasticity . . . 43

3.1 Constructing the Markov chain model of evolution . . . 55
3.2 Population dynamics defined by the matrix P. . . . 57
3.3 The 15 fitness landscapes from Mira et al. [228] . . . 61
3.4 Collateral sensitivity need not be repeatable . . . 65
3.5 Correlation of the fitness of genotypes between landscapes . . . 66
3.6 Schematic of the experimental derivation of collateral response . . . 68
3.7 Analysis of the possible collateral sensitivity matrices (CSMs) . . . 69
3.8 Collateral sensitivity likelihood . . . 70
3.9 Thresholded matrices of worst case collateral response . . . 71
3.10 Graphs of fixable mutations in genotype space . . . 72
3.11 The probability distributions for accessibility of the peaks of the AMP landscape for different steering regimes . . . 73
3.12 Histogram of peak numbers over 1000 resampled landscapes . . . 78
3.13 Collateral Sensitivity Likelihoods between landscapes with resampling (rCSLs) 81
3.14 Robustness of evolutionary steering to variance in fitness . . . 82

4.1 Schematic representation of the chemical reaction network model for determining phenotypes from genotypes . . . 94
4.2 Example molecular switches as GP–mappings . . . 99
4.3 The CRN model can closely approximate arbitrary switching probabilities . . . 101
4.4 Redundancy and robustness in the CRN model . . . 102

4.5 Invasion probabilities for a mutant with bet–hedging probability p_1 into a resident population of phenotype p_2 . . . 105
4.6 The switch–dependent relationship between genotype and fitness/invasion probability . . . 108
4.7 Schematic of the evolutionary model . . . 109
4.8 Convergence dynamics through genotype and probability space for the GP–maps defined by DC, DCx, DCy and AM . . . 111
4.9 Comparison of convergence times for increasing $g_{\max}$ . . . 113
4.10 Comparison of convergence times for increasing w_A, $w_B = 1.0$. . . . 114
4.11 The model for simulating treatment holidays . . . 115
4.12 Extinction dynamics for populations endowed with different molecular switches over different timescales of treatment holidays (cytostatic drug) . . . 118
4.13 Extinction dynamics for populations endowed with different molecular switches over different timescales of treatment holidays (cytotoxic drug) . . . 119
4.14 Extinction dynamics for populations endowed with different molecular switches over different timescales of treatment holidays (mixed drug dynamics) . . . 120

5.1 The different modes of bet–hedging . . . 129
5.2 Dynamics of internal molecular state through time . . . 133
5.3 The decay of molecular memory for phenotype A . . . 135
5.4 Dynamics of molecular decay and inheritence . . . 136
5.5 Timing Comparison for the Stochastic and Deterministic Models . . . 139
5.6 Dynamics of the epigenetic memory model . . . 142
5.7 Varying σ alters the the stationary distribution of intracellular states within the population . . . 143
5.8 Memory, molecular switch and genotype combine to determine fitness in constant environments . . . 149
5.9 The fitness landscape for molecular switch, inheritance and genotype in the hospitable environment . . . 150
5.10 The fitness landscape for molecular switch, inheritance and genotype in the drug environment . . . 151
5.11 Memory, molecular switch and genotype combine to determine fitness in fluctuating environments . . . 154
5.12 Maximum–dose–density (MDD) treatment fails for disease with bet–hedging–driven resistance . . . 157
5.13 Growth curves for a bet–hedging population under a range of metronomic therapies . . . 159
5.14 Efficacy of toxicity–restricted cycling strategies on a population with a small proportion of persisters (7day cycles, memory–free phenotype A probability 0.99) . . . 161

5.15 Efficacy of toxicity–restricted cycling strategies on a population with a large proportion of persisters (7day cycles, memory–free phenotype *A* probability 0.5) 162
5.16 Efficacy of toxicity–restricted cycling strategies on a population with a small proportion of persisters (1day cycles, memory–free phenotype *A* probability 0.99) 164
5.17 Efficacy of toxicity–restricted cycling strategies on a population with a large proportion of persisters (1day cycles, memory–free phenotype *A* probability 0.5) 165
5.18 The maximum equivalent dose (MEE) for each maximum–dose–density (MDD) therapy (memory–free phenotype *A* probability 0.99). 166
5.19 The maximum equivalent dose (MEE) for each maximum–dose–density (MDD) therapy (memory–free phenotype *A* probability 0.5) 167
5.20 The fitness effects of varying σ are dependent on the molecular switch (0.99 memory–free phenotype *A* probability) 169
5.21 The fitness effects of varying σ are dependent on the molecular switch (0.5 memory–free phenotype *A* probability) 170
5.22 Comparions of MDD, adaptive therapy and optimal strategies for restricted doses (memory–free phenotype *A* probability 0.99) 172
5.23 Comparions of MDD, adaptive therapy and optimal strategies for restricted doses (memory–free phenotype *A* probability 0.5) 173

6.1 Growth dynamics of *in silico* bet–hedging populations with carrying capacity 186
6.2 Growth rates and phenotype frequency for varying *c* 188
6.3 Schematic of the mutational spectra where drug resistance is genetically driven 197
6.4 Schematic of the mutational spectra where drug resistance is driven by bet–hedging 198
6.5 Mutational spectra of cytotoxic therapy 200
6.6 Mutational spectra of cytostatic therapy 201

B.1 Genotype-growth rate curves for varying parameter sets with the DC switch . 252
B.2 Genotype-growth rate curves for varying parameter sets with the DCx switch 253
B.3 Genotype-growth rate curves for varying parameter sets with the DCy switch 254
B.4 Genotype-growth rate curves for varying parameter sets with the AM switch . 255
B.5 Genotype-growth rate curves for varying parameter sets with the DC switch under a fluctuating environment 256
B.6 Genotype-growth rate curves for varying parameter sets with the DCx switch under a fluctuating environment 257
B.7 Genotype-growth rate curves for varying parameter sets with the DCy switch under a fluctuating environment 258
B.8 Genotype-growth rate curves for varying parameter sets with the AM switch under a fluctuating environment 259

List of tables

2.1 Summary of abstract models for the genotype–phenotype mapping 47

3.1 15 β–lactam antibiotics for which Mira et al. [228] derived fitness landscapes 60
3.2 Spearman correlation coefficients of fitness values between landscapes 67
3.3 Summary of evolutionary steering to avoid highly resistant genotypes 74
3.4 Summary of evolutionary steering towards low fitness genotypes 75
3.5 Exhaustive summary of all short steering strategies 77

Chapter 1

Introduction

1.1 Drug Resistance and Evolution

"Nothing in biology makes sense except in the light of evolution". This is the perspective taken by Theodosius Dobzhansky in a 1973 essay of the same title [74] as the foundation of an argument to reconcile creationism with Darwin's theory of evolution through natural selection [67]. Whilst Dobzhansky's essay is a theological one, his perspective has in recent years become one of practical *clinical* importance. The rapid emergence of antibiotic resistance, the failure of cancer therapies and the emergence of drug resistant viruses are all driven by the same evolutionary principles. As the pool of effective drugs in each of these cases is limited, and worse still in the case of infectious disease is in fact *shrinking* as multi–drug resistant strains continue to spread, it is becoming ever clearer that the evolution of drug resistance cannot be ignored or overcome through drug discovery programmes alone. Instead, we must understand the underlying evolutionary principles driving resistance and design therapeutic regimes that minimise the risk of drug resistance whilst maximising the likelihood of cure. In this thesis we develop evolutionary mathematical models of drug resistance to explore the potential efficacy of previously proposed treatment regimens. Through these models we will highlight the vital importance of understanding the, at present under–studied, relationship between genetic mutation and phenotypic change, the so–called *genotype–phenotype* map, in predicting the course of evolution.

1.2 Mechanisms of Drug Resistance

1.2.1 Antibiotic Resistance

Novel classes of antibiotics discovered since 1987 are few in number [305], and the discovery within known classes has slowed substantially. Pathogenic bacteria have displayed a staggeringly wide variety of adaptations that render them impervious to the cytotoxic effects of antibiotic treatment [319]. For example, the acquisition of genetic mutations that encode enzymes (e.g β-lactameses) that destroy antimicrobial agents, efflux pumps that remove antibi-

otics from the cell cytoplasm, mutations that alter metabolic pathways and prevent the drugs binding to the cell wall or even the loss of the cell wall entirely (L-form bacteria [87]). In addition to these genetically driven modes of adaptation, bacteria are able to survive drug treatment through non–genetic resistance mechanisms; phenotypic change in response to environmental cues (*phenotypic plasticity* [328]) or exploiting stochasticity to produce a phenotypically heterogeneous population wherein at least some individuals are drug resistant [204].

Such phenotypic heterogeneity has been observed within isogenic populations of a number of organisms and across many biological scales [21], from the unicellular – bacteria [327], fungi [203] or cancer cells [134] – through to insects [66, 151], plants [48] and even aspects of human development [322]. This phenotypic variation has been observed even in homogeneous and constant environments, suggesting that aspects of organismal phenotype may be stochastically determined. In environments that fluctuate unpredictably, this phenomenon can serve as a survival mechanism by increasing the likelihood that at least some individuals are well–adapted to future environments at the expense of fitness in the present one. Thus, stochastic phenotype determination has been termed *bet–hedging* as a species diversifies the phenotypes within the population in order to "hedge its bets" against future environmental change (Seger [292] provides justification for this naming).

An important clinical example of bet–hedging is that of *persister cells* that arise stochastically within isogenic populations of bacteria such as *Escherichia coli* [20, 204, 327]. These cells, which constitute a small fraction of the population ($< 1\%$[204]), have reduced metabolism and shut down non–essential cellular functions. In this dormant state the persister cells are tolerant to the effects of a number of antimicrobials. Although dormant, these cells can retain the ability to proliferate (although at a drastically reduced rate) and when antibiotic treatment ceases persisters will proliferate, producing non–persisters and driving the rapid re–emergence of the bacterial population. Hence bet–hedging, by creating a small sub–population impervious to those therapies that act on proliferating cells, proves to be an effective survival mechanism against antibiotic treatment. Indeed, bacterial persisters are thought to be a contributing factor to multidrug resistance in a number of diseases [204, 178, 237] and are implicated in the dormancy of chronic diseases, such as tuberculosis, which can be suppressed but not eradicated [365].

Drug discovery cannot keep pace with the quickening rate of antibiotic resistance and the pool of effective and potentially life–saving drugs is diminishing. Public health efforts have attempted to stem the emergence of drug resistance by reducing unnecessary prescription of antibiotics [301, 23, 200] and stopping the addition of sub-therapeutic antibiotics in livestock feed [219]. However, such policies require global adoption to be truly effective [229], which they have not yet achieved, and is likely to be infeasible. Thus, it is becoming ever clearer that we must find novel uses for existing antibiotics that allow us to avoid the emergence of drug resistance, through both the genetic and non–genetic mechanisms highlighted here, and to effectively treat resistant disease wherever it is encountered. The identification of these

novel therapeutic strategies is predicated on an understanding of the evolutionary principles underpinning the emergence of resistance.

1.2.2 Cancer Drug Resistance and Targeted Therapy

Hanahan and Weinberg [138] characterise cancer cells as those that undergo unregulated and malignant growth exhibiting six key physiological hallmarks: self-sufficiency in growth signals, insensitivity to growth–inhibitory signals, evasion of apoptosis, limitless replicative potential, and tissue invasion and metastasis. In 2011, these hallmarks were extended further to account for genomic instability, immune evasion, tumour promoting inflammation and deregulation of cellular energetics [139]. Cancer is generally viewed as a genetic disease wherein genetic aberrations induce cellular phenotypes displaying these hallmarks, causing malignant unchecked proliferation and ultimately death. The mutated genes can be characterised as either oncogenes, whose mutation causes a cell to take on a cancerous phenotype, or tumour–suppressor genes whose mutation abrogates evolutionarily acquired anti–tumour defence mechanisms [329]. The apparent genetic causes of cancer have driven researchers to focus their efforts on experiments and data analysis to identify those genes and mutations responsible. For twenty years this gene–centric view has dominated cancer research and molecular biologists have identified myriad genes associated with cancer and have carefully catalogued the complex pathways through which they act.

This molecular reductionist approach has provided major advances in both our understanding and treatment of cancer. Indeed, we are now able to pinpoint specific genes and pathways implicated in a number of cancers and specifically design drugs to target them [284]. These targeted therapies, molecules capable of promoting death or inhibiting growth in cancer cells by interfering with the action of specific molecular pathways, take on a number of forms including molecules that down–regulate the expression of specific genes, block intra–cellular signalling or induce apoptosis by upregulating alternative metabolic pathways. These therapies induce a substantial reduction of tumour burden in a number of cancers; however the response is rarely durable, with a drug resistant cancer often returning within five years. The ‘revolution’ of targeted therapy has followed a pattern that is all too familiar in drug discovery: the identification of a novel mechanism through which drugs can act, followed by the synthesis of such drugs and effective clinical trials, but ultimately failure owing to the evolution of drug resistance. In the past 20 years, this pattern was borne out previously by chemotherapies, then by targeted therapies, and there is evidence that immunotherapies could follow suit [267].

As with antibiotic resistance, there exists a diverse range of mechanisms through which resistance to cancer therapies can emerge. The therapeutic efficacy of chemotherapeutics or targeted therapies are negated by genetic mutations that induce similar resistance mechanisms to those observed in bacteria: efflux pumps, altered binding sites on the cellular membrane, activation of alternative metabolic pathways and resistance or repair of DNA damage etc [121]. Again not all mechanisms of resistance to cancer therapy are genetic; epigenetic or physiological adaptations can also occur. In the context of cancer, bet–hedging–like resistance

mechanisms have been minimally studied; however, a number of aspects of disease course suggest that bet-hedging may be important for understanding how tumours evade therapy. Significant regression of tumours post-therapy leads to a period of remission, followed by the regrowth of aggressive, therapy-resistant lesions. These dynamics are commonly explained by the clonal (i.e. genetic) model of cancer [128], wherein recurring drug–resistant cells are those that have stochastically acquired resistance mechanisms through genetic mutation either during or before therapy. However, the high frequency of tumour recurrence in many cancers suggests that therapeutic escape cannot be solely based on mutational chance. Experimental results have shown evidence of transitory resistance [192, 363] indicative of the existence of a small drug–resistant subpopulation that re–establishes a drug–sensitive cancer cell population. Recent experiments have identified the existence of such populations of 'cancer persister cells' in a cell line of EGFR+ non–small cell lung cancer [296] indicating that bet–hedging may play a role in the emergence of cancer drug resistance [265, 69]. Indeed, non-genetically arising quiescent cancer cells have been suggested as a possible driver of tumour dormancy, the clinical phenomenon wherein residual disease lies dormant for a period before driving recurrence [4]. However, other physiological mechanisms have also been hypothesised as the driver of this latency in recurrence, including immuno-surveillance or a delay in angiogenesis (the recruitment of vasculature). Thus, an understanding of bet-hedging in normal and abnormal (e.g. cancer) cells may help us understand why certain types of therapies fail while others succeed.

An alternative stochastic mechanism that acts to drive cancer progression is that of *stem-like* cancer cells that give rise to a population of heterogeneous phenotypes. In the traditional stem-cell model, a phenotypic hierarchy exists wherein the so–called *cancer stem cells* (CSCs) can produce the range of tumour cell phenotypes observed within a heterogeneous tumour [266]. Under this model, cancer stem cells divide to produce either more cancer stem cells (self–renewal) or cells with non–stem phenotypes further down the hierarchy. This cellular decision is often assumed to be stochastic and is thus an example of bet–hedging. This model of cancer progression, known as the *cancer stem cell hypothesis* [72, 77, 317], has important implications for the emergence of drug resistance as well as of other aspects of cancer progression, including the evolutionary dynamics of the initial cancer growth and dynamics of dissemination and metastasis [281].

A major difficulty in designing effective cancer therapy to mitigate drug resistance lies in the large number of resistance mechanisms that can occur (often simultaneously) and the difficulty in identifying those at work sufficiently early in the disease progression. For example, resistance to tyrosine kinase inhibitors (TKIs) in EGFR driven non–small–cell lung cancer is known to occur through at least twelve different mechanisms [83] and identifying which of these mechanisms is driving recurrence after initial therapy requires expensive and invasive biopsies.

Understanding the evolutionary progression of a disease is just as important in the treatment of cancer as it is in the treatment of bacterial infection. Of course, the diseases differ in that

cancers are not transmissible and in this sense cancers are unique in each instance: the evolution of drug resistance occurs *de novo* in each patient. From this perspective it is clear the treatment of cancer should be driven by a *personalised medicine* paradigm, wherein the patient–specific genetics, disease microenvironment *and* evolutionary history of the tumour are integrated to determine the optimal course of action.

Cancers are known to be genetically and phenotypically heterogeneous [300, 217, 216, 111], potentially harbouring distinct sub–populations that are resistant to different drugs [160]. This heterogeneity can arise as cancers are genomically unstable and because in an expanding population natural selection is not sufficiently strong to drive fitness neutral or weakly maladapted sub-populations to extinction [309, 354]. As such, provided they do not incur a substantial fitness cost, mutations that confer drug resistance need not be lost when they arise in the initial growth of a tumour. When multiple mutations conferring resistance to different drugs arise in the outgrowth of a single tumour, single drugs may be insufficient to entirely eradicate the disease, whereas multiple drugs used together may prove more effective.

A key difference between cancers and infectious disease lies in the availability of drugs. Whereas the discovery of antimicrobials is stalling, we are presently experiencing a golden age of cancer drug discovery driven by targeted therapies and, more recently, immunotherapies. However, targeted therapies almost invariably fail due to the evolution of resistance and combinations of multiple drugs are likely required to mitigate this resistance – a scenario mirroring the need for combination therapies to combat multi–drug resistant bacteria. In fact, despite a large and growing pool of potential therapies, cancer combination therapy is more restricted than the bacterial counterpart, as off–target effects render cancer therapeutics so toxic that they can rarely be prescribed together.

Reducing the toxicity of cancer drugs represents a key first step towards identifying combination therapies that are both effective against heterogeneous tumours and tolerable as a long-term therapeutic option. One approach to reduce toxicity to traditional chemotherapies is *chronotherapy* [201, 231]. Human tissues adhere to a 24 hour circadian rhythm and divide more readily during the day than they do at night. In cancerous cells, this circadian rhythm is often aberrant with mitosis occurring readily throughout a 24 hour period. Chronotherapeutic dosing schedules infuse chemotherapeutic drugs whilst cell divisions in normal tissue are minimised, with the aim to reduce toxicity. A multi-centre clinical trial found that this approach did indeed reduce toxicity whilst maintaining efficacy [202], potentially allowing the use of multiple drugs concurrently. However, it is presently unclear whether such an approach will prove effective in minimising the toxicity of targeted drugs, as unlike chemotherapies, the action of these drugs is not necessarily restricted to cells undergoing DNA synthesis.

An alternative to potentially toxic combination therapies is to prescribe *sequences* of drugs in order. This approach has been suggested as a potential therapeutic strategy to prevent drug resistance and cure disease in both bacterial infections and cancers [325, 270, 53, 228, 122]. Such sequences cannot be tested under a traditional clinical trial paradigm as the number of participants required to perform a trial of all sequences of even a modest sample of drugs is

intractable. Instead, we must rely on mathematical modelling and the phase *i* (for "imaginary") trial [291, 180] as a means to filter out those sequential drug therapies that show the potential to succeed in a traditional clinical trial. The results of these mathematical models have challenged mainstream clinical approaches to treating both cancers and infections and have formed the basis of a novel therapeutic approach — *the adaptive therapy paradigm*.

1.3 The Adaptive Therapy Paradigm

The traditional clinical perspective for treating disease, both infections and cancers, is the *maximum dose–density approach* — that drugs should be administered at high doses for as long as it takes for the therapy to cure disease, to become too toxic to continue, or to cease to be effective. This approach works on the assumption that inducing rapid cell death will cure the disease as quickly as possible but also minimise the number of cell divisions and thus the potential for resistance conferring mutations to arise [97, 294]. However, the argument that minimising cell divisions reduces the likelihood of drug resistance is only valid when that resistance occurs *de novo* during the course of treatment. When resistant phenotypes pre–exist within a heterogeneous population [112, 119], or when they emerge from environmentally–driven physiological adaptations, this approach is not necessarily effective. Indeed, the failure of the maximum dose–density approach in cancer therapy can occur because of the ecological principle of *competitive release* [2, 9]. Before treatment, cancer cells compete with one another for limited resources within a spatially constrained population. In this environment, pre–existing resistant cells are often less fit than sensitive ones and thus, owing to competition, do not grow to large numbers within the population. Selective pressures within the tumour are often not sufficiently strong for clonal sweeps (the fixation of a single genetic clone) to occur and for lower fitness phenotypes to become extinct [309, 271]. As a result, tumours are both genetically and phenotypically heterogeneous [104, 111, 364]. When this heterogeneous population is exposed to a maximum dose–density therapy the sensitive cells are killed, allowing the rapid outgrowth of the previously small and resistant population. This population then drives the recurrence of drug–resistant disease.

The pattern of fitness trade–offs associated with drug resistance mechanisms is a common one that occurs in both cancers and infectious disease [101, 302]. Exploiting these evolutionary trade-offs has formed the basis of an alternative paradigm called *adaptive* or *evolutionarily–informed* therapy, wherein multiple drugs are used, in combination, in sequence, with variable doses, or with breaks in therapy, to cure disease or extend life by accounting for the evolutionary drivers of drug resistance [100, 304]. This clinical paradigm has arisen as a potential treatment for both cancers [103, 102] and infectious disease [234, 228]. Recently, mouse model experiments of an adaptive therapy applied to breast [84] and ovarian [103] cancers, wherein the drug dose is lowered over the course of treatment, were able to stabilise tumour growth where the maximum dose–density approach resulted in death. Following the success of these preclinical experiments, the efficacy of adaptive therapy in treating prostate

cancer is currently being assessed in a clinical trial carried out at the Moffitt Cancer Center (clinicaltrials.gov ID: NCT02415621).

The difficulty of designing effective adaptive therapies, as highlighted by our brief overview of the myriad potential drivers of drug resistance, is that *de novo* evolution of drug resistance must be pre–empted whilst any pre–existing heterogeneity within the disease must be measured and accounted for. Population level statistics, for example those provided by traditional clinical trials, provide information regarding *average case* evolutionary behaviours for resistance to many therapeutics. For example, clinical trials indicate that patients with stage III/IV EGFR+ non–small cell lung cancer (NSCLC) respond to the tyrosine kinase inhibitor (TKI) erlotinib and that the median progression free survival (PFS) time is 9.7 months [277]. Such information allows us to design crude adaptive therapies by identifying average case estimates of the optimal time to perform imaging or biopsies or to change the treatment schedule. However, less is known about what patient–specific factors differentiate between those patients who respond badly (e.g. the 60% of NSCLC patients with PFS less than one year under erlotinib) from those who respond well (e.g. the 11% whose PFS extends beyond two years). To design true adaptive therapies we must be able to differentiate between patients who will respond well and those who will relapse, predict the causes of that relapse and determine effective follow up therapies.

It follows that designing adaptive therapies relies on the integration of multiple disciplines. Clinical expertise is required to identify and administer viable therapies and to monitor patient response. Molecular biologists and bioinformaticians must identify and quantify genetic heterogeneity within a disease and identify the genetic drivers of resistance. Imaging and immuno–histo–chemistry analysis must be performed to monitor disease progression and identify heterogeneity in phenotypes. The evolution of drug resistance is governed by complex interactions that bridge between these biological scales and disciplines. As these interactions are inherently non–linear, including regulatory feedback across biological scales, it is entirely infeasible to infer the consequence of genetic mutation, environmental change or therapy intuitively from data or observation alone [13]. It is for that reason that mathematics plays a central role in the design of adaptive therapies by allowing us, through close collaboration with clinicians and biologists, to identify the effects of perturbations to the complex dynamic system underpinning disease progression. This perspective, that mathematical models should be built in constant dialogue with biologists to aid in the design of adaptive therapies, is an emerging paradigm in oncology and has been coined *Integrative Mathematical Oncology* (IMO) [13]. The principles of the IMO approach, although originating from oncology, have just as much to offer communities studying infectious disease, where mathematical models have long served as a valuable tool. It is the IMO paradigm that we adhere to within this thesis.

1.3.1 Predicting Evolution and the Genotype–Phenotype Mapping

Evolution is not solely genetic mutation and nor is it simply natural selection but rather the directed adaptation of an organism (or population) to an environment that occurs from the

interaction of these two natural processes. Hence, evolution is inherently two–tiered, with genetic mutations changing phenotype and environmental selection of phenotype determining the survival or extinction of novel genotypes. If we are to predict how evolution will drive adaptation, as we must if we are to design effective adaptive therapies to mitigate drug resistance, then we must understand how these two biological scales interact. Specifically, we must understand how genetic mutations will manifest themselves as phenotypic change – the *genotype–phenotype (GP) map*. Understanding the relationship between genotype and phenotype has formed the basis of the modern evolutionary synthesis that reconciles Mendel's work on inheritance [224] with Darwin's theory of evolution by natural selection [67]. However, despite this role of the GP–mapping at the centre of our understanding of evolution and, at least from Dobzhansky's perspective, of evolution at the centre of our understanding of biology, the GP–mapping remains poorly understood and severely understudied.

The central dogma of molecular biology is that genes are transcribed to mRNA that is translated to proteins whose expression levels can be taken as a proxy for organismal phenotype. This perspective has driven a reductionist approach to biology that aims to catalogue the genes responsible for different biological phenomena. The reductionist paradigm ignores a number of important mechanisms implicated in the determination of phenotypes: epigenetics, development, phenotypic plasticity and stochasticity. This approach is also predicated on a gross oversimplification of the GP–mapping that assumes a one–to–one correspondence between genotypes and phenotypes that, owing to stochastic and environmental influences on organismal phenotype, is not justified.

Pigliucci [258] has argued, from the perspective of reconciling the modern evolutionary synthesis with developmental biology, that we must move away from the simplistic view of genes as "blueprints" for phenotypes towards a more holistic understanding of the GP–mapping. This perspective is equally valid in the design of adaptive therapies, as the GP–mapping is central to predicting the course of evolution. Pigliucci highlights that environment is as much of a determinant of the phenotype of an organism as its genotype; the non–genetic drivers of drug resistance that we have highlighted here provide evidence of this fact. In this thesis we present models of evolution and drug resistance from the unique perspective of the GP–mapping and argue that, if nothing in biology makes sense except in the light of evolution then, crucially, nothing in evolution makes sense except in the light of the genotype–phenotype mapping.

1.4 Thesis Overview

In this thesis we explore the potential efficacy of theoretical adaptive therapies for treating drug–resistant disease; both bacterial infections and cancers. We approach this problem through the construction of mathematical and computational models with a unique perspective that keeps the GP–mapping at the forefront of our approach. Justification for this approach, coupled with an extensive literature review of previous modelling, is presented in Chapter

2. In Chapter 3, we introduce the simplest possible model of the GP–mapping, a direct correspondence between genotype and fitness, as a means to formalise the intuitive concept that it is possible to *steer*, in an evolutionary sense, progression of disease to a state that is both treatable and from which drug resistance cannot emerge. Chapter 4 introduces a model of the GP–mapping in which phenotypes are determined from the stochastic simulation of molecular interaction networks. Through this model we are able explore the evolution of non–genetic, stochastically–driven resistance mechanisms such as bet–hedging, that have been largely overlooked by the molecular reductionism of modern cancer research. In Chapter 5, we present an extension of this model to account for epigenetic cell memory and highlight that the fitness contribution of memory mechanisms is dependent on the GP–mapping. In Chapter 6, we explore how properties of the GP–mapping are manifested at the population scale and suggest experimental approaches to identify non–genetic mechanisms of phenotypic heterogeneity in model biological systems. Finally, in Chapter 7, we present the potential extensions of our work, the most immediate steps towards experimental validation and the experimental and modelling extensions we are presently performing.

1.5 Scientific Contributions and Output

We conclude this chapter by first providing a brief summary of the scientific contributions that we believe are made by the work presented in this thesis. Finally, we provide an overview of any papers, preprints or manuscripts in preparation that are derived from the work in this thesis. Much of our work can be extended or applied to different biological questions, and for this reason, it is our hope that the following list will be further extended in future.

1.5.1 Scientific Contributions

We believe that the key scientific contributions of the work presented here are as follows

1. We have identified the potential for stochasticity in evolution that undermines the use of low-replicate-number experimental evolution to predict collateral sensitivity (See Chapter 3). Low-replicate-number experimental evolution has been suggested previously as a translational method to determine drug sequences for clinical use [161] and our work is the first to explicitly outline the pitfalls of this approach. This finding has changed the methodology for studying collateral sensitivity in my own future lab (Institute of Cancer Research, London) as well as those of our collaborators (Dr. Jacob Scott at the Cleveland Clinic and Dr. Robert Bonomo, Case Western Reserve).

2. We have highlighted the importance of drug ordering and, in particular, the potential for arbitrary drug ordering to promote drug resistance. Our work is not the first to do this, however, the quantification of the likelihoods of drug resistance has provided a means

to better communicate the importance of drug ordering with biological and clinical collaborators. As a result of our work, experimental collaborators at Case Western Reserve University (Dr. Robert Bonomo) and the Institute of Cancer Research (Dr. Andrea Sottoriva) are now studying sequential drug therapies.

3. We have explored how properties of the genotype–phenotype mapping can manifest as measurable population–scale properties. Through these model results we have suggested new experimental tests to help identify bet–hedging. At the Institute of Cancer Research, we will be working to validate the mutational spectrum test for bet-hedging in a model system of colorectal cancer thought to exhibit bet–hedging.

4. In Chapter 4, we have suggested a potential explanation for an apparent paradox of bet–hedging: why does it persist for long periods if it is necessarily deleterious? Through mathematical modelling, we argue that the biological mechanisms driving bet–hedging may slow evolutionary loss and hence preserve bet–hedging.

5. Further, our work on the mechanisms of bet–hedging has outlined the importance of understanding precise mechanistic drivers of bet–hedging if we are to make evolutionary predictions. For example, we find that phenotypic memory, modelled as the epigenetic inheritance of intracellular molecules that bias offspring fitness to that of the parent, can either be selected for or against depending on the switch that drives bet–hedging. Previously, many modelling studies encapsulated bet–hedging strategies as fixed probability values without explicitly considering how these values are biologically determined (as in our work on CRNs, Chapter 4) or whether they can change through time or vary between individuals (as in our work on phenotypic memory, Chapter 5). Through conversations with colleagues in the mathematical modelling community, we expect explicit considerations of the mechanisms driving bet–hedging to become more common place in bet–hedging modelling studies. Certainly, this avenue of mathematical modelling research will be continued by myself, as well as collaborators in the department of Integrated Mathematical Oncology at the Moffitt Cancer Center.

6. Mathematical prediction of optimal therapeutic dosing schedules is presently a broad field of which the design of adaptive therapy, and its subsequent evaluation in early stage clinical trials, is a notable success. If adaptive therapy becomes common clinical practice, then the robustness of this approach will need to be verified under a range of ecological disease conditions. Our work exploring the efficacy of adaptive therapy in bet–hedging populations is an example first step towards such verification, and to our knowledge the first study to consider bet–hedging in the context of adaptive therapy. Our colleagues at the Moffitt Cancer Center will continue to explore the efficacy of adaptive therapy for different disease conditions through a combination of mathematical modelling and biological experiments. We expect their work will build on the results presented in Chapter 5.

.

1.5.2 Papers, Preprints and Manuscripts in Preparation

A key contribution of this work is the introduction of a modelling paradigm that emphasises abstraction over reductionism. Specifically, we focus on abstract models of the genotype–phenotype mapping to help elucidate its role in the evolution of drug resistance. A brief summary of this modelling approach, as applied to the study of heterogeneity in cancer, will appear in an invited review for Cell Trends co–authored by Dr Alexander Anderson, Dr Mark Robertson–Tessi, Ryan Schenck, Casey Adam and myself. An extended review paper comparing abstract models of the GP–mapping and following the style of Chapter 2 is also in preparation.

The work of Chapter 3 exploring fitness landscapes, evolutionary steering and the therapeutic potential of drug sequences was the first to be completed. As such, this work has had the most time to mature. The model of bacterial evolution on fitness landscapes highlighting the potential for evolutionary steering was published in PLoS Computational Biology in September 2015 [234] and has received 28 citations (Google Scholar count) in the 18 months following publication. Although not the first to highlight the potential for sequences of drugs to overcome drug resistance, this paper was, to our knowledge, the first to highlight the correspondence between matrix non–commutativity and the importance of drug ordering. This observation, stemming from a simple and accessible mathematical model, has allowed us to communicate effectively with clinical and experimental scientists with whom we are now collaborating to test the efficacy of sequential drug therapies. The work regarding the non–repeatability of evolution and its effect on the experimental derivation of collaterally sensitive drugs was not included in our PLoS Computational Biology paper. Instead, this work, coupled with *in vitro* evolution experiments performed by collaborators in the lab of Dr Bonomo at Case Western Reserve University that verify the existence of non–repeatable evolution, is presently in preparation for submission to a biologically– or clinically–focused journal where it will have the largest impact. A preprint of this manuscript has been submitted to the bioRxiv [236].

Inspired by our work studying drug sequencing and collateral sensitivity in bacteria, similar evolutionary experiments were performed in ALK-mutant lung cancer by experimentalists at the Moffitt Cancer Centre (FL) and the Cleveland Clinic (OH). This work highlighted phenotypic plasticity in drug resistance that drives variable collateral response over a brief time period (1-21days) following the cessation of therapy. This work was published in Nature Scientific Reports in April 2017 [73] where I appear as second author for contributions to the design of the experiments and the production of the manuscript.

The work of Chapter 4 relating the structure of the genotype–phenotype mapping to the long term evolutionary dynamics of bet–hedging was published in the journal Genetics, in December 2016 [235], where it appeared as an issue highlight. To our knowledge, this work is the first to explore minimal mechanistic models of bistable switching, as opposed to

models derived from molecular pathways, in the context of the genotype–phenotype mapping. Through this minimal modelling approach we were able to demonstrate the central role that the switch structure plays in the evolution of bet–hedging, highlighting that small errors or uncertainty in the derivation of interaction networks could result in grossly incorrect evolutionary predictions. Further, by drawing on established chemical reaction network switching models, this work permits theory developed to describe network dynamics to be applied to the genotype–phenotype mapping and studies of evolution. The work of Chapter 5, introducing molecular inheritance and phenotypic memory, was undertaken as a direct extension of the published work of Chapter 4, with the aim to weaken the assumptions underlying our model. It is the results of Chapter 5, and in particular the observation of genotype–phenotype mapping dependent selection for or against phenotype memory, that will form the basis of the next manuscript we will produce from the work of this thesis.

Finally, in Chapter 6 we present two examples of the power of abstract modelling of the genotype–phenotype mapping: highlighting its role in the production of meta–phenotypic patterns such as phenotype frequency distributions and mutational spectra following therapy. Our aim is to verify our predictions from Chapter 6 through empirical experiments. Upon completion of this DPhil, I have been offered a position as a post–doctoral training fellow at the Institute of Cancer Research (ICR) in London where I will continue to study the GP–mapping in the multidisciplinary Centre for Evolution and Cancer. There, I plan to design model–driven experiments to verify the results of Chapter 6 and to further explore the role of non–genetic heterogeneity and the genotype–phenotype mapping in cancer.

Chapter 2

Genotype–Phenotype Mapping

To predict how (or indeed if) a population can adapt to a given environmental change it is necessary to understand how genetic alterations arise, how these alterations manifest themselves as phenotypic change and how viable the resulting phenotypes will be in the context of a specific environment. Johannsen [166] was the first to explicitly recognise a distinction between the mechanisms of inheritance, that is the *genes*, and the physical characteristics, or *phenotype*, of an organism. This distinction revealed that the relationship between the genotype and phenotype is anything but simple. Although a causal link between heritable “factors" and organismal characteristics had been established by Mendel [224], Johannsen observed that this relationship is not a simple one–to–one mapping. The complexity of this relationship, later named the *genotype–phenotype map* [5], was attributed to complex gene–gene interactions — epistasis [256] and dominance [224].

Watson and Crick’s [339] identification of the physical mechanisms of inheritance, coupled with the later international effort to sequence the human genome in full, provided some insight into the how genetic factors influence phenotype but failed to elucidate much of the process. The mechanisms of genetic transcription and translation have since been well categorised, as have in part, the complex cascades of molecular interactions that form cell signalling pathways. Indeed, understanding of these pathways has formed the basis of the targeted therapy revolution in cancer therapy [284]. However, a full mechanistic description of the genotype–phenotype map remains elusive, owing in part to its highly complex and inter–connected nature. Genes are expressed to create proteins which, together with environmental factors, interact to form complex signalling pathways and regulatory feedback systems that in turn alter gene expression levels to form a complex dynamic system. It is understanding the interactions within this system that form the basis of the emerging data–driven ‘-omics’ fields (metabolomics, transcriptomics, etc) [252, 168] that seek to determine how organismal phenotypes emerge.

The genotype–phenotype mapping is the subject of studies across multiple biological and analytic disciplines with much of the work falling into three (not necessarily distinct) categories – statistical analysis, mechanistic modelling and abstraction. Quantitative genetics and associated fields attempt to identify genes associated with categorical phenotypes or to

utilise statistical techniques to associate collections of genes, or specific loci of the DNA known as *Quantitative Trait Loci* (QTLs), with measurable phenotypic variation (i.e. quantitative traits). These correlations provide insight into the genotype–phenotype mapping as a “black box" process, by associating the inputs (genotypes) with the outputs (phenotypes) and, from a clinical perspective, prove valuable in identifying genetic drivers for a number of diseases [54]. However, statistical correlations provide little insight into the mechanisms through which genetic alterations are manifested as pathological phenotypic change.

The key to predicting the evolution of drug resistance, and to identifying therapeutic approaches that prevent it, is a mechanistic understanding of the genotype–phenotype mapping. For much of the 20th century the mechanism of the genotype–phenotype mapping was probed through reductionist techniques. Under the reductionist paradigm biological systems are organised into a hierarchy, with each layer comprising a basic biology entity, of which many are taken together to form a new entity at a higher level (Figure 2.1 shows an example of this hierarchy for cancer). Thus, one can in principle study biology at any scale by first understanding the individual atoms and exploring how they interact to first form biological molecules such as DNA, then sub–cellular structures and, moving up the hierarchy once more, the living cell. This process can be extended to study multicellularity, whole organisms, communities and eco–systems as the emergent behaviours of the sum of multiple parts. Through this methodology the genotype–phenotype map is studied by direct construction from the basic building blocks of biology, the genes, the DNA or the constituent nucleotides. This approach has successfully explained the structure of DNA, the nature of inheritance and mutation and the process of transcription and translation via RNA which forms the first step in a complex process from which phenotypes emerge.

Where reductionism fails is in bridging between levels of the biological hierarchy, and in particular, in failing to account for feedback mechanisms that act across different scales. Implicit in the assumption of a hierarchy is a directionality in which smaller components combine and “feed–forward" to build a whole. The most striking manifestation of this assumption can be seen in diagrams of cell–signalling pathways (for example as provided in [343]) where feedback mechanisms are rarely shown. In reality feedback is universal across all scales of biological systems. With regard to the genotype–phenotype mapping, we know that gene expression is regulated through both positive and negative feedback mechanisms [297, 147] by microenvironmental factors [220] or, in homeostatic systems, through inter–cellular signalling [98, 25]. It follows that gene expression, and in turn cellular phenotype, is modulated by feedback mechanisms that bridge *downwards* across levels of the hierarchy.

In response to this failing, systems biology emerged as an alternative paradigm wherein the focus is shifted from identifying biological components towards understanding their complex non–linear interactions. By combining system–wide quantitative data (genomics, proteomics, etc) with experimental and theoretical identification of intra–cellular interactions, systems biology attempts to build mathematical and computational models that provide quantitative

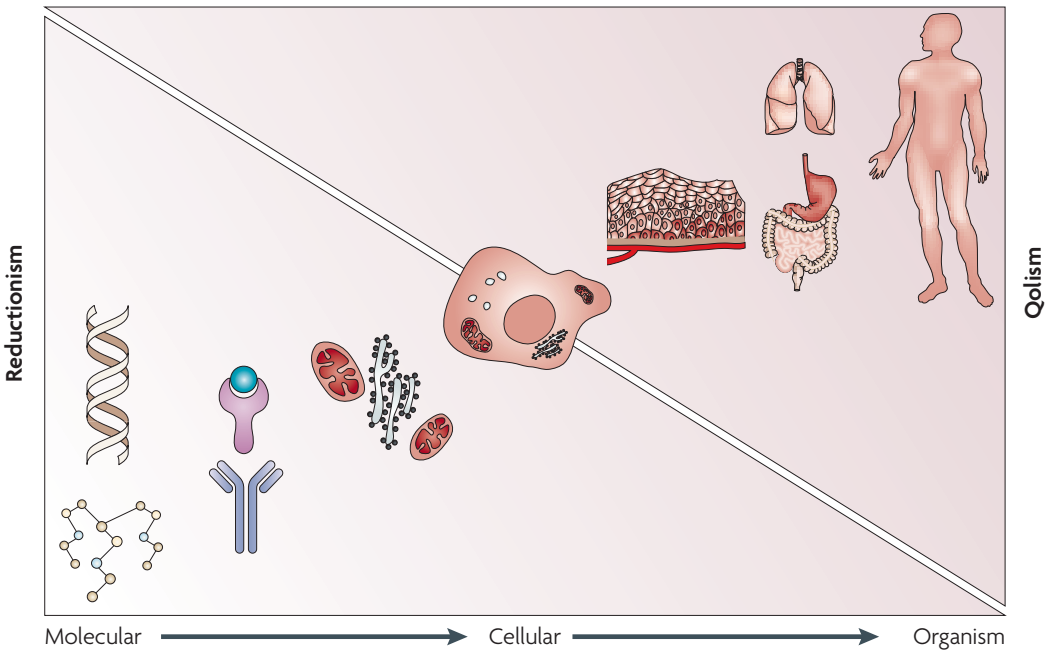

Fig. 2.1 **Cancer and the biological hierarchy.** Genetic alterations induce modified intra–cellular signalling and drive the emegence of cancerous cellular phenotypes. The cells aggregrate to form cancerous tissues (tumours) that eventually disseminate through the body. A reductionist approach suggests that this complex system can be understood by considering the basis (genetic) units. This approach fails owing to feedback mechanisms that bridge downward in the hierarchy. Evolution is an example of a such a mechanism as selection at the cellular, tissue or organ level determines which altered genotypes survive. Figure reproduced from Anderson and Quaranta [13] with permission.

and holistic biological predictions [283]. This paradigm has driven a revolution in modelling ecology and population dynamics and, since the early 2000s, is gaining traction as a tool to study molecular biology and biochemistry [324]. A notable success of this paradigm is the continuing improvement in cardiac modelling [238, 274]. At present the systems approach is utilized to model cancer–associated cellular signalling pathways to identify potential targets for novel therapies [189, 153, 150].

Both the reductionist and systems approaches struggle to account for the overwhelming complexity of biological systems. This complexity cannot be overcome through the collection of ever more data, and at present, proves an intractable problem for explicit computational simulation. Instead, we must rely on abstraction, the central tenet of computer science, to understand the genotype–phenotype map. Through abstraction we are able to strip away unnecessary complexity to reveal a tractable model of the underlying system. For example, the abstraction of nucleotides to a four character alphabet (ATGC) reduces the DNA to a string and permits an algorithmic treatment that forms the basis of phylogenomics. In systems biology intra–cellular molecular interactions are modelled as abstracted networks and analysed through graph–theoretic techniques [253, 218]. However, abstraction in this paradigm is seen

as a “necessary evil" and attempts are often made to maintain as much of the underlying complexity as possible.

An alternative abstract perspective is to instead study simple, tractable analogues to biological systems as a means to gain insight into properties of more complex systems [303, Preface]. From this perspective, biological systems are distilled to their base functional components (as opposed to the physical components of a reductionist perspective) and studied through the *algorithmic lens*. Understanding how these components interact serves as a conceptual tool in generating biological hypotheses and in guiding data collection, experimental design and further mathematical modelling. This paradigm has a rich history dating back to the work of Sewall Wright and Ronald Fisher who derived abstract models to explain aspects of evolution as early as the 1920s [89, 359]. It is this abstract and minimal modelling perspective that we use to explore the genotype–phenotype mapping and its role in predicting the evolution of drug resistance.

2.1 Evolution and the GP–Map through the Algorithmic Lens

To study genotype–phenotype mappings we begin by abstracting away the biological interpretation of evolution to reveal the underlying algorithmic structure. Darwinian evolution is not restricted solely to biological systems but occurs in any system exhibiting four key properties

1. Variation in the characteristics that individuals within a population can possess.
2. Competition amongst individuals causing those with greater fitness, as determined by their characteristics, to be more likely to survive and proliferate (i.e. *natural selection*)
3. The ability for individuals to reproduce or be replicated.
4. A means of heredity within this reproduction that can introduce variability within the population.

The individuals who comprise such a population need not be living in the usual sense[1]. In fact, the individuals comprising these populations need not physically exist at all; this is Dawkins' philosophy of Universal Darwinism [29]. For example, linguistics, economics and memetics have all been studied from an evolutionary perspective [59, 18] and Darwinian principles have been applied to solve optimisation and search problems in artificial intelligence [158]. We aim to study simple instances of the evolutionary process as a means to inform our knowledge of genuine biology.

[1]There is really no good definition of ‘living’ [288, 278] although John Maynard Smith argued that the capacity for evolution is as good a definition of ‘alive’ as any other [307].

2.1.1 The Algorithmic Structure of Evolution

To consider evolution through the algorithmic lens we first define a set of genotypes G as a set of strings over some alphabet of "alleles" Σ (i.e. $G \subseteq \Sigma^*$). Note here that the "alleles" in Σ need not correspond to the biological definition of alleles but rather some basic heritable unit from which genotypes are built. Such a definition permits epigenetic inheritance. This abstraction allows for study of evolving systems at different scales. For example, DNA can be studied from the perspective of base pairs, codons or (biological) genes. In the case of the human genome we could model the genes from the molecular perspective by taking $\Sigma = \{A,T,G,C\}$ and $N \approx 3 \times 10^9$ or we could take a simpler Mendelian approach by setting $\Sigma = \{0,1\}$ to indicate the presence or absence of specific factors or mutations of interest.

We endow the set of genotypes G with a probabilistic mutation mapping

$$\mu : G \times G \longrightarrow [0,1] \text{ satisfying } \forall g \sum_{h \in G} \mu(g,h) = 1 \tag{2.1}$$

which specifies the probability of a genotype $g \in G$ mutating to a genotype $h \in G$ during the replication of an individual. Here μ implicitly defines a mutation rate for each possible mutation. Where the set of possible genotypes is explicitly known, and all mutations are equally likely, we can specify the full mutation relation by a single value which we call *the* mutation rate. This value is often also represented by the symbol μ in a (slight) abuse of notation. Note that the mutation mapping assumes that reproduction is asexual and ignores the possibility of recombination associated with sexual reproduction or lateral gene transfer. An extension of the algorithmic structure presented here to account for recombination is possible, but beyond the scope of this thesis, as we choose to focus on drug resistance driven by asexually reproducing organisms/cells.

As we have not specified a given length for the genotype strings these mutations could be point mutations, deletions, insertions or larger structural mutations such as chromosomal duplications. Mutation rates are often altered by environmental factors such as radiation or the presence of mutagenic chemicals. In this sense μ has some dependence on the environment. For simplicity we assume that this specific environmental dependency remains fixed in all modelling and hence need not be accounted for.

The combination $\mathscr{G} = (G, \mu)$ is a *genotype space*[311]. Stadler [311] showed that, regardless of the choice of the mutation mapping, this definition of genotype space corresponds to a finite pretopology, or equivalently, to a directed graph in which the vertices represent genotypes and edges link genotypes to their mutational neighbours.

In contrast to the structured nature of genotype spaces, it is difficult to give an abstract definition for a space of phenotypes, as the definition of phenotype varies amongst biological disciplines. At one extreme a molecular biologist may define the phenotype of an organism simply as the stable profile of gene expression, whilst at the other, the phenotype can be defined as the collection of all behaviours and effects an organism has on its environment (this is the so–called "extended phenotype" of Dawkins [68]). Here we choose to define a

phenotype space, $\mathscr{P}$, as the collection of potential traits of interest. Phenotypes are thus the product of a number of characters that may be discrete or continuous.

Phenotypes are determined not solely from the genotype but are also modulated by environmental factors [257]. We represent an environment $\mathbf{e} = (e_1, \ldots, e_k)$ ($k \in \mathbb{N}$) as a list of values for environmental variables of interest. These values are drawn from domains $E_1, \ldots, E_k$ that may be continuous or discrete. As an example, in studying drug resistance, the e_i could be binary–valued to indicate presence or absence of a drug, or in a more complicated pharmacological model, be real values denoting the concentration of each drug. The space of possible environments is denoted $E \subseteq E_1 \times \cdots \times E_k$.

Finally, we define a *genotype–phenotype map* (GP–map) as a relation

$$\mathscr{R}_{GP} \subseteq (G \times E) \times \mathscr{P}$$

which associates a given genotype g to a set of possible phenotypes $P_g \subseteq \mathscr{P}$ modulated by an environment $\mathbf{e}$. We have taken the mapping to have an explicit dependence on the environment in order to account for phenotypic plasticity. Further, we emphasize that the mapping is relational and not functional as multiple phenotypes can arise from a given genotype, either non–deterministically or probabilistically.

It should be noted that this definition of the GP–map is a simplification of biological reality as there is no dependence on time. Thus, phenotypes are modelled to be determined instantaneously at birth. This assumption completely removes the process of biological development from our evolutionary modelling. Development is an important area of active research with a strong dependence on the GP–map [258] but is beyond the scope of this thesis.

2.1.2 Fitness and Population Dynamics

Evolutionary models require a set of (usually stochastic) rules governing the processes of death and replication/birth within a population. These rules are usually determined by *fitness* values that encode the survivability or fecundity of an individual with a specific phenotype in a given environment. The concept of fitness is an abstraction used to aid in understanding the process of adaptation through mutation and selection, however, there is no singular definition of fitness: multiple subtly different definitions exist across biological and mathematical disciplines [249].

We define fitness via a function f that assigns to each pair of phenotype and environment a real value

$$f : \mathscr{P} \times E \longrightarrow \mathbb{R}. \tag{2.2}$$

which encapsulates the factors of an individual's viability, fecundity and other factors governing reproductive success into a single value. Where a genotype–phenotype mapping is known, we may extend this definition to genotypes as the expected fitness of an individual

$$f(g, \mathbf{e}) = \sum_{p \in \mathscr{R}_{GP}(g, \mathbf{e})} \mathbb{P}(p | \mathscr{R}_{GP}, g, \mathbf{e}) f(p, e). \tag{2.3}$$

For a fixed environment, this mapping defines a *fitness landscape*. The combination of a genotype space $\mathscr{G}$, phenotype space $\mathscr{P}$, environment E, GP–map $\mathscr{R}$ and fitness function f gives an *evolutionary system* $\mathfrak{E} = (\mathscr{G}, \mathscr{P}, E, \mathscr{R}, f)$.

Simulation of an evolutionary system proceeds stochastically according to rules governing the population dynamics. A number of population dynamics models exist and which is appropriate depends on the underlying problem being studied. Generations can be considered as non–overlapping (as is appropriate for, say, annual plants) or overlapping (appropriate for animals) and time can be treated as continuous or discrete for simplicity.

Population Size and Mutation Rate

In many models of population dynamics, and in biological reality, the fate of a new genotype (and the associated phenotypes) within a population is not deterministic. Beneficial mutants can be lost through genetic drift or deleterious mutations can fix in small populations [62]. Under certain conditions, simplifying assumptions can be made. Specifically, the fate of a new mutant arising in a population, and the maintenance of genotypic heterogeneity within that population, will be dependent on the relationship between population size, N, mutation rate, u, and the *fitness benefit*, s, defined as the value such that if the resident population has fitness normalised to 1, the mutant has fitness $1+s$. Here, fitness can refer to the fecundity (expected number of offspring) or growth rate, depending on the underlying dynamics. For example, neutral mutations (those which do not affect fitness) fix with probability $\frac{1}{N}$ [141] in fixed size populations that evolve according to a Moran process. Beneficial mutations fix with a higher probability than this and deleterious ones with a lower probability. Provided that cells with genotypes of interest can be engineered, typical values for N and s can be estimated from cell culture experiments, however, estimations of the mutation rate parameter are more difficult owing to confounding factors such as the cell death rate [354].

If one assumes that all mutations are either beneficial or deleterious (selection is *strong*) and that $Nu \ll \frac{1}{\log(Ns/2)}$, known as the *Strong Selection Weak Mutation* (SSWM) assumptions, then the population is almost always isogenic. This is because under this regime each new mutant will either *fix*, replacing the genotype shared by the whole population, or become extinct before another arises. Deleterious mutations fix with such low probability that this may be assumed never to happen. Thus, the population remains isogenic with a genotype that is periodically updated by a neutral or fitter mutational neighbour.

Where the SSWM assumptions do not hold, for example due to high mutation rate, large populations or neutral mutations, the population can be genotypically heterogeneous. This is because mutations, even when beneficial, are unable to fix or become extinct before new mutations arise. Such population dynamics give rise to genotypically heterogeneous populations, for example, those encountered in tumours prior to therapy [309, 354, 104]. The existence of these distinct modes of population dynamics is a classical result of population genetics [141]. A more thorough review of the relationship between the strength of selection,

the rate of mutation and the dynamics of adaptation in asexually reproducing populations is provided by Sniegowski and Gerrish [308].

When considering the evolution of drug resistance, both of the modes of population dynamics presented above have some relevance. For example, prior to treatment the population size, N, is likely be much larger than during treatment and thus, it is more likely that the population is genotypically heterogeneous. This scenario is common in cancers where mutation rates are very high and the selective advantage of beneficial mutants has been observed to not be sufficiently strong to induce *clonal sweeps* and fix in the population [242, 225, 128, 309, 271]. In viruses a similar pattern of large population sizes and high mutation rates is also observed, contributing to genotypically heterogeneous populations [40]. When drugs are administered a strong selective pressure is imposed and a large number of cells die. This both lowers the populations size, N, and increases the strength of selection s significantly and can push the population dynamics towards the monomorphic SSWM type. For this reason, simulations of monomorphic/SSWM type population dynamics have seen use in predicting the *de novo* evolution of resistance during therapy [122].

2.2 Mathematical Models of Evolution and the GP–Map

We now present a brief overview of a number of model genotype–phenotype mapping that have previously been used to study evolution. We present each of these models as an instance of an evolutionary system, $\mathfrak{E}$, as presented above. In doing so, we highlight a common evolutionary structure between models developed for different purposes which are, at first glance, seemingly quite different. We find that some common features of evolution emerge time and time again in these theoretical models, often with some empirical corroboration. We discuss a number of these features that are pertinent to the design of evolutionarily–informed therapies.

We begin by considering theoretical models developed by two of the earliest researchers to apply mathematics in studying evolution and population genetics: Ronald Fisher and Sewall Wright. These early theoretical models, Wright's fitness landscape metaphor and Fisher's geometric model, ignored the complex relationship between genotypes and phenotypes and instead explored evolution either from a purely genomic perspective, assuming a direct correspondence between genotype and fitness, or a purely phenotypic perspective, ignoring the genetic basis of changes to the phenotype. Despite these restrictive assumptions, these models yielded a number of relevant results that remain true in more complex models. A more thorough review of the historical impact of these models on evolutionary thinking was presented by Orr [248].

2.2.1 Fisher's Geometric Model

The geometric model of Fisher [89] forgoes modelling the genotype and focuses on evolution at the phenotypic scale. In this model a phenotype space, in the sense of our evolutionary

system $\mathfrak{E}$, is defined as a product of a number of continuous and independent traits,

$$\mathscr{P} = \mathbb{R}^M \text{ for } M \in \mathbb{N} \tag{2.4}$$

where M can be considered a measure of organismal complexity. The genotypes G are taken to be equal to the phenotypes of $\mathscr{P}$ and the GP–mapping is taken to be the identity $\mathscr{R}_{\text{GP}}(\mathbf{p}) = \mathbf{p}$. This definition differs from our presentation of an evolutionary system above, as genotypes are not discrete, a difference which introduces a subtle inconsistency into Fisher's model of evolution.

Mutations, as encapsulated by a mutation relation $\mu : \mathscr{P} \times \mathscr{P} \to [0,1]$ of an evolutionary system, are defined by assuming the existence of universally *pleiotropic* phenotype changes (a single mutation can affect all parts of the phenotype). Thus, mutations correspond to vectors $\mathbf{m} \in \mathbb{R}^N$ which change a phenotype $\mathbf{p} \in \mathscr{P}$ (equivalently a genotype $\mathbf{p} \in G$) according to

$$\mathbf{p} \xrightarrow{\mathbf{m}} \mathbf{p} + \mathbf{m}. \tag{2.5}$$

These mutations should be assigned a likelihood by μ. Fisher's work omits any discussion of what this probability distribution should be and focuses on the probability that a mutation of a *given magnitude*, $r = |m|$, confers a fitness advantage. The aim of this focus is to determine the likely magnitude of each successive phenotypic change that results during adaptation. It is here that we meet the inconsistency of Fisher's work, as without an explicit statement of the probability of such changes the question is ill–formed. Indeed, Fisher refers to mutations as *random*, apparently invoking a uniform distribution, however no such distribution can exist on $\mathbb{R}^M$. From this we infer that each mutation of a given magnitude r is taken to be equally likely, but how this probability changes as $r \to \infty$ is not stated. Although this observation does not impact on Fisher's results, it does serve to highlight the power of using the algorithmic lens and the abstracted evolutionary system $\mathfrak{E}$ as a tool to ensure correctness in evolutionary modelling.

The environment in Fisher's model is defined as a single point in phenotype space, $\Theta_E \in \mathscr{P}$, corresponding to the optimally adapted phenotype. Thus $E = \mathscr{P}$ in the evolutionary system. The fitness, f, of an individual with a non–optimal phenotype $\mathbf{p}$ is defined in terms of a function, w, of the Euclidean distance ($z = \|\mathbf{p} - \Theta\|$) from the optimal phenotype

$$f : \mathbf{p} \mapsto w(z). \tag{2.6}$$

The function w determines how fitness decreases as phenotypes move away from the global fitness optimum. In Fisher's original formulation w is taken to be a univariate Gaussian function

$$w(z) = \exp\left(-\frac{z^2}{2}\right). \tag{2.7}$$

A graphical representation of Fisher's model is shown in Figure 2.2.

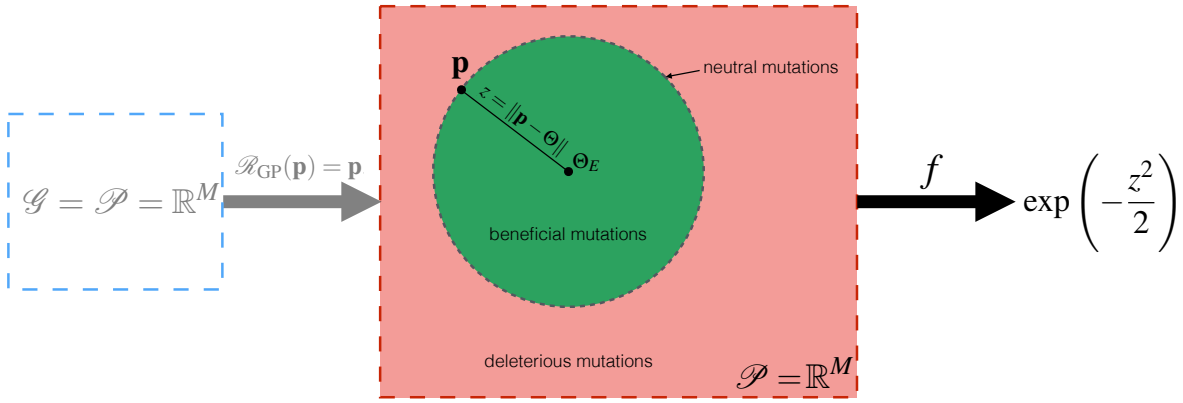

Fig. 2.2 **A graphical representation of Fisher's geometric model.** The genotype space $\mathscr{G}$ and phenotype space $\mathscr{P}$ are both $\mathbb{R}^M$ and the GP–mapping is the identity. Mutations are vectors from $\mathbf{m} \in \mathbb{R}^M$ that are added to a phenotype $\mathbf{p} \in \mathbb{R}^M$. The environment is determined by a globally optimal phenotype $\Theta_E \in \mathbb{R}^M$ and fitness as a function of the distance ($z = \|\mathbf{p} - \Theta_E\|$) of a phenotype from this global optimum. Beneficial mutations are those for which $\mathbf{p} + \mathbf{m}$ lies closer to Θ_E than $\mathbf{p}$ (inside the green circle). Deleterious mutations are those that generate a phenotype outside of this circle (shown in red) and neutral mutations those that generate a phenotype an equal distance from Θ_E (dashed line).

Using this model, Fisher investigated the relationship between the magnitude of a phenotypic change and the likelihood of that change conferring a fitness advantage. Fisher determined that the probability that a mutation of magnitude $|\mathbf{m}| = r$ is beneficial is given by $1 - \Phi(x)$, where Φ denotes the cumulative distribution function of a standard Gaussian distribution and $x = r\sqrt{\frac{N}{2z}}$ is a standardised mutation size. A proof of this result was not provided by Fisher himself but has since been re-derived (for example in Kimura [183]).

Fisher's result highlights three important features of adaptation on the phenotypic scale. Firstly, the probability that a mutation is beneficial approaches 50% from below as the magnitude approaches zero (provided the phenotype is not the optimal and that all mutations of a given magnitude are equally likely). From this result Fisher concluded that the genetic basis for adaptation is the accumulation of a large number of very small phenotypic changes – the *micromutationist* view favoured by Darwin himself [67]. However, as Kimura [183] demonstrated, this conclusion is partially flawed. To be adaptive, a mutation needs not only to be beneficial but must also not become randomly extinct when it does occur. Whilst smaller mutations are more likely to be beneficial, they also confer a smaller fitness advantage when they are beneficial and thus are more likely to be stochastically lost through genetic drift. Kimura derived the distribution of sizes amongst mutations substituted at each step of adaptation, concluding that it is mutations which induce *intermediate*, and not infinitesimal, changes in phenotype that are the most likely drivers of adaptation.

Orr [245, 246] performed a computational and analytic study to determine the distribution of sizes of adaptive mutations that occur successively during adaptation, as opposed to Kimura's distributions of size at individual stages. These results show that the distribution of mutational size is approximately exponential and that adaptation in Fisher's model is

characterised by a relatively small number of large magnitude mutations and many more small ones. This work also shows that the expected size of an adaptive mutation substituted at each step of an adaptive trajectory (first mutation to fix, second mutation to fix etc) diminishes by a constant proportion at each step, forming an approximately geometric sequence. Remarkably, simulations show this result to be robust to assumptions about the precise shape of the fitness function w [245, 246].

Fisher's model also predicts that the likelihood of a mutation being beneficial decreases as M increases, suggesting that adaptation proceeds at a slower rate for more complex organisms. This observation was expanded upon by Orr [247] who described this phenomenon as the *cost of complexity* and estimated its extent, showing that the rate of adaptation declines at least as fast as M^{-1}. This cost is reduced if the phenotype is modular, comprising phenotypic modules

$$\mathbf{p} = (\mathbf{p}_1, \ldots, \mathbf{p}_k) \in \mathbb{R}^{M_1} \times \cdots \times \mathbb{R}^{M_k}$$

of which only one of the $\mathbf{p}_i$ is affected by each mutation. An alternative argument for evolutionary advantage of modularity, and thus that selection may favour modularity within the GP–map, was laid out by Wagner [335, 336]. The geometric model has also been used to study other aspects of evolution including the evolutionary advantages of sex [254], development [269], compensatory mutations [142], mutation–selection–drift balance [143], mutation load [262] and hybridisation [24].

2.2.2 Fitness Landscapes

The fitness (or adaptive) landscape metaphor was first introduced in the 1930s by Wright [359, 360] as a model to account for *epistasis*. Epistasis refers to any genetic interaction in which the phenotypic (or fitness) impact of a mutation at a specific loci is modulated by the genetic background on which it occurs. Wright was particularly interested in the possibility that, owing to epistasis, genotypes may exist in which all single mutations are deleterious but for which a fitter genotype exists. (It is precisely this phenomenon, the existence of *multipeaked* landscapes, that we will exploit to design optimal sequential drug therapies (Chapter 3).)

For a static environment, all GP-mappings induce a fitness landscape over the genotype space according to Equation 2.3. However, we can consider the fitness landscape as the GP–map itself. For an arbitrary genotype space $\mathscr{G}$ (often taken as $\{0,1\}^L$ for some $L \in \mathbb{N}$) we can assign a phenotype to each genotype which is simply the fitness of individuals with that genotype. In the context of the evolutionary framework presented in Section 2.1, we assume that E is fixed (and thus unimportant), that $\mathscr{P} = \mathbb{R}$ and define a functional GP–mapping $\mathscr{R} : \mathscr{G} \to \mathbb{R}$. The fitness function, f, is then simply the identity on $\mathscr{P}$. A representation of this construction, together with the traditional schematic view of a fitness landscape introduced by Wright, is presented in Figure 2.3. Under this construction epistatic interactions are mathematically quantifiable and their effects on landscape topography can be measured.

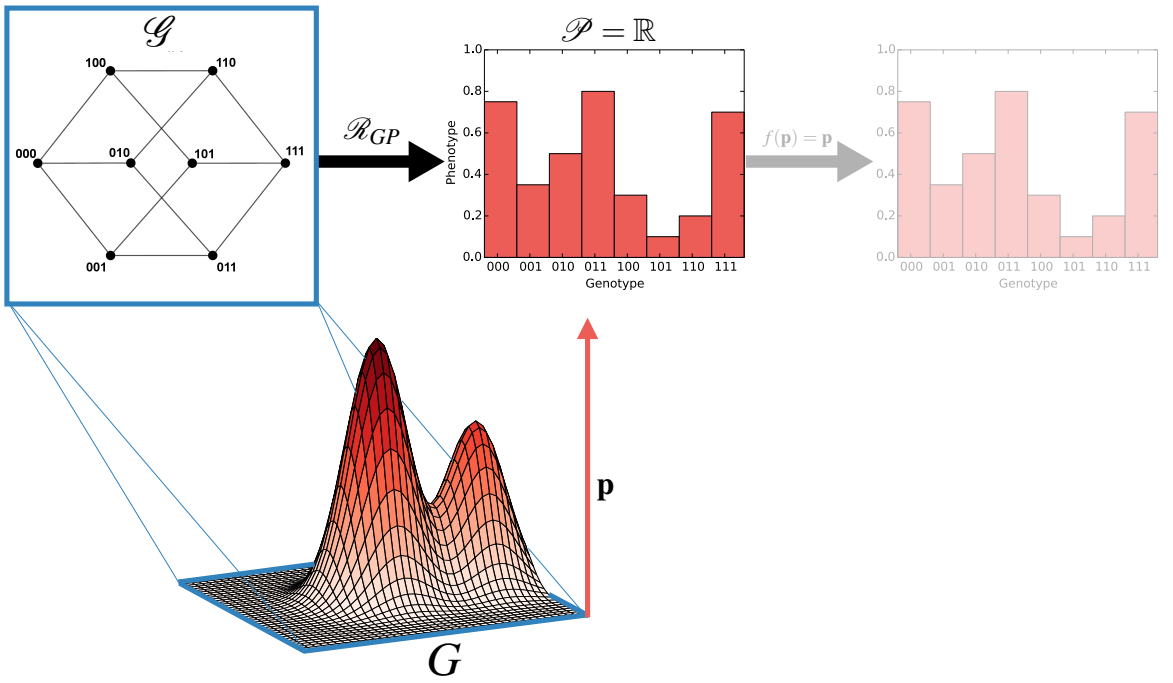

Fig. 2.3 **Fitness landscapes as an evolutionary system.** The genotype space forms a hypercube graph wherein the vertices represent genotypes and the edges mutations. The fitness landscape GP–mapping is an assignment of a fitness to each genotype. Under this mapping the fitness mapping is taken to be the identitity. In the traditional schematic view of the landscape, the hypercubic genotype space is projected into the x-y plane and fitness represented as the height above the plane. This representation emphasises the dynamics of evolution as an 'up–hill' walk.

Coupled with empirically–derived landscapes, this mathematical formulation provides insight into the evolutionary implications of the structure of the GP–map.

The landscape model is particularly useful in studying the accessibility, repeatability and predictability of evolution; three properties of evolution that are particularly important in designing treatment schedules to avoid drug resistance. Maynard Smith [306] introduced the concept of adaptive trajectories in discrete sequence spaces under the assumption that point mutations are sufficiently rare that one can assume no two occur simultaneously. These trajectories can be encoded as a sequence of point substitutions that each increase fitness. Maynard Smith likened these trajectories to a "word ladder" wherein one word must be transformed to another by substituting one character at a time: "COLD $\rightarrow$ CORD $\rightarrow$ CARD $\rightarrow$ WARD $\rightarrow$ WARM". Each intermediate string of characters must also form a valid word. This rule is a weaker analogue of a fitness requirement in evolution: each substitution must form another valid word, just as with evolution, each substitution must increase fitness (or if the population dynamics permit neutral mutations [183], not decrease fitness). When the Strong Selection Weak Mutation criterion for population dynamics is satisfied (Section 2.1.2), these trajectories are an accurate description of how evolution proceeds. A common visualisation of a fitness landscape is to view the $x-y$ plane as a genotype space $\mathscr{G}$ with a surface above that indicates fitness on the z axis (shown in Figure 2.3). Evolutionary trajectories are then

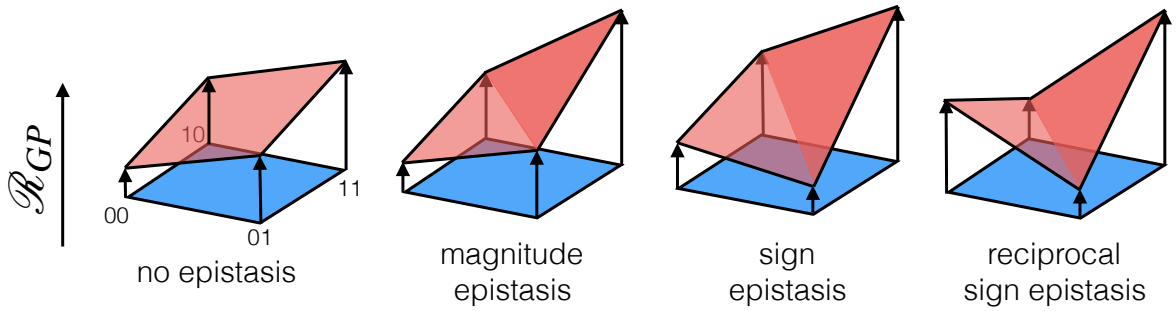

Fig. 2.4 **Forms of pairwise epistasis.** The fitness (height above the plane) effects of different mutations, modelled as flipping a single bit of the genotype, can differ depending on the genetic background. These forms of epistasis can inhibit evolutionary trajectories between two genotypes.

viewed as “uphill" walks on this surface . This metaphor has received some criticism [105] as in reality the genotype space is extremely highly dimensional, a property noted by Wright himself [359].

Theoretical Studies of Fitness Landscapes

Epistasis has a significant impact on how adaptation to a new environment proceeds. Fitness landscapes, by associating a simple phenotype measure (fitness) with each genotype, offer a natural model in which to study the effects of epistatic interactions on landscape topography. The simplest form of epistasis is the pairwise interaction between two genetic loci. Pairwise interactions take many forms (Figure 2.4). Of specific interest is *sign epistasis*, wherein the fitness effect of a pointwise mutation, $a \to A$, changes sign dependent on the allele at a second locus (B or b). It has been shown that sign epistasis can severely restrict the number of accessible evolutionary trajectories between a low fitness genotype and a higher one [349, 260]. This restriction makes evolution somewhat predictable, permitting us to design adaptive therapies that pre–empt the emergence of drug resistance.

A stronger form of epistasis is *reciprocal sign epistasis* in which each of a pair of mutations $a \to A$ and $b \to B$ are individually deleterious but offer a fitness advantage together. This pattern of interaction induces a “fitness valley" between genotypes ab and AB which cannot be crossed in many regimes of population dynamics. This motif is useful in the design of adaptive therapies (for example, in Chapter 3) as by steering evolution to the opposite ‘side’ of a fitness valley from a highly resistant genotype, the uncrossable valley prevents the evolution of drug resistance.

Poelwijk et al. [261] showed that the existence of a pair of mutations exhibiting reciprocal sign epistasis is a necessary condition for a landscape to have multiple optima of fitness. This result highlights a relationship between epistasis and ruggedness in a landscape that is still not fully understood. Many empirically–derived landscapes (see below) are multipeaked suggesting that reciprocal sign epistasis is a common feature of biological genotype–phenotype mappings. Poelwijk et al. [261] also show that there is no sufficient locally identifiable

condition on gene interactions that guarantees a landscape is multipeaked (in the sense that finding a set of genes satisfying that condition guarantees the landscape to be multipeaked).

Epistasis can occur at higher orders than gene pairs, and in general the epistatic interactions amongst n genes ("n^{th} order epistasis") can be considered. Characterising higher order epistasis becomes increasingly difficult as for L loci there are $\binom{L}{k}$ subsets of size k that may interact. Weinreich et al. [348] introduced a mathematical formulation for higher order epistasis. By analysing empirically–derived landscapes using this technique, Weinreich finds that higher order epistasis is prevalent in nature. The importance of this finding in predicting evolutionary trajectories remains an open problem and an active area of research.

The LK–Model

Until recently, empirically deriving fitness landscapes was not possible and experimental studies focused on quantifying the interactions between pairs of genetic loci. As such, theoretical landscape models were required in order to study topography and the effects of epistasis and ruggedness. Kauffman and Weinberger [174] devised the LK model[2] which defines a landscape on L loci in which the degree of 'ruggedness' is tunable by the parameter $0 \leq K < L$. In this model the fitness of a genotype $g = g_1 g_2 \dots g_L$ is determined as the sum of the fitness contribution of each locus

$$F(g) = \sum_{i=1}^{L} F_{g_i}. \tag{2.8}$$

The fitness contribution of each locus g_i is determined by its epistatic interactions with K other loci, $g_1^i, g_2^i \dots g_K^i$. In the original formulation of the LK model these interacting loci are either chosen as adjacent to the locus g_i (with periodic boundaries) or uniformly at random. Thus, the fitness contribution of g_i is determined by a scalar function $F_{g_i} = F(g_i; g_1^i, g_2^i \dots g_K^i) \in \mathbb{R}$. To generate a new fitness landscape, the values of this function are drawn randomly from some population distribution. Following their introduction in the 1980s, and owing to the (until recent) absence of genuine empirically–derived landscapes, LK landscapes were the subject of intense theoretical study to highlight the effects of ruggedness, (as encapsulated by K) on evolutionary trajectories [344, 345, 90, 173, 8].

For $K = 0$, each locus contributes fitness without interaction with others and there can only be two possible values of $F(g_i)$ for each i. It follows that the fitness contributions of each locus are additive and the landscape necessarily has a single globally accessible peak. The expected length of an adaptive walk to this peak is $\frac{L}{2}$ mutations, although there can exist trajectories of length exponential in the number of loci [176]. In such landscapes similar genotypes (those at a short Hamming distance [137]) have highly correlated fitness values. When $K = L - 1$, the fitness function F is equivalent to a random assignment of fitness to each genotype. As such, similar genotypes are uncorrelated. Random landscapes have an expected number of $\frac{2^L}{L+1}$ local

[2] originally presented as the NK but renamed to avoid confusion with N as denoting population size [96]

optima. A summary of theoretical landscape statistics for $K = 0$, $K = L - 1$ and intermediate K, along with a more generalised LK–like landscape construction, is provided by Altenberg [8]. The expected number of local optima, the expected fitness of those optima and the expected length of adaptive walks, all of which have important implications in predicting the evolution of drug resistance, are contingent on the degree of epistasis or ruggedness in the landscape.

Empirical Fitness Landscapes

In the previous ten years, novel experimental techniques have permitted empirical measurements of the genotype to fitness map [71]. Landscapes are measured by engineering strains of an organism with each possible combination of alleles at some genetic loci of interest. For example, in studying antibiotic resistance in gram–negative bacteria these loci may correspond to mutations in the genes coding for beta–lactamese. Where there are L loci of interest a total of 2^L strains must be engineered and these strains are then assayed in a drug–treated environment to determine resistance as a measure of fitness. This fitness measure is commonly determined as the minimum inhibitory concentration (MIC) of drug, or as the average growth rate in clinically–relevant drug concentrations. In Chapter 3, we introduce a mathematical model of evolution to predict evolution on empirically–derived fitness landscapes. Here, we provide a brief overview of empirical studies that derive fitness landscapes.

De Visser and Krug [71] note that by 2014 there had been fewer than 20 empirical studies to derive fitness landscapes and that of these the maximum number of loci considered was $N = 8$. Drug induced fitness landscapes have been derived for *E. coli* [228, 179, 318, 287, 122, 346], *Saccharomyces cerevisiae* [136], *P. falciparum* [209] and type 1 Human Immunodeficiency Virus [64]. A meta–analysis of these landscapes shows that, on average, these landscapes show a substantial level of ruggedness which is suggestive of substantial epigenetic, but not random, gene interactions [316]. For example, of the 15 drug landscapes derived by Mira et al. [228] all but one have multiple local optima of fitness and for 12 of these landscapes there exist evolutionary trajectories from the wild–type genotype to more than one of these local optima. This suggests that evolution is likely not *repeatable*, in the sense that drug resistance can arise through different mechanisms in different cases.

Weinreich et al. [346] derived an empirical landscape for *E. coli* exploring combinations of 5 mutations in the TEM gene (specifically A42G, E104K, M182T, G238S [10] and g4205a [340]) under the antibiotic cefotaxime by using the minimum concentration of drug required to arrest growth (the *minimum inhibitory concentration*, MIC) as a proxy for fitness. Corroborating the theoretical result that sign epistasis restricts evolutionary trajectories, Weinreich finds that only 18 of 120 mutational trajectories from $g = 00000$ to $g = 11111$ are accessible. This work suggests that although rugged fitness landscapes may render evolution non–repeatable, it may be sufficiently constrained to be *predictable*. It is precisely this predictability that forms the basis of evolutionarily–informed sequential therapy.

De Visser and Krug [71] (also Orr [248]) demonstrated that accessible evolutionary trajectories in empirical landscapes are often short, despite the mathematical result showing

that evolutionary trajectories in fitness landscapes can have length exponential in the number of loci [176]. This finding suggests that the genotype–phenotype mapping induces a specific structure of fitness landscapes in which high fitness is accessible in relatively few mutations — a powerful property for rapid evolution of drug resistance. Empirical datasets also show that beneficial mutations induce diminishing increases in fitness as the fitness of the genetic background increases [287, 50, 179], a finding that corroborates the theoretical results from Fisher's model [245, 246] on the expected magnitude of successive adaptive mutations. Thus, the initial mutations in an evolutionary trajectory are, on average, expected to induce the largest increases in fitness. This result was further corroborated by *in vitro* experimental evolution of two bacteriophage viruses under the selective pressure of inhibitory temperatures [149].

From the perspective of the early stages of cancer growth, this result may explain observations (the so-called *big bang*) of colorectal cancer growth, wherein numerous subclones grow together from the early stages. These distinct subclones may have diverged later than the appearance of a tumourogenic phenotype and not differ in fitness sufficiently for clonal sweeps (the extinction of all but a single genotype) to occur. This pattern of diminishing fitness increases also favours the emergence of drug resistance, by allowing the rapid evolutionary escape of a pathogen from the toxic effects of chemotherapeutics.

Rather than focus on the full mapping from genotype to fitness, some empirical studies instead derive empirical landscapes from a small part of the GP–mapping. For example, Aguilar-Rodríguez et al. [3] measured the binding affinity of a transcription factor to all possible short DNA sequences. This study generated over 1,000 empirical landscapes, finding an intermediate degree of ruggedness and epistasis. However, it is not clear whether these binding site affinity landscapes will correspond to the fitness landscapes that arise under different evolutionary circumstances (e.g. the drug–resistance landscapes described above). Further, the landscapes in this work were derived through purely experimental means, without consideration of a model GP–mapping. Other empirical approaches to construct adaptive landscapes from aspects of the GP–mapping include the prediction of RNA secondary structures or protein folding. In these models, which are discussed below, empirical observations are coupled with model GP–mappings derived from physical laws.

Limitations of Fitness Landscape Models

As mentioned above, to empirically derive a fitness landscape in L loci, a total of 2^L strains must be engineered. As such, these experiments become intractable for even modest L. It follows that combinatorically complete empirical landscape studies can only probe a small portion of the fitness landscape. It is possible that this small subspace misses important topographical features of the landscape. For example, whilst one genotype may be separated from another by a fitness valley in the measured landscape, there could exist a ridge of un–modelled mutations that connect the two. To fully understand the landscape topography, as we must to design effective multiple–drug therapies, we must improve data collection and better understand the genotype–phenotype map. This approach has been taken by Hinkley

et al. [146] who used generalised kernel ridge regression (GRR) to estimate fitness landscapes from an incomplete data set of fitnesses for different genetic strains. Their work derives an approximation to the fitness landscapes of HIV-1 under 15 different drugs, for a genotype space totalling 200 genetic loci, from only 70,081 isolates. This landscape was later analysed by Kouyos et al. [188] who identified a high degree of ruggedness as well as large networks of genotypes (neutral spaces) over which fitness varies very little. To our knowledge, no attempt has yet been made to corroborate any small subspace of the statistically derived landscape with the landscape derived from empirically engineering each strain and assaying it for fitness.

Despite the success of recent fitness landscape studies in predicting the evolution of bacteria, and the clear importance of evolutionary principles in driving drug resistance to cancer therapies, no fitness landscape has been derived for mammalian or cancer cells. This lack of empirical landscapes can be explained, in part, by the relative difficulty of *in vitro* experiments utilising cancer cells in comparison to equivalent experiments in bacteria. However, as experimental techniques continue to improve, and as the value of the fitness landscape model and of evolutionary theory in general continues to be demonstrated in oncological research, we can expect that empirical fitness landscapes for cancers will soon be derived.

2.2.3 Secondary Structure Prediction

As empirical studies have shown, ruggedness and epistasis in fitness landscapes are common. However, understanding *where* this epistasis comes from is not possible with the oversimplified fitness landscape model. This is because under the landscape model, one either decides on some arbitrary degree of epistasis/ruggedness (the K of the LK–model) or measures it from empirical datasets; treating the GP–map as a 'black–box' without considering *how* phenotypes and fitness emerge from the genotype. Thus, the mechanistic causes or evolution of epistasis itself, in addition to other properties of the GP–map, cannot be studied through the fitness landscape model as presently formulated.

By contrast, secondary structure prediction, wherein a prediction is made for the structure formed when a single–stranded sequence of nucleotides (or amino acids) 'folds' onto itself, has emerged as a directly–computable model GP-mapping. Owing to its tractable computation, the most common secondary structure model is the prediction of folded RNA secondary structures. This model is appealing for a number of reasons. First, because folded RNA constitutes an important part of the full mapping from genotype to phenotype, for example, mediating translation and transcription of genes, regulating gene expression, and playing a key role in signal recognition at the cell membrane. Secondly, RNA has been suggested as the first error prone, self–replicating (i.e. evolving, Section 2.1) molecule, and thus plays an important role in theories of abiogenesis [55]. Finally, RNA models have been used as a tractable, abstract model of evolution [290, 289]. Here, we present the pertinent results from abstract modelling of RNA folding as a GP–mapping and discuss their relevance both to evolution in general and to evolution of drug resistance.

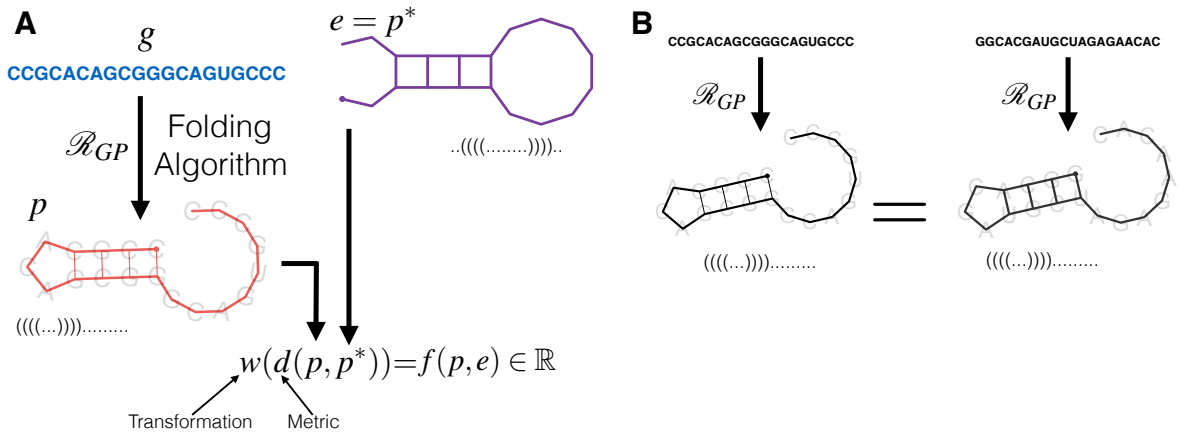

Fig. 2.5 **The RNA folding model**. **A.** RNA folding as an evolutionary system $\mathfrak{E}$. The GP–mapping takes a genotype g to a minimum free energy secondary structure p through an RNA folding algorithm (ViennaRNA [208]). The environment is specified as a target structure $e = p^*$ and fitness is determined as a transformation (w) of the distance (d) between a structure and this target. **B.** Phenotypes/secondary structures are topologically defined, thus two RNA strands can fold into the same target structure.

The RNA Secondary Structure Model

In the context of our evolutionary system, the genotypes of the RNA secondary structure model are strings representing a single linear strand of RNA comprising Adenine (A), Cytosine (C), Guanine (G) and Uracil (U). Thus, the genotypes are given by $G = \{A,U,G,C\}^*$. As with all abstract models, the mutation relation is chosen depending on the problem being studied, however the most common choice is to only permit point mutations — a change of one nucleobase to any of the other three. Thus, μ specifies $3L$ mutational neighbours of an RNA strand $g \in G$ of length L. The genotype space is then $\mathscr{G} = (G,\mu)$. Although the physical properties of the nucleobases result in some substitutions being more likely than others [182], the mutation relation is commonly simplified such that each of these mutations is considered equally likely.

The bases of an RNA molecule pair to form hydrogen bonds (A-U, G-C and, less commonly, G-U) and thus an RNA strand will 'fold' onto itself to form a complex structure [93]. This folding can be conceptualised as having two stages. The base pairs first bond to form (a planar shape known as) a *secondary structure* and then distant parts of this structure come together to form a three–dimensional *tertiary structure*. It is this three–dimensional structure that determines the biological properties of the folded RNA, however it is difficult to accurately predict owing to the complexity in simulating the thermodynamic constraints. Secondary structures are more easily computed and, given knowledge of the secondary structure, many of the biological properties of the full tertiary structure can be approximated. It is for this reason that they are taken as the phenotypes in the RNA model.

An RNA secondary structure is defined as a list of pairs, $p = \{(i_1,j_1),\ldots,(i_k,j_k)\}, (k \in \mathbb{N})$ of positions in the RNA strand g, each with $i < j$ and satisfying that for any two $(i,j),(k,l) \in p$

1. $i = k$ if and only if $j = l$ (a given nucleotide can appear in at most one base pair)
2. $k < j \implies i < k < l < j$ or $k < l < i < j$.

This second condition prevents the existence of *pseudoknots*, a structural feature of folded RNA that makes prediction considerably more difficult. In an evolutionary system, the phenotype space $\mathscr{P}$ is taken as the space of all such secondary structures. An example of this folding is shown in Figure 2.5(A).

RNA secondary structures can be represented in a number of ways including base pair lists as above, outer–planar graphs, as a tree structure, or in "dot–parenthesis" notation as strings over the alphabet $\{(,.,)\}$ where the symbols "(", "." and ")" represent base pair openings, unpaired bases and base pair closings respectively [91]. Note that RNA secondary structures are *topologically* defined and do not depend on the underlying nucleotides of the RNA strand. Thus, different RNA strands can fold to form the same secondary structure (Figure 2.5(B)).

RNA secondary structures can be uniquely decomposed into combinations of *stacks*, double–helical runs of base pairs, and *loops*, sequences of unpaired nucleotides occurring between stacks. Secondary structures can be assigned a free energy value by considering the thermodynamic properties of the base pairs – loops of unpaired bases increase free energy whilst stacks lower it. An RNA strand can fold into many secondary structures, but those with lower free energy are most stable and thus the most likely to be formed when the RNA strand folds.

The simplest instance of a GP–mapping, $\mathscr{R}_{\mathrm{GP}}$, using the RNA folding model is to assign to each RNA strand $g \in G$ the *minimum free energy* (MFE) secondary structure $p \in \mathscr{P}$. As the formation of a stack necessarily creates a loop, the energy trade–off between these two features induces a rugged energy landscapes for secondary structure folds. Due to this rugged energy landscape, determining the MFE secondary structure, and thus mapping genotypes to phenotypes, is a non–trivial task. Indeed, it is not sufficient to simply maximise the number of base–pairings in the secondary structure (an example is provided by Wuchty et al. [361]). The free energy of a secondary structure is determined additively from the contributions of the loops within the structure and hence the MFE secondary structure can be determined through a dynamic programming algorithm [148, 367]. This algorithm provides a tractable means to map genotypes to phenotypes within the RNA model and permits a statistical analysis of their relationship.

The common approach to simulating evolution within the RNA model is to define a fitness function, f, on a secondary structure p, as determined by its "distance" from a given target structure that represents the environment ($e = p^*$),

$$f(p,e) = w(d(p,p^*)) \tag{2.9}$$

where d is a metric determining the similarity between two secondary structures and w is a monotonic transformation of this value. This definition is chosen to represent the concept that RNA molecules with similar fully–folded tertiary structures should have similar mechanistic

action. In the context of an evolutionary system, a target structure p^* plays the role of the environment, **e** (thus, $E = \mathscr{P}$). It is difficult to define biologically–relevant metrics for the comparison of secondary structures and a number of choices have been proposed. For example, measuring the edit distance between the two secondary structures represented in tree form [93] or the Hamming distance between the two secondary structures represented in dot–parenthesis notation [91]. These metrics are somewhat artificial, however some studies suggest that the results of the RNA model are robust to changes in the distance metric defining fitness [93, 11]. The transformation function w also varies between different studies. As RNA structures are often highly conserved between species, it has been suggested that slight deviations from the optimal structure can induce large changes in fitness. For this reason a hyperbolic function is often chosen for w [11, 92, 57].

This definition of environment in terms of optimal target structure ($E = \mathscr{P}$), and fitness, f as some function of the distance (in phenotype space) from this target structure, completes the definition of the RNA secondary structure model as an evolutionary system, $(\mathscr{G}, \mathscr{P}, \mathscr{R}_{\text{GP}}, f, E)$, as presented in Section 2.1. A schematic representation of this evolutionary system is presented in Figure 2.5. Here, we present some of the important results of the RNA model and interpret them in the context of drug resistance. A more comprehensive review of *in silico* studies of RNA folding is provided by Cowperthwaite and Meyers [58].

Neutrality, Robustness and Evolvability

Perhaps the most fundamental finding of the RNA model is that the mapping from genotypes to phenotypes is extremely degenerate. As an example, consider RNA strands of length $L = 20$. There are 4^{20} such RNA sequences but an exhaustive computational search reveals only 2741 distinct minimum free energy structures that emerge when these RNA strands are folded [334, p. 44]. This property has also been observed in longer RNA sequences [132, 133] and highlights the existence of large neutral sub–spaces of the genotype space, wherein every genotype corresponds to the same phenotype. The first conclusion to be drawn from this finding is that a high degree of genetic heterogeneity need not necessarily correspond to the same degree of variation in phenotypes. This phenomenon is well understood in the study of cancer, where *synonymous* mutations are throught to occur at a fixed rate and can be counted as a *molecular clock* to determine the age of a tumour [299, 191].

Given that the RNA folding model suggests that the GP–mapping is highly degenerate, a number of questions arise regarding the likelihood and accessibility of phenotypes. The RNA model has been used to answer these questions. Exhaustive analysis of the space of short RNA sequences and their MFE secondary structures indicates that not all secondary structures are equally likely to be formed. Indeed, the majority of RNA strands fold into one of a small number of frequently occurring structures. For example, for sequences consisting solely of G and C there are more than 10^9 possible sequences of length 30, but approximately 2×10^5 MFE secondary structures formed when they fold. Thus, each MFE secondary structure accounts for, on average, 5000 sequences. If a structure is defined as *frequent* when it corresponds to

more sequences than this average, then only 10.4% of MFE secondary structures are frequent but over 93% of sequences fold into them. Despite the extremely highly dimensional genotype space, the phenotype space is considerably smaller, and further, amongst those phenotypes relatively few make up the majority we might expect to find.

This degeneracy in the GP–mapping is not simply an interesting fact, it has profound implications for how evolution proceeds. In particular two aspects of the GP–mapping, *robustness* and *evolvability*, are heavily dependent on this degeneracy and the shape of the neutral spaces within G. Robustness refers to the capacity of an organismal phenotype to remain unchanged when genetic mutation occurs and evolvability refers to the capacity of a population to generate phenotypic variability through genetic mutation. These two properties of the GP–mapping are at first glance inversely related, yet both are critical for the survival of a species. Consider an isogenic population with population genotype g of length L corresponding to a given MFE secondary structure p. There are $3L$ possible mutants of g. If R of these mutants are neutral (having MFE secondary structure p) and S are not, then $R+S=3L$. Wagner [333] refers to R as the *genotype robustness* of g and S as the *genotype evolvability* and notes that they are inversely related. Here we see the apparent trade–off, high robustness (large R) necessarily reduces the number of mutations, S, that induce phenotype change.

Neutral spaces in the genotype space, which are induced by the GP–mapping, permit a population of phenotypically identical individuals to overcome the robustness–evolvability trade-off through the generation of genetic heterogeneity. To see how this occurs we must consider how evolution occurs on a neutral subspace of G. If the Strong Selection Weak Mutation criterion is satisfied (Section 2.1.2), then the population is necessarily isogenic and evolution proceeds as a random walk through the neutral space. In this case the robustness–evolvability trade-off cannot be overcome.

If instead the population can be genotypically heterogeneous, then the existence of neutral spaces increases both mutational robustness and evolvability. Consider a collection of genotypes $G' \subset G$, wherein each genotype corresponds to the same MFE secondary structure, p, and any pair of genotypes is connected by a sequence of nucleotide substitutions that do not alter p. G' is then a contiguous neutral subspace of G or a *neutral network*. Suppose that the phenotype p is optimal such that selection acts to preserve it (i.e. $\mathbf{e} = p = p^*$) . From an initially isogenic population, say with genotype g, the accumulation of genetic mutations will cause the genotypes within the population to diversify and to resemble a 'cloud' in G' known as a *quasispecies* [78, 79, 240]. When non–neutral mutations occur, the resulting individuals will be less fit than those of phenotype p and will be removed from the population by natural selection. Van Nimwegen et al. [326] demonstrated through graph–theoretic techniques and simulation that through this process the average robustness of an individual within the population, as measured by the number of neutral mutation neighbours (R), increases as the population evolves. This is because the quasispecies does not spread randomly through the neutral network but instead becomes centred where there is a high density of neutral mutations. In short, robustness *evolves*.

The genotypically heterogeneous population will still contain individuals with genotypes that have low mutational robustness. These individuals lie at the boundary of the neutral space for p and the neutral space for a different phenotype. Thus, they have higher sequence evolvability S than other individuals and, critically, allow the population to generate phenotypic variation despite the majority of individuals exhibiting high robustness. Exhaustive computational analysis suggests that the neutral networks for frequent secondary structures can span the whole genotype space, permitting the mutation of *every* nucleotide (in some order) whilst maintaining phenotype [93, 290]. Mathematical analysis from the theory of percolation can provide statistical constraints for when a neutral network will span the genotype space in this way [127]. These large neutral networks can border neutral networks for a large number of other phenotypes and thus, once the quasispecies spreads to contain genotypes from the whole neutral space for p, increase evolvability. It follows that neutral spaces for a given phenotype, which are induced by the GP–mapping, can increase both robustness and evolvability. This resolution of the apparent trade–off between robustness and evolvability was presented by Wagner [333] who refers to it as a 'paradox' resolved.

Critically, Wagner [334] highlights that this argument for the evolution of robustness and evolvability is not solely limited to RNA secondary structures, but could apply in any instance where evolution occurs on genotypes drawn from a discrete sequence space. Comparing this perspective with our abstract evolutionary framework (Section 2.1) wherein the genotypes of G are *necessarily* discrete, we see that a similar study to that undertaken by Van Nimwegen et al. [326] could be applied to *any* model GP–mappings. The reasons for choosing the RNA secondary structure model are purely pragmatic, to balance biological complexity with computational tractability.

Evolutionary Simulations with the RNA Model

Consider now evolution towards some target optimal structure p^* from another phenotype p. A number of *in silico* studies have utilised the RNA secondary structure model, coupled with the weak selection/strong mutation ($N\mu \gg 1$) population dynamics, to explore the effects of neutral networks on evolutionary dynamics [11, 93, 92, 159]. The key finding of these studies is that the evolutionary dynamics follow a pattern of *punctuated equilibria*; long periods of phenotypic stability followed by rapid and large phenotypic change. This pattern emerges as the population will diversify as neutral mutations accumulate and the quasispecies spreads to cover the neutral network. For some time this spreading will only encounter genotypes corresponding to p or deleterious mutants of these genotypes. Thus, the phenotype will remain unchanged whilst genotypic heterogeneity increases. Eventually (if one exists), a fitter phenotype p' bordering the neutral network for p will be found and a selective sweep will occur. p' will come to dominate the population and genetic heterogeneity will decrease. The process then repeats until a local fitness optima p is found. The pattern of punctuated equilibria is common over evolutionary timescales [124].

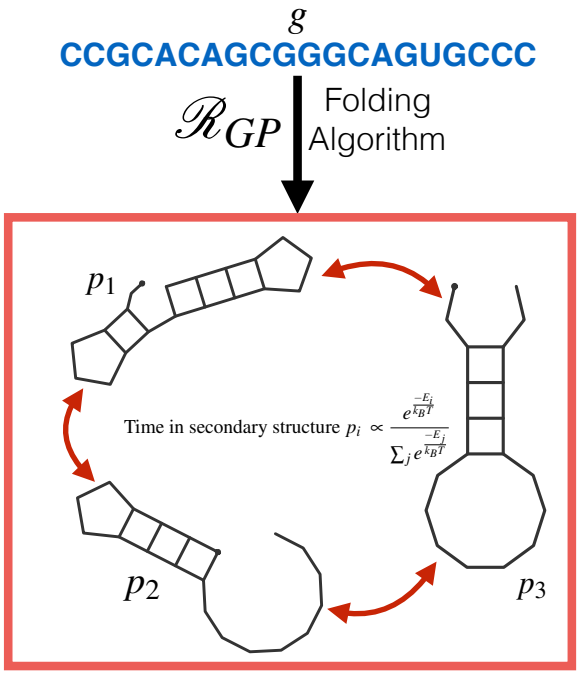

Fig. 2.6 **The stochastic RNA model.** Thermal fluctuations cause RNA strands to spontaneously unfold and fold again into different secondary structures. In this model the phenotype is not fixed but switches between secondary structures with free energy below some threshold. The time spent in each secondary structure is dependent on an environmental factor, the temperature T. This phenotypic switching is an example of bet–hedging in the GP–mapping.

Fontana and Schuster [93], and later Stadler et al. [310], characterised the phenotypic transitions within these evolutionary trajectories as either *continuous*, a change of the dominant population phenotype to another *phenotypically* closely related one (as defined by the metric function d in Equation 2.9), or *discontinuous*, large phenotypic changes. They found that the initial stages of adaptation are characterised by the fast accumulation of continuous phenotypic change and that this stage is followed by the pattern of punctuated equilibria; long periods of phenotypic stasis followed by discontinuous phenotypic change. These studies show that it is in the first (continuous) stage of the evolutionary trajectory that most of the progress (with regard to the metric d) towards the target structure is made, and thus in this stage that most of the increase in fitness happens. This corroborates the findings from studies of Fisher's model, that rapid fitness increase occurs during the early stages of adaptation to a new environment.

Evo–Devo and Phenotype Switching in RNA Models

The structure of folded RNA strands are only metastable. Owing to thermal fluctuations, an RNA strand will spontaneously unfold and refold into any of a number of low–free–energy secondary structures (Figure 2.6). The time spent in a secondary structure, p_i, with free energy,

E_i, satisfies,

$$\text{Time in secondary structure } p_i \propto \frac{e^{\frac{-E_i}{k_B T}}}{\sum_j e^{\frac{-E_j}{k_B T}}}, \tag{2.10}$$

where $k_B = 1.4 \times 10^{-23}\text{JK}^{-1}$ is the Boltzmann constant, T is the temperature in kelvin and the sum is over all other secondary structures. Algorithms that calculate all secondary structures within some temperature range of the MFE structure have been derived [221]. Thus, it is possible to approximate the stochastic switching between a number of low energy secondary structures computationally.

We note that this stochastic switching, in the context of our evolutionary framework, is precisely phenotype switching or *bet–hedging* as discussed in Chapter 1. It follows that the RNA secondary structure model has the potential as a tool for studying the evolution of bet–hedging by considering a population of individuals (with genotypes determined by the RNA strand) wherein each has a secondary structure that can spontaneously switch. Further, as the switching probabilities are dependent on an environmental factor, the temperature T, this bet–hedging is environmentally (or *plastically*) modulated. This model was studied by Ancel and Fontana [11] who considered the effects of phenotypic switching on robustness, evolvability and modularity. They note that, under the selective pressure to evolve towards a target structure, bet–hedging is costly and is minimised by evolution. Interestingly, this reduction in phenotype switching results in an evolutionary dead–end wherein no further progress is made. This dead–end occurs as the neutral networks of MFE secondary structures display a distinct pattern of heterogeneity. Those sequences with the most thermodynamically stable MFE secondary structures (i.e. those that spend the most time in their MFE structure) lie at the 'centre' of the neutral networks for that structure. Thus, the population quasispecies becomes centred in these neutral spaces because sequences near to the boundary of a neutral space for a different MFE secondary structure display greater switching rates and thus lower fitness. These sequences are removed from the population by natural selection and evolvability is reduced. In this case bet–hedging is an impediment to adaptation and and the accessibility of a globally optimal phenotype. The evolutionary consequences of stochastic switching in fluctuating environments, where bet–hedging has been demonstrated to offer a fitness advantage, have not been studied in the RNA model.

A second, more realistic, model of secondary structure folding considers the temporal kinetics of the fold. The dynamic programming algorithm to find a minimum free energy secondary structure ignores many of the physical properties of an RNA fold and instead assumes an instantaneous fold into the most thermodynamically stable structure. In reality, this assumption is flawed as some parts of the structure fold into place before others. This can cause an RNA strand to reliably fold into a secondary structure that is not the MFE predicted by the dynamic programming algorithm [91]. The partially folded RNA strand can be considered to exist on a rugged energy landscape and to move 'down–hill' to more stable structures. This energy landscape is an analogue to the *epigenetic landscape* introduced by Waddington [330]

as a metaphor to describe the role of development pathways in the *canalisation* (robustness to genetic or environmental variability) of phenotype. In Waddington's work, the developmental pathways are viewed as inducing a landscape and development as a ball 'rolling' down this landscape. Highly stable phenotypes are then seen as occurring due to high 'ridges' in the landscape that require large perturbations to knock the 'ball' over. The physical energy landscape of the kinetic folding model is in direct correspondence with the metaphorical model of Waddington, suggesting that the RNA model is an ideal candidate for studying the interplay between development and evolution. To our knowledge few studies have taken this perspective, perhaps because the RNA model misses a key aspect of evolutionary developmental biology: that the developmental pathways (i.e. the shape of the landscape) can also evolve, whereas in the RNA model they are fixed by the laws of physics.

Protein Folding and Combinatoric Limits

Alternative secondary structure models can also be used to study the GP–mapping. The prediction of folded protein structure has seen extensive study, as understanding of these structures is essential in the study of a number of biological processes. However, whereas RNA sequences constitute strings built from four nucleotides, the protein folding problem seeks to find the three dimensional MFE structures formed when chains built from a possible 20 amino acids fold. As such, the space of possible amino acid chains is astronomical. Further, prediction of protein folding requires the minimisation of free energy that is determined by the interactions between amino acids in 3D space. These predictions are computationally complex, preventing a statistical analysis similar to that performed for RNA.

Through simplifying assumptions and abstraction some progress has been made. Studies of *lattice protein folding*, a simplified protein folding problem in which the amino acids in the chain must occupy positions on a (2D or 3D, square, triangular etc) lattice, have yielded a number of results that corroborate those of RNA studies [45]. In this model, genotypes are taken as sequences of amino acids, or in further simplified models, sequences of two kinds of monomer: hydrophobic (H) or polar (P) amino acids. When folded, these monomers are restricted to a lattice structure. This lattice structure allows for the explicit enumeration of all folded structures, and thus the identification of the minimum free energy fold.

Where the protein folding model, or at least the simplified analogues, are in agreement with the RNA model is that the GP–mapping is massively degenerate. The space of possible amino acid chains is astronomical with 20^L possible chains of length L and yet the space of tertiary structures is miniscule in comparison. For example, studies of $L = 100$ suggest that there are less than 10,000 folded protein stuctures [56]. Just as with the RNA model, these tertiary structures are not evenly distributed amongst the chains and a small number of *frequent* structures account for the majority of amino acid chains [185]. This suggests that degeneracy and neutrality may not be simply an interesting property of the RNA model, but a key property of GP–mappings in nature.

There are at present a number of open problems with respect to protein folding as a model GP–mapping [334, p. 77]. Taking the results of the RNA secondary structure as a guide, we find a number of results that are intractable to verify in the protein–folding model. For example, whilst RNA secondary structures exhibit a 'shape–space covering' property wherein each secondary structure lies within a small number of point mutations of any other, the equivalent property for folded proteins has not been established. In fact, empirical evidence suggests that folded protein structures may be closely clustered in sequence space and that some structures may lie 'far', with respect to the number of point mutations, from one another [35, 36].

A major complication of the protein folding model is that the folding action can be approximated at several levels of complexity. For example, restrictions of the amino acid repertoire, restriction of the amino acids to either hydrophobic or polar, the imposition of a lattice structure, or the omission of space entirely, have all seen use as simplifications of the folding process. More complex models, for example those considering the kinetics of folding opposed to solely the thermodynamic properties of the final structure, have also been proposed. Many of these models give contradictory results [39], suggesting that the future for protein secondary structure modelling lies in deriving tractable, accurate models.

Protein Quaternary Structure

Once folded into complex tertiary structures, proteins combine together to form large, complex structures known as *quaternary structures*. This process too forms the basis of a model of a GP–mapping. In this model, individuals proteins take the role of genes and are represented by *Polyominos* with genotype–specific constraints on how they can be tiled in the plane [129, 130]. Tilings of the genetically determined Polyominos represent phenotypes and fitness is determined by comparison to a target phenotype. The results of the quaternary structure model mirror those of other structure prediction models. Namely, a high degree of robustness and redundancy is observed [130], as well as a strong bias towards a small set of phenotypes. The space–shape covering property also holds under this model [129]. The preservation of these properties between biologically–inspired structure–prediction models suggests that they could be universal properties of biological evolution.

2.2.4 Network Models

A key observation from quantitative genetics is that not all genes are equal in their potential to adaptively affect phenotype. Mutations occur with (approximately) equal probability throughout the genome, and yet the genetic drivers of phenotypic differences between, or within, species have been observed to accumulate in so–called *hot spot* areas of the genome [312]. This non–uniform distribution of *evolutionarily–relevant* genetic variation can be attributed to the differing regulatory actions of some genes on others that drives epistasis and pleiotropy. Mutations in highly pleiotropic genes, those contributing to many aspects of the phenotype, are considerably less likely to be adaptive than mutations in genes contributing to a single phe-

notypic trait, and thus are much less likely to reach evolutionary fixation. This phenomenon, wherein genes regulate one another such that mutations at different genomic sites induce substantially different phenotypic effects (with respect to fitness), is poorly accounted for in the fitness landscape or secondary structure model.

A second biological phenomenon that is absent from these models stems from developmental biology. Consider the following. A critical process in embryonic development is the folding of the sheet of early cells to form more complex tissues (specifically, the precursor to a central nervous system). A cluster of cells called the *Spemann Organiser* has been identified as responsible for this process both in mammals and some invertebrates, for example tunicates, and those genes expressed within the mammalian Spemann Organiser during development are present in the tunicate genome. What is remarkable is that these genes within the tunicate genome play no role in the development of the Spemann Organiser [51, 198] (although they are implicated in other aspects of development) – evolution has found an alternative solution to this complex developmental process. From the perspective of our abstract evolutionary modelling, we have two species in which, in a restricted sense, the genotype, g, and phenotype, p, are equal, but the GP–map *itself* differs. Even if the genotype and phenotype are identical for two species, a different GP–mapping means that the same mutation to g may manifest itself in distinct phenotypic differences – an important consideration if we are to predict evolution. In the worst case, the existence of many possible GP–mappings may render evolution *un*predictable. However, this need not be the case, for example, if the many different mappings capable of mapping g to p share a common structure or restriction that permits evolutionary predictions. This scenario of differing GP–maps cannot be captured by the secondary structure model as it is the laws of physics (as interpreted through a dynamic programming algorithm) that determine the folded structure. It is for this reason that network models encoding gene–gene or protein–protein interactions have arisen as a means to study the GP–map in both evolutionary and developmental biology.

Gene Regulatory Models

The gene regulatory network (GRN) model of the GP–mapping emerged in the 1970s as an abstract model to study regulatory feedback between genes and their role in the development of an organism [171]. The GRN model assumes the existence of L mutually regulatory genes and defines an initial conditions for the activation/expression of these genes by

$$\mathbf{S}(0) = (S_1(0), \ldots S_L(0)) \in \{-1, 1\}^L \tag{2.11}$$

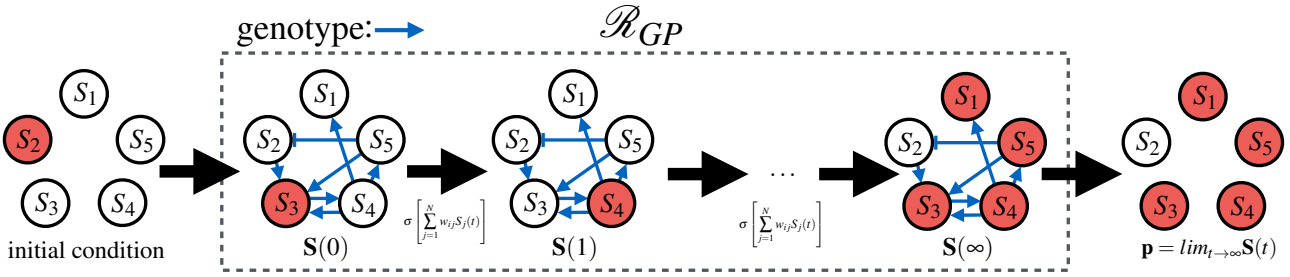

Fig. 2.7 **Schematic of the gene regulatory model for a GP–Mapping.** The genotype (blue arrows) defines the regulatory actions of genes on one another. The phenotype, p, is determined by the stable state for the network, under the update rule specified in Equation 2.16. The initial condition of the network can be taken to be genetic, environmental, or arbitrary, depending on the specifics of each study.

The cross– and auto–regulatory interactions of these genes are defined by a network $W = [w_{ij}]_{L\times L}$ where,

$$w_{ij} = 1 \text{ if gene } i \text{ upregulates gene } j, \tag{2.12}$$

$$w_{ij} = -1 \text{ if gene } i \text{ downregulates gene } j, \tag{2.13}$$

$$w_{ij} = 0 \text{ if there is no regulatory effect of gene } i \text{ on } j. \tag{2.14}$$

The expression of the L genes at each time t is defined by values

$$\mathbf{S}(t) = (S_1(t), \ldots S_L(t)) \in \{-1, 1\}^L \tag{2.15}$$

and subject to synchronous updates over a time step τ according to

$$S_i(t+\tau) = \sigma\left[\sum_{j=1}^{L} w_{ij} S_j(t)\right]. \tag{2.16}$$

In the simplest, and most common, version of the GRN model σ is the sign function and the phenotypes are taken as stable expression profiles,

$$\mathbf{p} = lim_{t\to\infty}\mathbf{S}(t). \tag{2.17}$$

Where no stable expression profile exists, the phenotype is assumed to be non–viable (having fitness equal to zero)[3]. The phenotype space in the GRN model is thus given by $\mathscr{P} = \{-1, 1\}^L$, although more complex GRN models permit the genes to take real values and the function σ to take alternative forms. A schematic of the gene regulatory network GP–mapping, along with example updates for the gene expression values as determined by W, is shown in Figure 2.7.

[3]We note that oscillatory dynamics of protein expression are in fact common, for example in the cell cycle [244], and thus it is not clear whether this definition of phenotype is too restrictive.

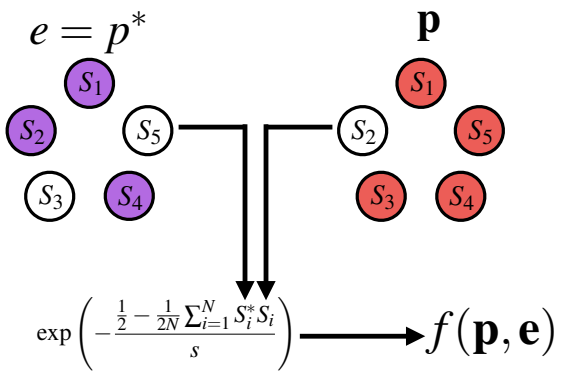

Fig. 2.8 **The fitness function for the GRN model.** A metric value (in this case, Hamming distance) between a phenotype, $\mathbf{p}$ (red), and a target phenotype $e = \mathbf{p}^*$ (environment, purple) is calculated and a transformation of this value is used to determine the fitness, $f(\mathbf{p},\mathbf{e})$, of an individual with phenotype $\mathbf{p}$ in the environment $\mathbf{e}$.

Evolutionary modelling has utilised GRNs as a model for development and focused on how changes to the underlying network, W, affect evolutionary phenomena such as evolvability or robustness. From the perspective of our abstract evolutionary modelling, the genotypes, G, specify the *network* and not simply the L *genes* whose expression is being modelled. Thus $G = \{-1,0,1\}^{L^2}$ and the mutation relation μ corresponds to changing an entry within the matrix W specifying the regulatory action of one gene on another. The environment, $\mathbf{e}$, is determined by the specifics of different studies, but a common approach is to define a target phenotype, p^*, that is globally optimal and measure the fitness as a function of the distance from this globally optimal phenotype (Figure 2.8). For example, Wagner [331] used a Gaussian fitness function of the Hamming distance between the two phenotypes

$$f(\mathbf{p},\mathbf{e}) = \exp\left(-\frac{\frac{1}{2}-\frac{1}{2N}\sum_{i=1}^{N} S_i^* S_i}{s}\right) \tag{2.18}$$

where $s > 0$ denotes the strength of selection. The initial conditions for gene expression differ between studies and can be genetically, environmentally, or arbitrarily defined depending on the purpose of the study.

This model has also been used to study how regulatory networks organically grow through mutations that duplicate genes, delete genes or alter regulatory interactions [331]. Wagner [332] demonstrated that for GRNs evolving under stabilising selection, evolution tended towards genotypes with few phenotypically different mutational neighbours, highlighting that robustness to mutation is an evolvable property. Ciliberti et al. [52] extended this work by introducing a geometric description of the 'meta–network' of neutral spaces and found that mutational robustness and evolvability are negatively correlated, but that for phenotypes with a fixed measure of robustness, the shape of the corresponding neutral space in G can increase evolvability. These results suggest that both robustness and evolvability can evolve as properties of the GP–mapping and mirrors the results of the RNA folding model. Crombach and Hogeweg [60] simulated the evolution of GRNs in fluctuating environments by periodically

switching the environment between two target phenotypes, p_1 and p_2. This work found that evolution converged to genotypes on the 'boundary' between two neutral spaces in G corresponding to the two phenotypes, a result which was interpreted as the evolution of evolvability.

Recent studies have used the GRN model to explain the apparent paradox of the evolution of evolvability by drawing parallels with learning theory [186, 187]. These studies highlight the correspondence between the generalisability of a learning algorithm and the capacity for organisms to develop phenotypes well adapted to previously unseen environments. In these studies, experimental *in silico* evolution of recurrent GRNs was performed with fitness taken as a measure of similarity between the organismal fitness and an environmentally–determined target phenotype. Where the environment was taken to fluctuate, and a cost imposed on the complexity of the GRN, evolvability (or generalisability) was observed to evolve.

As a conceptual tool to explain non–genetic heterogeneity of phenotypes, Huang [154] considered the stable expression profiles (i.e. phenotypes) induced by gene regulatory networks as attractors in a high–dimensional space [38, 155]. In this model, gene expression profiles are assigned a "potential" corresponding to the stability of the expression profile to stochastic molecular interactions. This assignment creates an *epigenetic landscape* (in the sense of Waddington [330] – arising not from single genes but the interaction of many). The local minima of potential in this landscape correspond to stable expression profiles (or equivalently phenotypes) and those gene expression profiles with higher potential will move down–hill through regulatory feedback mechanisms until a stable expression profile is found. Huang argues that non–genetic heterogeneity can then be explained by the existence of multiple accessible stable expression profiles, and that phenotypic switching is due to stochastic fluctuations causing jumps between stable states. These fluctuations can be either intra–cellular, for example due to intrinsic thermal fluctuations of molecules, or extra–cellular and environmentally driven. Thus, Huang's model of phenotypic heterogeneity can account for both bet–hedging and phenotypic plasticity, but no distinction is made between these two phenomena.

Experimental evidence for the attractor states model was found by monitoring the expression of over 2700 genes simultaneously during neutrophil differentiation [156]. In these experiments, the gene expression profiles were perturbed and the trajectory of the expression recorded over seven days. Remarkably, from different (perturbed) gene expression profiles, the trajectory converged to the same stable expression state, indicating the existence of a phenotype with a large 'basin of attraction' in gene–expression space as posited in the GRN model. These findings lend credence to the epigenetic landscape model and, crucially, suggest that the abstract GRN model could be coupled with empirical datasets to form a predictive model of the GP–mapping.

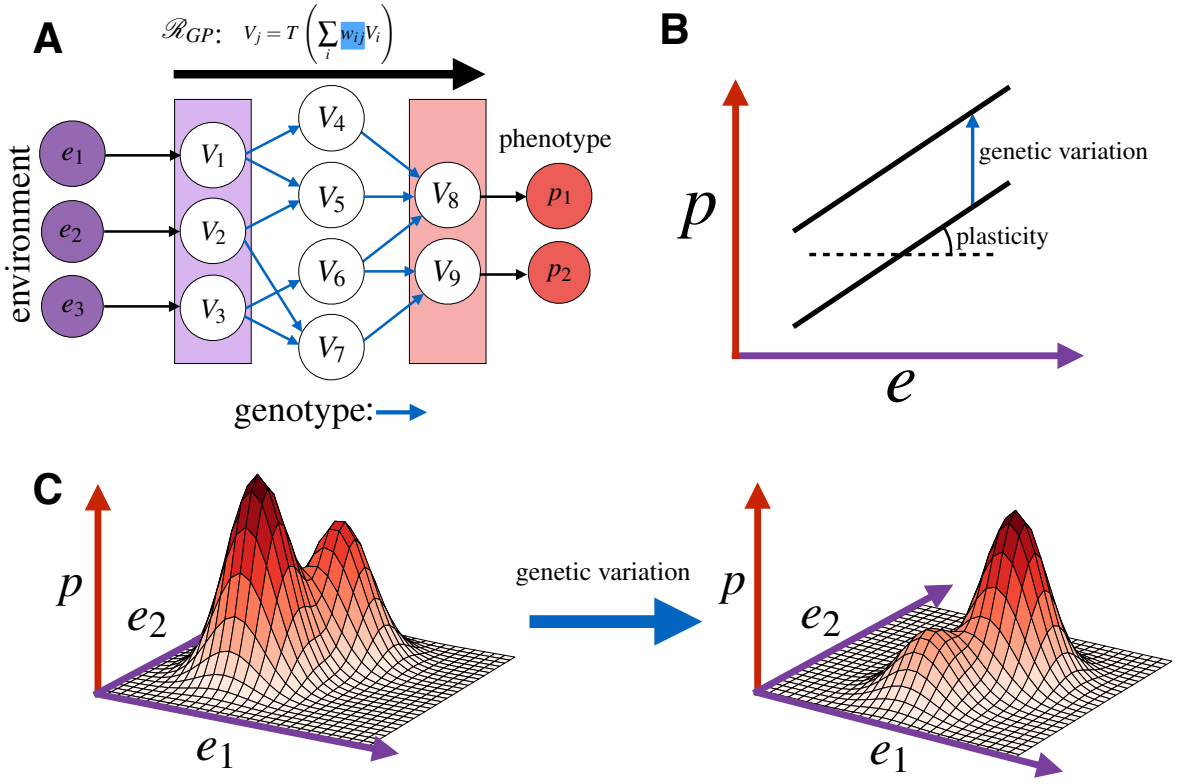

Fig. 2.9 **The neural network model and phenotypic plasticity. A)** Schematic representation of the (feed–forward) NN model. Environments, $\mathbf{e} = (e_1, e_2, e_3)$, form the input (purple), the genotype determines the internal weights of the NN (blue) and the phenotype (red) is computed by evaluation of the NN. **B)** The traditional (linear) representation of a reaction norm mapping an environment, e, to a phenotype p. Genetic variation results in changes to this relationship. **C)** A higher dimensional and non–linear reaction norm, these more complicated forms of the reaction norm can be represented by the NN model.

Phenotypic Plasticity and the Neural Network Model

Gerlee and Anderson [106] introduced an alternative network model for the GP–mapping in order to study the effects of phenotypic plasticity and spatial constraints in growing tumours [108, 109]. In this model, genotypes are taken to define the internal wiring of a neural network that maps r real–valued environmental inputs ($E = \mathbb{R}^r$) to s phenotypic traits ($\mathscr{P} = [0,1]^s$). Thus, $G = \mathbb{R}^{k^2}$ where k is the number of nodes within the artificial network (including the input and output nodes). Here, we denote the nodes that comprise the neural network, including the s input nodes and r output nodes, by variables V_i for $i \in \{1, \dots k\}$. The GP–mapping, $\mathscr{R}_{\text{GP}}$, is the evaluation of the neural network defined by $g = [w_{ij}]_{k\times k}$ on an input vector $e \in E$ (Figure 2.9(A)).

The original formulation of this model was a *feed–forward* neural network, wherein the nodes are organised into layers and the internal wiring determined by $g \in \mathscr{G}$ only permits a node to influence those in the next layer (Blue arrows, Figure 2.9(A)). The phenotype is determined by setting the input nodes equal to some environmental state and evaluating the internal nodes in order of layer according to

$$V_j = T\left(\sum_i w_{ij} V_i\right) \tag{2.19}$$

where $T : \mathbb{R} \to [0,1]$ is a thresholding function. In their work Gerlee and Anderson [106] take $T(x) = \frac{1}{1+e^{-2x}}$.

Gerlee and Anderson [106] embedded this model into a hybrid cellular automaton [12] to demonstrate the evolution of glycolytic or acid–resistant phenotypes in cancer in response to differing local concentrations of acid, glucose and oxygen, as well as spatial constraints. Harsh environments, for example, dense extra cellular matrix (ECM) or low oxygen concentrations, drove the evolution of a glycolytic phenotype. This study also demonstrated that the network can induce different tumour morphologies, for example, a 'fingered' morphology evolved in low oxygen environments and a more compact morphology in environments with dense ECM. Gerlee and Anderson [107] studied this phenomenon in a more general framework to derive an approximate relationship between the parameters governing the dispersal of nutrients and the morphology of a growing colony. These studies are of particular interest as they clearly demonstrate that the impact of the genotype–phenotype mapping stretches beyond the cellular phenotype and shapes the aggregate behaviour of a population. This *meta*–phenotypic behaviour can be much more complicated than simple morphology and could, for example, encapsulate developmental processes or self–organising behaviour. It is from this perspective that the vital role of the GP–mapping in reconciling the modern evolutionary synthesis with developmental biology is most clear [258].

In a later study, Gerlee and Anderson [109] considered how the genetic and phenotypic heterogeneity of population changed over the course of evolution. The genetic heterogeneity, as measured by the Shannon index of the genotypes in the population, was observed to

increase throughout the simulation whilst the phenotypic heterogeneity decreased. This result corroborates those derived from the RNA secondary structure model that suggest a high degree of degeneracy in the GP–mapping. Interestingly, the degree of genetic heterogeneity was also higher in harsh (e.g. low oxygen) environments. This result bears a strong resemblance to that of the RNA model suggesting that a more diverse genetic quasispecies can serve to increase evolvability. However, no attempt has been made to verify if this is the case as the relationship between the heterogeneous population of genotypes and their phenotypes, and in particular the mutational 'closeness' to other phenotypes (an analogue to the evolvability parameter S of the RNA model) was not measured.

More recently, an extension of the neural network model to include *recurrent* networks was used to compare the efficacy of theoretical treatment strategies targeting single proteins, whole pathways or organismal phenotype [110]. In this model the genotypes are unrestricted with respect to the influence one node may have on another. The genotype–phenotype mapping, $\mathscr{R}_{\mathrm{GP}}$, is then determined by solving a dynamical system defined by

$$\frac{dV_j}{dt} = T\left(\sum_i w_{ij}V_i(t)\right) - \lambda V_j(t) \tag{2.20}$$

Where λ is a decay term for the artificial 'proteins' inside the model. Under this more complex model the phenotype need not be stable but could exhibit oscillatory or even chaotic behaviour. Because the evaluation of a recurrent neural network is dependent on the initial values for all of the nodes, Gerlee et. al. required that the initial values of non–environmentally determined nodes are zero. However, an alternative approach would be to permit nodes to retain their values (subject to decay) throughout the cell–cycle. This alternative modelling assumption could permit the study of epigenetic cell memory which has been demonstrated to confer an evolutionary advantage [164, 44] but which is, at present, rarely studied in the context of the GP–mapping. In Chapter 5, we will introduce a model GP-mapping that exhibits such epigenetic memory in bet–hedging.

At present the studies of the neural network model have been restricted to simulations of cancer progression or colony growth. In fact, the authors have missed an important parallel between their work and previous biological modelling that becomes clear when the GP–mapping is considered from our abstract perspective. In 1985 Via and Lande [328] introduced the concept of the *reaction norm* — a mapping from a single environmental parameter to a phenotypic value as a simple model of phenotypic plasticity (Figure 2.9(B)). The work of Gerlee and Anderson [106] is a natural extension of this concept to higher–dimensional functions that, as recurrent neural networks can closely approximate any function on bounded input space, can approximate any biologically occurring reaction norms (Figure 2.9(C)). Further, as neural networks can be trained without knowledge of the underlying mechanism and from limited datasets, the neural network model has the potential to *predict* non–genetic adaptation to previously unseen environments.

2.3 The Future of GP–Mapping Models

Our survey of the literature regarding mathematical and computational models of the GP–mapping has highlighted a number of important, and in many cases unintuitive, evolutionary phenomena. By considering these different models under the umbrella of an abstract evolutionary framework, we have been able to identify a number of similarities between models that would be missed by considering these models on an *ad hoc* basis. To conclude this chapter we present a summary of the important properties of the GP–mapping predicted by theoretical modelling (See Table 2.1). We pay special attention to those properties that are unintuitive or contradictory to the 'central dogma' of evolutionary biology (genes $\rightarrow$ mRNA $\rightarrow$ proteins $\rightarrow$ phenotype) or the 'genes as blueprints/programs' perspective. These properties serve as a cautionary warning against the reductionist approach of single gene knockouts with the assumption of *ceteris paribus* (all other things held equal). We are not the first to make this argument, Both Stephen Jay Gould and Denis Noble did so over 10 years ago [123, 239], but as the reductionist paradigm is still pervasive in molecular biology there is no harm in making it again. Finally, we highlight those aspects of the GP–mapping that have been overlooked in previous theoretical modelling and provide a partial overview of evolutionary questions that could form the basis of future theoretical studies of evolution. It is some of these questions, specifically those relating to non–genetic influences on organismal phenotype, that form much of this thesis.

2.3.1 Properties of the Genotype–Phenotype Mapping

Neutrality and Degeneracy

A key property of models which include explicit GP–mappings, for example secondary structure prediction and network models, is a high degree of degeneracy in the relationship between genes and phenotypes. Specifically, many genotypes will map to a single phenotype, inducing large neutral subspaces of the genotype space. Perhaps the most thorough studies of this phenomeon have been undertaken with the RNA–secondary structure model, wherein the degree of neutrality has been quantified, the distribution of phenotypes in genotype space estimated and the evolutionary impact of this neutrality explored. From the perspective of evolving drug resistance this neutrality can have important consequences. For example, neutrality can induce a two–stage pattern of evolution, wherein fitness first rapidly increases before following a pattern of punctuated equilibira, and may explain many aspects of disease progression [61]. The rapid increase in fitness associated with the early stages of evolution may explain why evolutionary escape from therapy is so prevalent, whilst the later stages of punctuated equilibria can explain why diseases that appear stable for many years may suddenly worsen — a common scenario in cancers and HIV infections.

Understanding the relationship between neutral mutations, the associated genetic heterogeneity, and the potential for pre–existing or *de novo* evolution of drug resistance is of

Model	Fisher's Geometric Model	Fitness Landscapes	Secondary Structure Model	Gene Regulatory Networks	Neural Network Model
Genotype Space	$\mathbb{R}^M$	Σ^L	Σ^L	$G = \{-1,0,1\}^{L^2}$ (Regulatory action of L genes)	$G = \mathbb{R}^{k^2}$ (specifying a neural network)
Mutation Relation	$\mathbf{p} \to \mathbf{p} + \mathbf{m}$, $\mathbf{m} \in \mathbb{R}$	Point mutations	Point mutations	Alterations to the internal connections	Alterations to the internal weights
Phenotype Space	$\mathbb{R}^M$	$\mathbb{R}$	$\mathscr{P} = \{'(',',',')'\}^*$ (RNA secondary structures)	$\mathscr{P} = \{-1,1\}^L$ (stable expression profiles)	$\mathscr{P} = [0,1]^s$ (final values for s output nodes)
GP mapping	Identity	$f : \Sigma \to \mathbb{R}$	Prediction of mfe secondary structure	$lim_{t\to\infty}\mathbf{S}(t)$ where $S_i(t+\tau) = \sigma\left[\sum_{j=1}^{L} w_{ij}S_j(t)\right]$	$V_j = T\left(\sum_i w_{ij}V_i\right)$ (evaluation)
Environment	Target phenotype $e \in \mathscr{P}$	Fixed	Target phenotype $e \in \mathscr{P}$	Target phenotype $e \in \mathscr{P}$	$E = \mathbb{R}^r$ (inputs)
Fitness	$f : \mathbf{p} \mapsto \exp\left(-\frac{\|\mathbf{p}-\Theta\|^2}{2}\right)$	Identity	$f(p,e) = w(d(p,p^*))$ (d a metric, w a transformation)	$f(\mathbf{p},\mathbf{e}) = \exp\left(-\frac{\frac{1}{2}-\frac{1}{2N}\sum_{i=1}^{N} S_i^* S_i}{s}\right)$	Not specified / emergent

Table 2.1 **Summary of abstract models for the genotype–phenotype mapping.** For each of the models discussed in this chapter, the key components of an evolutionary system are shown: The genotype space and mutation relation, phenotype space, GP–mapping, environment and fitness mapping.

particular importance to understanding the progression of cancers and remains an active area of research [364, 104]. It is clear from our survey of the literature that a focus on the GP–mapping is essential in furthering our understanding of the impact of intra–tumoural heterogeneity and developing effective diagnostic and therapeutic techniques. As outlined above, one of the most convincing arguments for this perspective does not originate from cancer research at all, but rather from studies of the secondary structure model and other theoretical GP–mappings. Here we see the value of the abstract modelling approach. By building minimal abstract models, we can identify phenomena that may be clouded by the complexity of biological reality and, with the clarity offered by these simple models, generate hypotheses and design experiments to test them empirically.

Evolvability and Robustness

A second major contribution of theoretical GP–mappings is the resolution of the apparent paradox between robustness and evolvability. Neutral mutations can induce genetically heterogeneous but phenotypically homogeneous populations in which both robustness and evolvability are jointly increased owing to the distribution of genotypes. This phenomenon was clearly demonstrated by computational studies of the secondary structure model and

corroborated by studies of network models, suggesting that evolvability and robustness may themselves be evolvable. This observation suggests a potential strategy for an adaptive therapy that may be overlooked if the GP-mapping is ignored. It is not solely the pathogenic aspects of organismal phenotype that can be evolutionarily steered by careful selection of drugs, it is any property of an individual subject to evolution. Thus, if robustness or evolvability are themselves evolving then it may be possible to reduce them through careful choices of drugs. Of course, this theoretical approach is likely infeasible, especially in highly aggressive diseases where time is critical, but provides powerful motivation for the approach of abstract modelling. By understanding evolution as an abstract system, we are able to generate new hypotheses that would otherwise remain unexplored.

Epistasis

Studies of empirical fitness landscapes have highlighted the prevalence of epistasis in biology. A mathematical argument has shown that reciprocal sign epistasis induces rugged fitness landscapes, limiting the number of accessible evolutionary trajectories, and rendering evolution potentially predictable. Higher order epistasis, wherein the fitness contribution of a given mutation is dependent on the alleles at more than one other loci, is also commonly observed, although the implications of this, especially with respect to the predictability of evolution, have yet to be determined [348, 347]. These empirical findings demonstrate the importance of understanding gene–gene interactions in predicting evolution. These interactions have served as motivation for the gene regulatory network and neural network models of GP–mappings. These models have yielded valuable insight into the evolutionary process and it is clear that future models should also account for gene–gene or protein–protein interactions. However, this does not mean that the more simple models presented here should be discarded entirely. Indeed, many results regarding epistasis have been derived from the fitness landscape model within the past 20 years and there is still much to learn. In Chapter 3, we will use the landscape model to present a concrete argument for the intuitive concept that evolution is a ‘steerable’ process. As all evolutionary systems induce a fitness landscape (at least of *expected* fitness where the GP–mapping is non–functional), some results derived from fitness landscape studies will still apply when a more complex GP–mapping is present.

Modularity and Pleiotropy

Model GP–mappings have provided insight into the evolution of modularity and pleiotropy. Studies from Fisher’s geometric model suggest that mutations to highly pleiotropic genes are less likely to be beneficial and, hence, that evolution will proceed more slowly when organismal phenotype is less modular [247]. This phenomenon has been called the ‘cost of complexity’ and there is some empirical evidence to support this hypothesis [337]. Future theoretical studies of this phenomenon will require a much deeper understanding of organismal phenotypes. Whilst molecular reductionism has generated a wealth of genetic (and proteomic, metabolomic

etc.) data, the nature of cellular phenotypes is, comparatively, poorly understood. This represents a problem for all studies of evolution, as phenotypes are the ultimate determinants of fitness. However, the problem is compounded for studies of modularity and pleiotropy, as the simplifying assumptions made for specific studies to focus only on aspects of phenotype relevant to that study, serve to reduce the complexity that is central to the hypothesis that modularity can evolve.

Non–Functionality and Bet–hedging

A major assumption underpinning many models of the GP–mapping, and indeed much of mainstream biological thought, is that phenotypes result from genotypes in a deterministic way. The appeal of this assumption is that we can ignore the phenotypes and characterise diseases purely at the genetic or molecular level; as we have done for much of the past 20 years. However, there is clear evidence that the genes are not the sole determinant of phenotype and that the 'genes as blueprints' model is flawed [258]. Bacterial persisters [204], and their recently–discovered analogue in cancer [134, 145], suggest that some aspects of organismal phenotype may be stochastically determined. This phenomenon, known as bet–hedging, has been demonstrated to have important implications for the progression of disease and the evolution of drug resistance.

There presently exist very few models that consider non–functional or stochastic determination of phenotypes. Those that do generally consist of a deterministic process of phenotype development coupled with an associated concept of stochasticity that modifies the process. For example, Huang conceptualises development through the metaphor of an epigenetic landscape on which a 'ball' representing phenotype (as modelled by gene expression) rolls 'down–hill' to a stable configuration [156, 154, 155]. Stochasticity is viewed as random perturbations to the motion of the ball, but no attempt is made to quantify the form of these perturbations. Further, no distinction is made between intra–cellular and extra–cellular (i.e. environmental) sources of stochasticity.

A second candidate model for bet–hedging is RNA folding with spontaneous re–folding [221]. This model is a poor one for studying the evolution of bet–hedging, as the GP–mapping is fixed by the laws of physics. Although this is obviously true of the GP–mapping in genuine biology, it is demonstrated by network models that an abstracted layer of gene–gene regulatory interactions, themselves subject to evolution, provides a powerful tool in conceptualising evolution. New models that directly account for the stochasticity in the GP–mapping, and in which the magnitude of this stochasticity is subject to evolution, are needed if we are to understand the evolution of bet–hedging. It is for this reason that we develop such models from the theory of chemical reaction networks in Chapter 4 and go on to derive novel evolutionary insights from them.

Phenotypic Plasticity

Phenotypic differences amongst genetically identical individuals can also be by driven by environmental influence; endowing a species with robustness to environmental change within certain limits. This phenotypic plasticity has been observed as a driver of drug resistance in both cancers and bacterial infections, for example, through the development of efflux–pumps in response to drug treated environments [259]. Our survey of mathematical models demonstrates that plasticity is presently poorly represented and the environment is often over–simplified. In Fisher's model, the secondary structure models, the GRN model and others, the environment is taken as a single target 'optimal' phenotype and fitness as some measure of difference from this phenotype. This approach is clearly flawed. Firstly, we know from genuine biology that a single optimal solution to any 'problem' rarely exists. Examples include the Spemann Organiser discussed earlier [51, 198], or the convergent evolution of eyes or wings [194]. The second major flaw in this definition of environment is that it has no influence on the determination of phenotypes and, thus, phenotypic plasticity cannot be modelled.

The neural network model of Gerlee and Anderson [106] partially overcomes this problem by considering a complex environment of diffusible factors that explicitly influence the phenotypes of cells within a growing solid tumour. These factors are spatially heterogeneous and, coupled with constraints on growth driven by crowding effects, are also responsible for driving natural selection and determining the fitness of individual cells. The drawback of this more realistic evironmental modelling is that fitness itself cannot explicitly be defined but is rather emergent from simulation of the entire system. Such simulations are driven by a hybrid cellular automaton model and are computationally expensive. This complexity prevents the analysis of large numbers of the drug combinations or timing strategies required to design adaptive therapies through the phase *i* paradigm [291]. Thus, the future of modelling phenotypic plasticity will lie in balancing complex environmental influence on phenotype with tractable models of fitness and evolution. One current approach to this problem is the development of statistical methods for the analysis of discrete cellular automaton methods [195].

2.3.2 Conclusions

Our survey of mathematical models of the GP–mapping highlights some aspects of the mapping that are unaccounted for in many models. In particular, the non–functional aspects of the mapping, wherein a single genotype can give rise to a number of phenotypes either stochastically (bet–hedging) or through environmental modulation (plasticity), are presently understudied, despite the observed role of these phenomena in the evolution of drug resistance. The need for models displaying phenotypic plasticity was highlighted by Pigliucci [258] as a necessary step towards reconciling developmental biology with the modern evolutionary synthesis and moving away from the reductionist 'genes as blueprints' perspective.

Our survey of the literature of theoretical evolutionary models through a common algorithmic lens has highlighted the power of models of the GP–mapping in clarifying seemingly

paradoxical evolutionary phenomena. For example, the robustness/evolvability trade–off or the evolution of modularity. We believe that similar modelling will shed light on similar unintuitive properties of bet–hedging and plasticity. If we are to be able to predict evolution, as we must to design adaptive therapies, then it is precisely these properties that will be key in interpreting experimental or clinical observations. In Chapter 4 we present models of the GP–mapping wherein bet–hedging naturally arises as an evolvable property and in Chapter 5 extend this model to account for epigenetic inheritance. Using these models we explore the role of the structure of the GP–mapping in the evolution of bet–hedging and the implications for the success of adaptive therapies for disease with bet–hedging–driven resistance.

Chapter 3

Steering Evolution with Sequential Therapy

The fitness landscape has seen extensive use as a tool to understand the predictability, repeatability and accessibility of evolutionary trajectories (See Section 2.2.2). Recently, novel experimental techniques have permitted the derivation of fitness landscapes over small combinatorically–complete sub–spaces of the genotype space [71]. These empirical landscapes, coupled with the mathematical theory of population genetics, offer an ideal starting point in exploring the potential for adaptive therapies to steer evolution and prevent the emergence of drug resistance. In this chapter, we introduce a simple model for the evolutionary dynamics of an asexually reproducing population that exists under the selective pressure of an environmentally determined fitness landscape. We go on to parametrise this model with empirically derived landscapes for *E. coli* in the presence of antibiotic drugs. Using this simple evolutionary model, we highlight the potential pitfalls of sequential drug therapy and outline model–derived clinical guidelines for effectively designing multi–drug treatments.

3.1 A Mathematical Model of Evolution on Fitness Landscapes

Consider an asexually reproducing population with population size N and mutation rate u. Denote by s the *selection coefficient* (or *fitness benefit*) of an invading mutant, defined as the value such that if the fitness of a resident genotype is normalised to 1, the fitness of the invader is $1+s$. Under the Strong Selection Weak Mutation (SSWM) assumptions, $Nu \ll \frac{1}{\log(Ns/2)}$, introduced in Section 2.1.2, the population can be assumed to be isogenic (each individual has the same genotype) and thus we may consider the *population genotype* [116, 117]. Under SSWM dynamics, new mutants of this population genotype that arise in the population will either become extinct or reach fixation before a second mutant arises in the population. Thus, we may ignore the dynamics of mutation and the stochastic nature of fixation in favour of simplified evolutionary dynamics: the periodic replacement of the population genotype by

a fitter genotype chosen stochastically from the mutational neighbours of that genotype. In fact, this model can be applied when the SSWM assumptions are relaxed, for example when $Nu \gg 1 \gg Nu^2$ [71], but breaks down owing to stochastic tunelling (the crossing of fitness valleys) when $Nu^2 \approx 1$ [241, 184, 163].

Following previous studies of the fitness landscape model (Section 2.2.2), we assume that individuals within the population take genotypes from a (hyper–)cubic genotype–space, $G = \{0,1\}^L$. Recall that a single genotype in this space corresponds to a bit string and mutations are modelled as the flipping of a single bit in this string. The collection of mutants of a single genotype is then the size L set of all genotypes that exist at a Hamming distance of one from that genotype. We assume that the fitness landscape defines a functional fitness mapping in which each genotype determines a single phenotype which has a singular real–valued fitness in the environment. Thus, we consider fitness as directly assigned to genotypes through a fitness function $f : G \to \mathbb{R}$ and assume that a different fitness function exists for each different environment. A schematic representation of this model is shown in Figure 2.3. Thus, we study the evolutionary dynamics of a population with SSWM dynamics and with GP–mapping defined by the fitness landscape model.

We begin by presenting a useful simplification of the evolutionary dynamics of this population. First note that there exist only a finite number of genotypes (2^L) in the fitness landscape model. Further, under SSWM dynamics a singular population genotype can be modelled as periodically replaced by a fitter mutational neighbour. This assumption is justified because deleterious and neutral mutations can be shown to fix with such low probability that we may assume that they never achieve fixation [116, 117]. It follows that the evolutionary process must terminate at a genotype that has no fitter mutational neighbours. Such a genotype is known as a *local optimum of fitness*. Under the SSWM model, the population can be viewed as occupying a single vertex on a directed graph over the set, G, of 2^L genotypes in which there exists an edge from vertex $g_1 \in G$ to a neighbouring vertex $g_2 \in G$ if, and only if, $f(g_2) > f(g_2)$ (see Figure 3.1). The population can then be considered to follow a stochastic walk on the graph in which the population genotype is replaced by a fitter adjacent genotype with some probability. Owing to the existence of local optima, this stochastic walk is necessarily an absorbing one.

This model of evolution as a stochastic walk requires us to determine a rule for which of the fitter mutational neighbours of the present population genotype will be stochastically chosen to replace it. Several such 'move rules' have been proposed [248] and which of these rules is most accurate depends on the relationship between the population size and mutation rate [71]. Common move rules include selecting the fittest neighbour [172, 94], selecting amongst fitter neighbours at random [214, 213, 90] or selecting from fitter neighbours with probability proportional to the fitness increase conferred [116–118]. We encapsulate each of these variants of the SSWM model within our model.

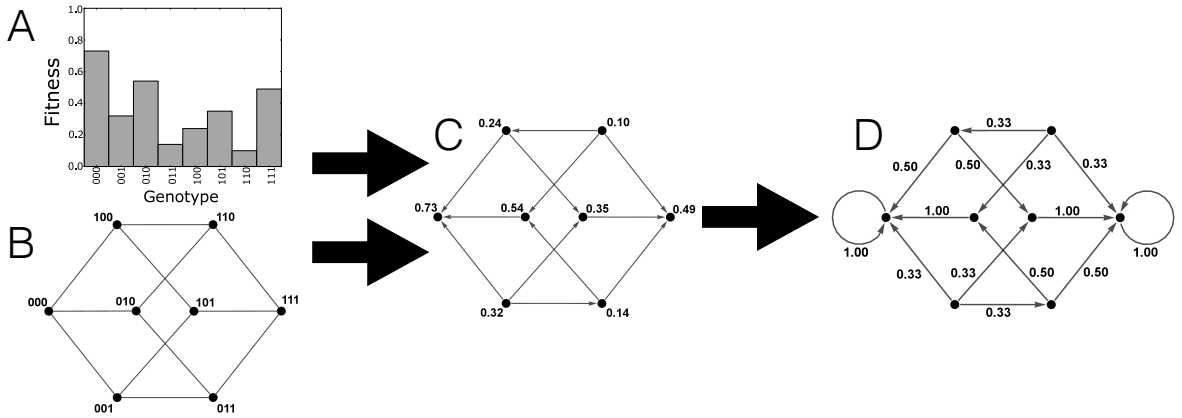

Fig. 3.1 **Constructing the Markov chain model of evolution A)** An example fitness landscape. **B)** The space of genotypes comprising bit strings of length $L = 3$. The vertices represent genotypes, and edges connect those genotypes which are mutational neighbours. **C)** The directed evolutionary graph according to the landscape in (A) where the vertices still represent genotypes but are labeled by the associated fitness (from (A)) for clarity. The directed graph edges are determined by the fitness function and represent those mutations which can fix in a population (those which confer a fitness increase). (D) The Markov chain constructed for the same landscape according to Equations 3.3 and 3.4 with $r = 0$.

3.1.1 Encoding Evolution as a Random Walk

As the SSWM assumptions reduce the evolutionary process to a random walk on a directed graph, the evolutionary process can be modelled as a Markov chain. Assume, without loss of generality, that fitness is defined by a positive value and consider a fitness landscape defined by $f : \{0,1\}^L \to \mathbb{R}^{\geq 0}$. Consider a population with genotype $i \in \{0,1\}^L$. To encode the random walk described above as a Markov chain we must determine for each $j \in \{0,1\}^L$, the probability that j replaces i as the population genotype. Denote this probability by $\mathbb{P}(i \to j)$.

First, we note that

$$\mathbb{P}(i \to j) = 0 \text{ if } \mathrm{Ham}(i,j) \neq 1 \tag{3.1}$$

owing to the definition of mutation in the genotype space $\mathscr{G}$ (here, Ham denotes the Hamming distance between two binary string genotypes, i and j). Further, owing to the Strong Selection assumption,

$$\mathbb{P}(i \to j) = 0 \text{ if } f(j) \leq f(i). \tag{3.2}$$

Where j is a fitter mutational neighbour of i, the probability of j fixing as the next mutational neighbour is dependent on the number of other fitter mutational neighbours, the fitness benefit associated with that mutation and the strength of selection. Rather than make assumptions

regarding the strength of selection, we assume a general form for the probability,

$$\mathbb{P}(i \to j) = \begin{cases} \dfrac{\Big(f(j)-f(i)\Big)^r}{\displaystyle\sum_{\substack{g\in\{0,1\}^L,\ \mathrm{Ham}(i,g)=1 \\ f(g)-f(i)>0}} \Big(f(g)-f(i)\Big)^r} & \text{if } f(j) > f(i) \text{ and } \mathrm{Ham}(i,j) = 1 \\ 0 & \text{otherwise} \end{cases} \tag{3.3}$$

where the parameter $r \geq 0$ determines the extent to which the fitness benefit of a mutation affects its likelihood of being the next population genotype. In the case $r = 0$, each fitter mutation is equally likely to be the next population genotype. This is equivalent to the assumption that the first fitness–confering mutant to arise will be the next population genotype. Thus, taking $r = 0$ corresponds to the random move SSWM model as taken by Macken and Perelson [214], Macken et al. [213] and Flyvbjerg and Lautrup [90]. If instead we consider the limit $r \to \infty$, then only the fittest mutational neighbour can fix and the random walk becomes the steepest gradient ascent SSWM model considered by Kauffman and Levin [172] and Fontana et al. [94]. For intemediate r, the likelihood of a fitter mutational neighbour fixing as the next population genotype is biased by the fitness benefit conferred, with increasing r resulting in a stronger bias towards fitter neighbours. In the case $r = 1$, the probability is directly proportional to fitness increase and corresponds to the model presented by Gillespie [116, 117].

Equation 3.3 alone does not define a Markov chain, as for a local optimum of fitness genotype, i, we have $\mathbb{P}(i \to j) = 0$ for all $j \in \{0,1\}^L$. As such, for local optima of fitness we further define

$$\mathbb{P}(i \to i) = \begin{cases} 1 & \text{if } i \text{ has no fitter one-step mutational neighbours} \\ 0 & \text{otherwise} \end{cases}. \tag{3.4}$$

From Equations 3.3 and 3.4, we can define a time–homogeneous absorbing Markov chain encoding the evolutionary process as

$$P = [\mathbb{P}(i \to j)]_{i,j\in\{0,1\}^L}. \tag{3.5}$$

A schematic representation of the construction of this Markov chain is shown in Figure 3.1. Note that in our construction we ensure that neutral mutations (i.e. those for which $f(i) = f(j)$) can never fix. Thus, the model derived here differs from the Markov chain model used by Sella and Hirsh [295] to study the neutral theory of evolution.

Through the Markov chain construction we can explore the accessibility and likelihood of evolutionary trajectories to a target population genotype under different fitness landscapes f. Define a collection of population row vectors $u^{(t)}$ for ($t \in \mathbb{N}$) wherein $u^{(t)}$ is a length 2^L vector with k^{th} component denoting the probability that the population has the k^{th} genotype at time t. Here, we consider the genotypes to be ordered numerically according to their binary

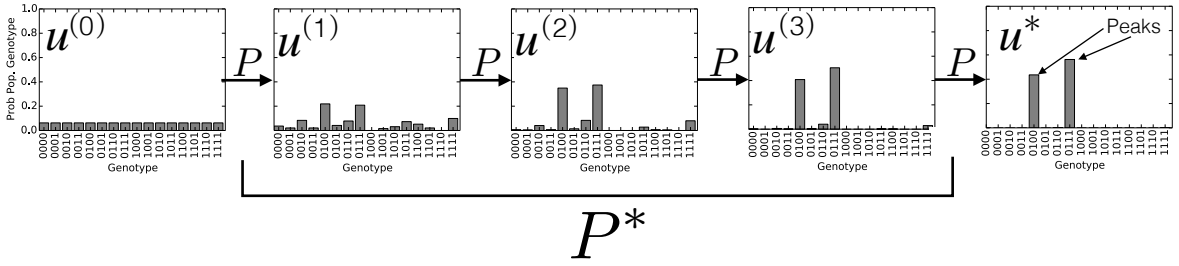

Fig. 3.2 **Population dynamics are defined by the matrix P.** Population distribution updates corresponding to the fixation of one beneficial mutation can be computed by multiplication by P. After finitely many matrix multiplications a stationary distribution, u^*, is found. In this distribution only genotypes corresponding to local optima of fitness have non–zero probability.

value. The discrete time steps are taken as an abstraction measuring the event of a beneficial mutation occurring *and* fixing in the population. For this definition of time, a one time step update of the population distribution can be computed by matrix multiplication as

$$u^{(t)} = u^{(t+1)}P. \tag{3.6}$$

The actual time taken for the population to transition between model time steps t and $t+1$ is not constant. Instead, it may be considered drawn from a distribution parameterised by the mutation rate, reproductive rate and the number of beneficial mutations that can occur. This distribution will be dependent on the precise population dynamics and could be estimated by simulating an idealised population according to classical models such as a Moran [230] or Wright–Fisher [359, 89] process. If the population has a genotype corresponding to a local optimum of the fitness landscape at time t, then there are no beneficial mutations that can occur, and this definition of a time step is not well defined. In these cases there can be no further changes to the population genotype under the SSWM assumptions. Thus, for mathematical convenience we can choose timestep t to $t+1$ to take some fixed arbitrary time in this case.

The distribution of a population at time t is related to an initial distribution, $u^{(0)}$, by

$$u^{(t)} = u^{(0)}P^t. \tag{3.7}$$

As we know that the Markov chain is absorbing, there necessarily exists some k for which $P^kP = P^k$ [131]. Consequently, we know that the matrix

$$P^* = \lim_{t\to\infty} P^t \tag{3.8}$$

exists. This limit is necessarily reached after only finitely many matrix multiplications. To intuitively see that this limit is reached for finite t, note that all paths through the Markov chain are strictly increasing in fitness and that there are only finitely many states (corresponding to the 2^L genotypes). It follows that a given initial population distribution $u^{(0)}$ will converge to a stationary distribution u^* after a finite number of steps in our model. Furthermore, if P^* is

known then we can compute this stationary distribution u^* according to

$$u^* = u^{(0)} P^*. \tag{3.9}$$

The matrix P can be estimated from empirically derived fitness landscapes, for example, those derived from the experiments described in Section 2.2.2. Here, we use the matrix P^* to predict the likelihood that a specific optimum of fitness in the landscape is found by an evolutionary trajectory from a starting genotype. We note that Markov chains are memoryless, and so under the model outlined here, evolutionary predictions can be made from any measured genotype without the need to know the evolutionary history. In predicting the evolution of diseases such as bacterial infections or cancers this property is a useful one, if it indeed it holds in biological reality. This is because genomic profiling will only be performed after the disease becomes symptomatic, which occurs some time after infection (for bacteria or viruses) or malignant transformation (for cancers), and may have evolved in the time prior to diagnosis.

The simplified discrete time steps in this model allow us to consider evolution in broad parameter ranges of population size and mutation rate without the need to know these parameters precisely. Such a simplification is particularly useful in modelling the evolution of drug resistance, as drugs induce large changes in population size and may, particularly in the case of cancer therapies, be mutagenic themselves [88, 282]. The drawback of this approach is that predicting the population genotype, or even the number of mutants that have arisen and fixated, after a specific time interval is not possible. However, if we assume that each drug is applied for sufficiently long to ensure that the disease population genotype reaches evolutionary equilibrium (a local optimum genotype), then we can explore the effects of applying multiple drugs sequentially by multiplication of the matrices P^*. The assumption of evolutionary convergence under each drug appears to be a strong one, particularly as evolution is never truly 'converged'. However, here we are considering a restricted subset of common fitness–conferring mutations for which studies of empirical landscapes indicate that the evolutionary trajectories are often short [248, 71]. This suggests that evolutionary convergence could happen rapidly.

3.1.2 Model Parametrisation

Recent empirical studies have derived fitness landscapes by engineering strains of *E. coli* with all combinations of a small number of mutations and then assaying the fitness in the presence of a number of antibiotic drugs. We can use these empirical landscapes to parametrise the model and to predict evolutionary trajectories under the associated drugs. Two different proxies for organismal fitness have been used previously in deriving empirical fitness landscapes. The first is the minimum inhibitory concentration (MIC) of a drug. This MIC is defined to be the smallest drug concentration that will inhibit the visible growth of an organism after overnight incubation [16]. This measure was used as a proxy for fitness in the studies performed by Weinreich et al. [346] and Tan et al. [318] to explore (respectively) the accessibility and reversibility of evolution. The second fitness measure, which was used by Mira et al. [228],

is to determine the average growth rate of the population under a specific concentration of the drug. The average growth rate is determined by applying linear regression to a sequence of optical density measures of the growing population [135]. Deriving a direct relationship between the MIC of a drug and the average growth rate under a given concentration is not possible. Ideally, a concentration–growth rate curve would be derived to provide both the MIC and growth rates under varying drug concentrations. However, deriving such curves introduces additional complexity into the experimental design and reduces the number of drugs that can be investigated. As we aim to explore drug sequences here, a dataset deriving single fitness values (i.e. one drug concentration) for a high number of drugs is preferable to a dataset deriving the concentration–growth curve for a smaller number of drugs.

Mira et al. [228] derived small combinatorically–complete subspaces of the fitness landscape of *E. coli* under the selective pressure of 15 β–lactam antibiotics, using the average growth rate as a proxy for fitness. Sixteen strains were engineered to account for all combinations of four resistance conferring mutations, M69L, E104K, G238S and N276D, in the gene TEM. This gene codes for enzymes known as beta-lactamases that provide protection from the action of beta-lactam antibiotics (a large class of broad spectrum antibiotics including penicillin derivatives, cephalosporins and others) by destroying the molecular structure of these drugs.

The engineered strains were then assayed to determine an average growth rate as a proxy for fitness in the presence of the different antibiotics. Details regarding the drugs and drug concentrations used in this analysis are shown in Table 3.1. Fitness value were determined as the average growth across 12 replicates. Figure 3.3 shows the fitness of each genotype under each of the 15 drugs with local optima genotypes highlighted in red.

We parametrised our model using these average growth rates as fitness values for the 15 antibiotic drugs. Mira et al. [228] did not report the variance of these growth rate estimates, but did report those adjacent genotypes with significantly different change in fitness, as determined by one-away ANOVA. However, in their analysis, Mira et al. [228] assumed that a mutation can achieve fixation in the SSWM model if it results in an increase in average growth rate, regardless of whether this increase was reported as significant. We make the same assumption in our analysis, as it is consistent with the assumption of strong selection, wherein slightly advantageous mutations (that are unlikely to be reported as significant in a small number of replicates) can achieve fixation. We note that in future our results should be corroborated by the analysis of more precisely measured fitness landscapes, but that such landscapes are presently not available.

Determining the matrix P from an empirical landscape requires us to estimate a value for the parameter r determining the extent to which the increase in fitness biases the likelihood of that mutation fixing. For the remainder of this work we take the assumption that $r = 0$, a value that corresponds to the definite fixation of the first beneficial mutant that randomly arises. We make this assumption as fitness values determined by average growth rate do not provide sufficient information about birth and death rates of individuals to derive invasion

Antibiotic	Abbreviation	Group	Concentration
Ampicillin	AMP	Aminopenicillin	2,048 μg/ml [227]
Amoxicillin	AM	Aminopenicillin	512 μg/ml
Cefaclor	CEC	Cephalosporin	1 μg/ml
Cefotaxime	CTX	Cephalosporin	0.05 μg/ml
Ceftizoxime	ZOX	Cephalosporin	0.03 μg/ml
Cefuroxime	CXM	Cephalosporin	1.5 μg/ml
Ceftriaxone	CRO	Cephalosporin	0.045 μg/ml
Amoxicillin + Clavulanic acid	AMC	Penicillin derivative + β-Lactamase inhibitor	512 μg/ml (Amoxicillin) and 8μg/ml (Clav)
Ceftazidime	CAZ	Cephalosporin	0.1 μg/ml
Cefotetan	CTT	Cephalosporin	0.312 μg/ml
Ampicillin + Sulbactam	SAM	Penicillin derivative + β-Lactamase inhibitor	8 μg/ml (Ampicillin) and 8μg/ml (Sulbactam)
Cefprozil	CPR	Cephalosporin	100 μg/ml
Cefpodoxime	CPD	Cephalosporin	2 μg/ml
Pipercillin + Tazobactam	TZP	Penicillin derivative + β-Lactamase inhibitor	12μg/ml (Pipercillin) and 8 μg/ml (Tazobactam) [227]
Cefepime	FEP	Cephalosporin	0.0156 μg/ml

Table 3.1 **15 β–lactam antibiotics for which Mira et al. [228] derived fitness landscapes.**

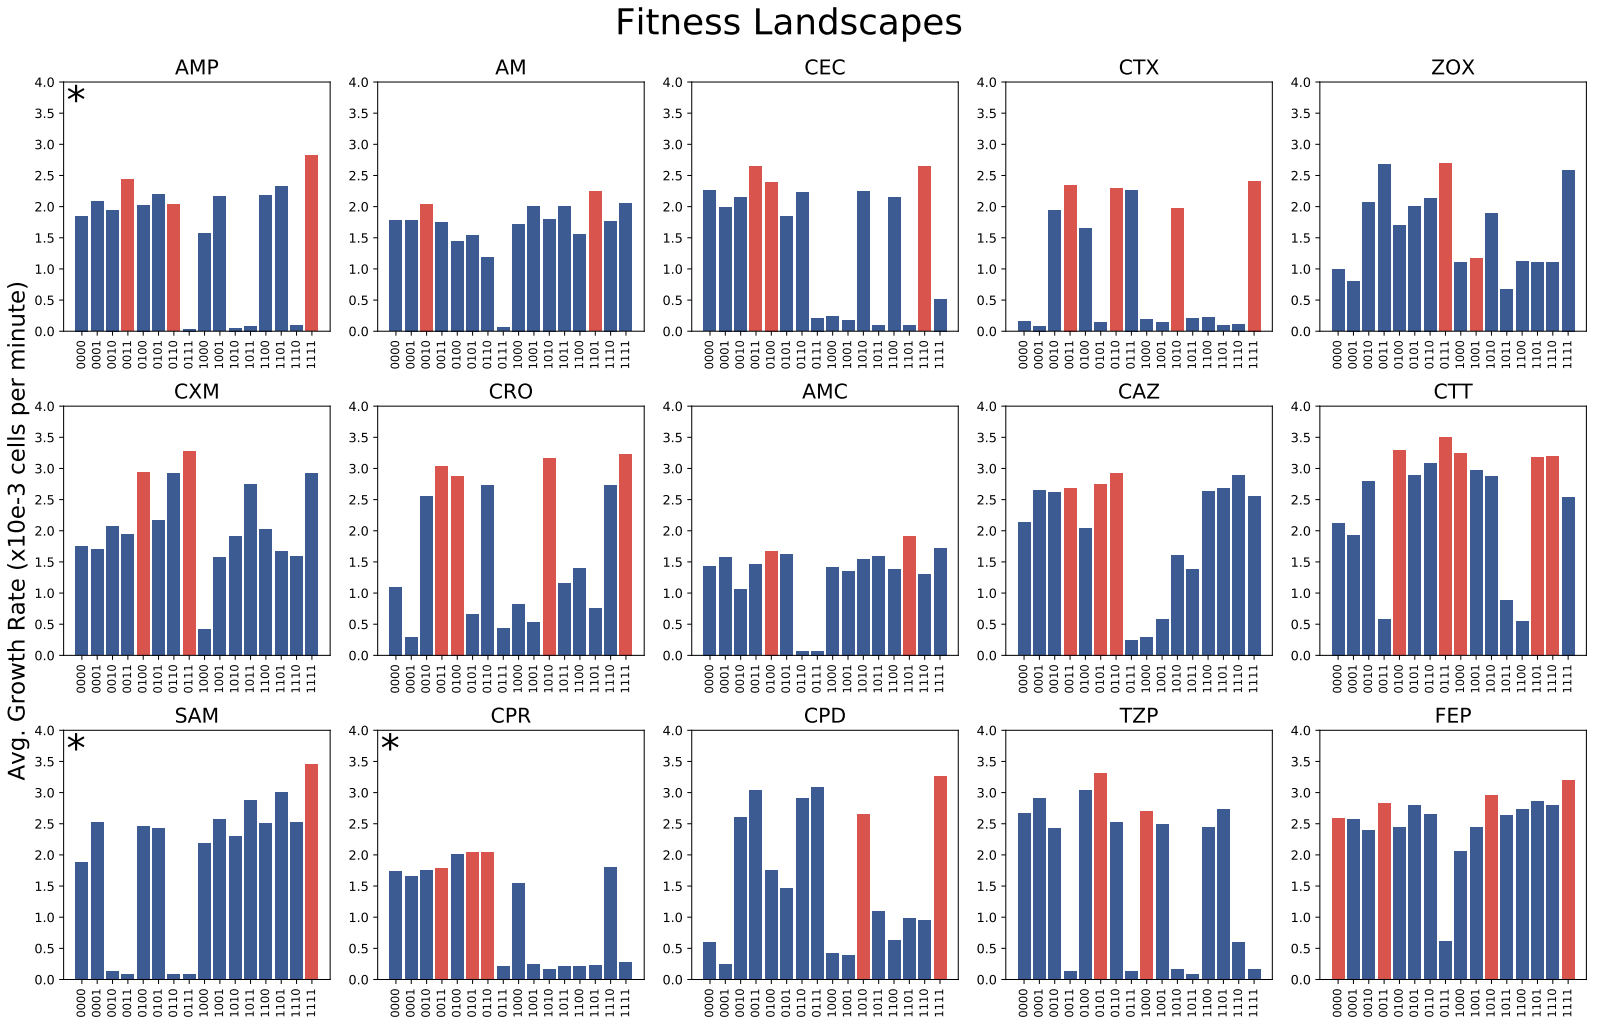

Fig. 3.3 **The 15 fitness landscapes from Mira et al. [228].** The fitness landscapes derived by Mira et al. [228] replotted as bar charts. Bars indicate the fitness (measured as an average per minute growth rate taken from 12 replicates) of each of the 16 genotypes under each of the drugs listed in Table 3.1. Red bars indicate local fitness optima. Note that adjacent bars in the graph need not be mutational neighbours. Inset stars indicate the three landscape graphs shown in Figure 3.10. Error bars are not shown as neither the variance nor raw data were provided by Mira et al. [228].

probabilities or the associated r values. Note that changing the value of r does not change the accessibility of an evolutionary trajectory. Thus, taking a different value of $r \geq 0$ will change the results quantitatively, the probabilities of trajectories may change, but not qualitatively, as accessibility (non-zero probability) is preserved. Where this quantitative change could impact our results we corroborate the results for two alternative commonly used r values: $r = 1$ and $r \to \infty$.

3.2 Simple Algebraic Results for Evolution

Before exploring the implications of the parametrised model for sequential drug therapies, it is informative to consider what we can infer about evolution in this model purely from its algebraic structure. Here, we present arguments for the existence of two clinically–important properties of evolution that are independent of the fitness landscapes. First, we demonstrate that the order in which drugs are applied has a significant impact on the accessibility of evolutionary trajectories. Second, we show that without knowledge of the population genotype, there is no way to ensure that an evolutionary trajectory taken under one drug will be reversed by a second. When evolution is considered in a theoretical and algorithmic framework, for example the framework introduced in Chapter 2, these properties of evolution seem obvious. However, they are not considered obvious for many physicians and play no role in clinical decision making. As such, a simple disease–independent and drug–independent derivation of the non–commutativity and irreversibility of evolution is valuable as a starting point from which to make the case for the clinical potential of sequential drug therapies.

3.2.1 Non–Commutativity of Evolution and the Importance of Drug Ordering

We can formally prove that, for a large class of fitness landscape pairs, there is non-commutativity in the evolutionary process when modelled by the Strong Selection Weak Mutation model. Suppose that there exist two drugs, X and Y, with corresponding fitness landscapes f_x and f_y, and that we wish to determine what, if any, difference there is between applying X followed by Y to a population as opposed to applying Y followed by X. We can construct the Markov chain transition matrices, P_x and P_y, corresponding to f_x and f_y according to Equation 3.3 and take the limits, P_x^* and P_y^*, of these matrices under repeated multiplication. Consider now an initial population with genotype $i \in \{0,1\}^L$. Denote by u_i the population vector with i^{th} component equal to one and all others equal to zero. The model predicts that drug ordering makes no difference to the final population distribution if, and only if,

$$u_i P_x^* P_y^* = u_i P_y^* P_x^*. \tag{3.10}$$

Suppose now that the initial population genotype is unknown. We can only guarantee that the order of drug prescription is irrelevant when the outcome is the same regardless of the starting genotype. Thus, drug ordering is irrelevant only when

$$u_i P_x^* P_y^* = u_i P_y^* P_x^* \qquad \text{for each genotype } i\,. \tag{3.11}$$

As the unit vectors u_i form a basis of $\mathbb{R}^L$, this can occur if and only if $P_x^* P_y^* = P_y^* P_x^*$. It follows that drug application will commute precisely when the corresponding limit matrices commute. As we know that matrix commutativity is not guaranteed, it follows that there is no reason to assume that the order in which drugs are prescribed is irrelevant.

Amongst the 15 drug landscapes, we found no commuting pairs under the SSWM model. We next tested 10^6 pairs of random fitness landscapes with varying ruggedness generated according to Kauffman's LK model for generating "tunably rugged" fitness landscapes [174, 172]. The landscapes were generated using a random neighbourhood Boolean function for determining the fitness contributions of each locus. We fixed $L = 5$ and generated each landscape by first drawing K uniformly from $\{0,\ldots,L-1\}$ and then using Kauffman's model to build a landscape. We found that 0.132% of the landscape pairs generated had limit matrices which commuted, suggesting that commutativity is rare and that the order in which drugs are prescribed will be have some effect in almost all instances.

We note that the SSWM model is a strong simplification of the evolutionary process. However, we believe this serves to strengthen the argument for non–commutativity of evolution – given that non-commutativity is present in this highly simplified model, it is unlikely that commutativity will emerge as more of the complexity present in the evolutionary process is introduced.

3.2.2 Irreversibility and Cycling of Evolutionary Trajectories

Evolutionary cycling, wherein a sequence of drugs is given that drive evolution through a 'cycle' in genotype space back to the original genotype, has been suggested as a potential strategy to maintain the efficacy of antibiotics [30, 255, 34, 368]. Cycling strategies were studied using a collection of empirically derived landscapes by Goulart et al [122]. In this work, a depth first search of the (directed) evolutionary graphs (similar to those shown in Figure 3.3) was used to identify sequences of drugs for which the evolutionary trajectory is guaranteed to return to the starting genotype.

The Markov chain encoding of the evolutionary trajectories simplifies this search. Suppose that the initial genotype is known to be $i \in G$. The initial population distribution will be u_i and a sequence of drugs $X_1,\ldots X_k$, with fitness landscapes $f_{x_1},\ldots f_{x_k}$, will constitute a cycling strategy (driving the trajectory through a cycle in G that returns to genotype i) precisely when

$$u_i P_{x_1}^* \ldots P_{x_k}^* = u_i. \tag{3.12}$$

This criterion will be satisfied precisely when u_i is a left 1–eigenvector of the product $P^*_{x_1} \dots P^*_{x_k}$. Thus, we can find cycling strategies amongst a collection of drugs efficiently by direct calculation and avoid the graph–search technique used by Goulart et al [122].

However, our model also demonstrates that evolution is, without knowledge of the initial genotype, irreversible. This finding has important consequences for evolutionary cycling. Consider a fitness landscape, f_x, corresponding to a drug, X. Unless f_x is an entirely flat landscape (taking equal values for all genotypes) there must exist at least one genotype $j \in G$ whose fitness is a minimum and which has a fitter neighbour. Such a genotype satisfies $\mathbb{P}(i \to j) = 0$ for all genotypes $i \in G$. Hence if f_x is not flat, the limit matrix P^*_x has at least one column of all zeros and is singular. Thus, there cannot exist a second landscape f_y (i.e. a drug Y) for which $P^*_x P^*_y = I$. For any second landscape, f_y, there must exist a unit row vector, u_i, for which $u_i P^*_x P^*_y \neq u_i$. Hence, evolution in our model is irreversible in the sense that for a given (non-flat) landscape we cannot find a second that is guaranteed to reverse its effects, unless we first measure the population genotype. This result precludes the existence of a *general* cycling strategy in which a sequence of drugs is guaranteed to return the disease population to its starting genotype regardless of what genotype that is. Further, if drugs are prescribed in an order that promotes resistance, there can be no guaranteed way to reverse the emergence of drug resistance without prior genomic sequencing. Worse still, evolutionary trajectories may be irreversible even when the population genotype is known [318].

3.3 Implications for Collateral Sensitivity Analysis

Recent empricial approaches have attempted to determine effective sequential drug therapies through *in vitro* experimental evolution [197, 161, 196, 233, 275]. In these experiments bacteria are cultured under the selective pressure of a specific drug to promote the evolution of resistance. After a set time, and under the assumption that drug resistance has evolved, the sensitivity of that population is then assayed against a panel of alternative drugs. The sensitivity of the evolved population to these alternative drugs is compared to the sensitivity of the wild–type population to determine a *collateral response*. Where the evolved population is more resistant to the second drug than the wild–type population, the first drug is said to induce *collateral resistance* (or cross–resistance) in the second. Where the evolved population is more susceptible, *collateral sensitivity* has occurred. By performing the experimental evolution for each drug and then assaying the resistance to each of the others, a matrix of collateral responses is derived. From this matrix, sequences of drugs in which each induces collateral sensitivity to the next can be identified. Identifying such sequences has been suggested as a means to determine effective therapeutic strategies that mitigate drug resistance. Here, we use the fitness landscape model to explore how the existence of multiple accessible trajectories in the fitness landscape impacts on the efficacy of collateral sensitivity analysis.

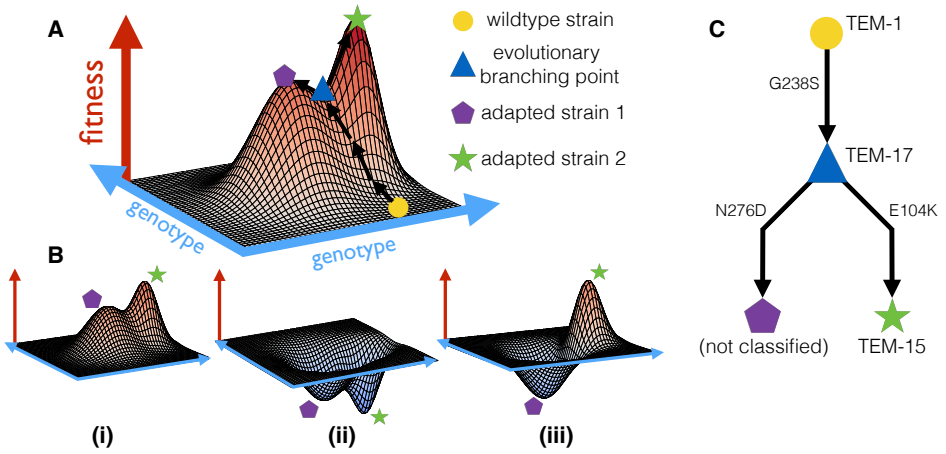

Fig. 3.4 **Collateral sensitivity need not be repeatable. A)** Evolution from a wild–type genotype (yellow circle) need not be repeatable owing to the potential of branching points in the fitness landscape (blue triangle). **B)** Patterns of collateral response that can arise from evolutionary divergence: **i)** guaranteed cross–resistance, **ii)** guaranteed collateral sensitivity, **iii)** collateral response determined by which evolutionary trajectory occurs. **C)** An example evolutionary branching point in the TEM gene under the antibiotic cefotaxime derived from the fitness landcape dataset of Mira et al. [228].

3.3.1 Collateral Response is not Necessarily Repeatable

We begin by noting that 14/15 of the landscapes shown in Figure 3.3, all except the landscape for Ampicillin/Sulbactam (SAM), are multi–peaked, and for 12/15 more than one peak is accessible by a trajectory of fitness conferring–mutations from the wild–type genotype $g = 0000$. It follows that an evolutionary trajectory to resistance is not necessarily repeatable, as there can exist genotypes from which the trajectory can branch stochastically. Where collateral response is determined from a small number of experimental replicates, rare trajectories may be missed. Further, it is possible that the different local optima genotypes of one landscape have significantly different fitnesses under a second. In particular, one peak genotype may be sensitive to a second drug whilst another is resistant. It follows that collateral resistance could stochastically occur on some occasions where collateral sensitivity has been reported by experimental evolution. A schematic of this hypothesis is shown in Figure 3.4.

The impact of this non–repeatability is dependent on how the similarity of fitness between genotypes in one landscape correlates with similarity of fitness in others. If the correlation is high then we can expect genotypes with high fitness in one landscape to have similar fitness in a second. In this case, the existence of multiple accessible fitness optima in one landscape will have a smaller effect on the validity of the predictions made by experimental evolution – as collateral sensitivity shown by one peak genotype means other peak genotypes will show collateral sensitivity with higher probability. However, there is no *a priori* reason to assume any correlation in fitness between genotypes, beyond the observation that many drugs act on similar targets (cell wall, DNA damage etc).

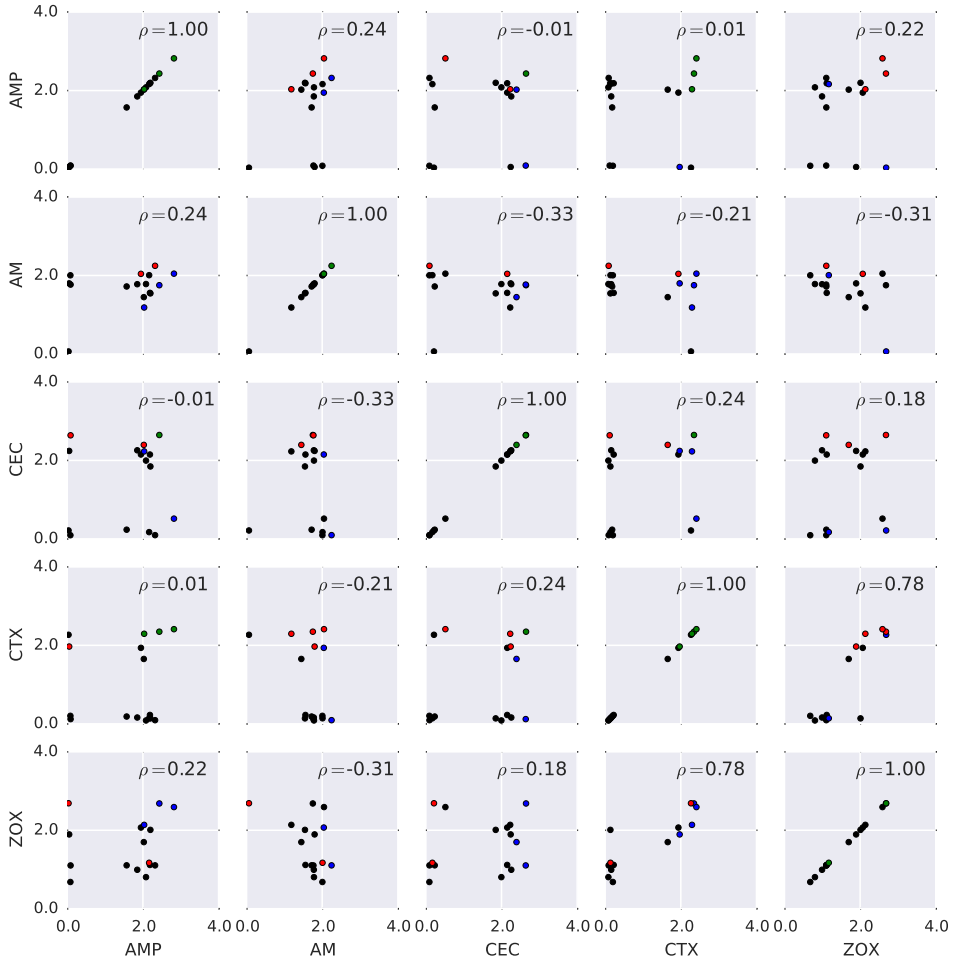

Fig. 3.5 **Correlation of the fitness of genotypes between two landscapes.** Each point represents a genotype from $\{0,1\}^4$ plotted according to its fitness in the two landscapes (axis labels). Blue points are local optima of the landscape labelled on the abscissa, red points are local optima of the landscape labelled on the ordinate and green points are local optima of both landscapes. The Spearman correlation coefficient, ρ, is given for each pair of landscapes.

For each pair of antibiotics listed in Table 3.1 we compared the fitness of each of the 16 genotypes between the two landscapes in a scatter plot. Figure 3.5 shows, for 25 of the pairs, examples of such plots. Scatter plots for the remaining 200 pairs are presented in Appendix A. We recorded the Spearman correlation coefficient between fitness in the two landscapes in Table 3.2. From this table it is clear that an assumption of correlation of fitness between landscapes is not well founded although drugs within a class, for example the cephalosporins, show more correlation than those between groups. An ideal pair of drugs for use in an alternating fashion would have a high negative correlation such that the evolution of resistance under one drug would induce sensitivity to a second. Unfortunately, no such drug pair exists and we must employ more sophisticated methods to identify viable sequential drug strategies.

	AMP	AM	CEC	CTX	ZOX	CXM	CRO	AMC	CAZ	CTT	SAM	CPR	CPD	TZP	FEP
AMP	1.00	0.24	-0.01	0.01	0.22	-0.02	0.09	0.46	0.50	-0.36	0.31	0.31	0.05	0.26	0.49
AM	0.24	1.00	-0.33	-0.21	-0.31	-0.38	0.09	0.39	-0.07	-0.33	0.63	-0.35	-0.16	-0.18	0.35
CEC	-0.01	-0.33	1.00	0.24	0.18	0.03	0.59	-0.17	0.45	-0.11	-0.48	0.54	0.18	0.01	0.16
CTX	0.01	-0.21	0.24	1.00	0.78	0.64	0.64	-0.20	-0.14	-0.09	-0.45	-0.08	0.84	-0.59	0.11
ZOX	0.22	-0.31	0.18	0.78	1.00	0.55	0.33	-0.25	0.04	0.19	-0.51	0.13	0.82	-0.28	0.08
CXM	-0.02	-0.38	0.03	0.64	0.55	1.00	0.27	0.04	-0.03	0.03	-0.22	0.07	0.75	-0.24	-0.03
CRO	0.09	0.09	0.59	0.64	0.33	0.27	1.00	0.15	0.22	-0.16	-0.07	0.15	0.57	-0.40	0.52
AMC	0.46	0.39	-0.17	-0.20	-0.25	0.04	0.15	1.00	0.07	-0.19	0.63	-0.02	-0.01	0.29	0.52
CAZ	0.50	-0.07	0.45	-0.14	0.04	-0.03	0.22	0.07	1.00	-0.21	-0.06	0.65	0.09	0.23	0.53
CTT	-0.36	-0.33	-0.11	-0.09	0.19	0.03	-0.16	-0.19	-0.21	1.00	-0.15	0.14	0.09	0.30	-0.38
SAM	0.31	0.63	-0.48	-0.45	0.51	-0.22	-0.07	0.63	-0.06	-0.15	1.00	-0.31	-0.34	0.07	0.35
CPR	0.31	-0.35	0.54	-0.08	0.13	0.07	0.15	-0.02	0.65	0.14	-0.31	1.00	0.02	0.48	-0.02
CPD	0.05	-0.16	0.18	0.84	0.82	0.75	0.57	-0.01	0.09	0.09	-0.34	0.02	1.00	-0.50	0.30
TZP	0.26	-0.18	0.01	-0.59	-0.28	-0.24	-0.40	0.29	0.23	0.30	0.07	0.48	-0.50	1.00	-0.17
FEP	0.49	0.35	0.16	0.11	0.08	-0.03	0.52	0.52	0.53	-0.38	0.35	-0.02	0.30	-0.17	1.00

Table 3.2 **Spearman correlation coefficients of fitness values between landscapes.** The scatter plots for the shaded region are shown in Figure 3.5. The remainder of the scatter plots are presented in Appendix A. Note that there exists no pair of drugs with large negative correlation. Such a pair would represent an ideal drug pair for an alternating/cyclic strategy. Instead, we must rely on longer sequences of drugs to minimise the risk of a highly resistant genotype arising.

We next performed an *in silico* derivation of tables of collateral response, or *collateral sensitivity matrices* (CSMs), by simulation of evolution with the model described by Equation 3.3. These simulations mirror the experimental techniques used to derive empirical collateral response, for example those used by Imammovic and Sommer [161] to determine drug cycling protocols (Figure 3.6). Evolutionary trajectories in each drug fitness landscape, f_x, were stochastically simulated from the wild–type starting genotype ($g_0 = 0000$) by sampling the associated Markov chain defined by P_x (Equation 3.3). The simulation was terminated when the evolutionary trajectory encountered a local optimum genotype, g_x^*. The fitness of this final genotype in the second drug landscape, f_y, was then recorded and collateral response was calculated as

$$\text{Collateral response of } Y \text{ to } X = \log_2\left(\frac{f_y(g_x^*)}{f_y(g_0)}\right)$$

These collateral response measures were collated to form a table.

We can count the total number of CSMs that can be generated through this simulation. There exist 3 landscapes with only one peak accessible from the genotype $g_0 = 0000$, 6 in which two peaks are accessible, 4 in which three peaks are accessible and 2 in which four peaks are accessible. Assuming that for each landscape, f_x, evolution is simulated from g_0 a single time to determine g_x^* and then the collateral sensitivity of g_x^* in each of the other landscapes, f_y is recorded. Then there exist

$$1^3 \times 2^6 \times 3^4 \times 4^2 = 82944$$

potential collateral sensitivity matrices that can occur. As we know the Markov chain used to perform this simulation, we can also determine the most likely of these CSMs by taking the

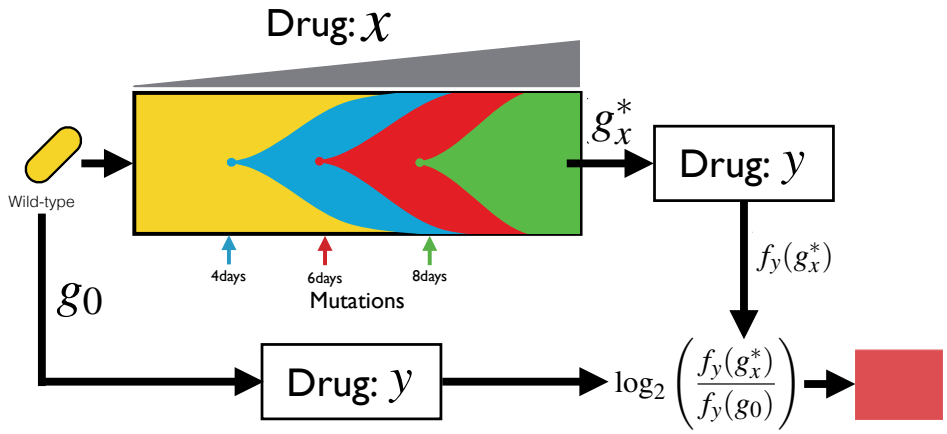

Fig. 3.6 **Schematic of the experimental derivation of collateral response.** To empirically derive a collateral response, a wild–type strain is cultured in increasing concentrations of drug, x. During this culture period, fitness conferring mutations arise and fix in the population. To determine the collateral response of a drug y to x, the wild–type sensitivity is compared to the sensitivity of the the evolved strain. To simulate this process *in silico*, we sample a trajectory from the Markov model for the drug x and compare the fitness values under the two drugs for the resulting genotype.

most likely g_x^* for the landscapes and recording the collateral response. This CSM is shown in Figure 3.7 and occurs with probability 0.0023 (assuming $r = 0$ in Equation 3.3). Given the strong underlying assumptions of the model it is unlikely that this probability is accurate, but the fact that it is small suggests that empirically derived CSMs need not be, and in fact likely are not, repeatable.

Suppose now that a new g_x^* is determined through simulation for each drug pair, rather than deriving a single g_x^* for each row of the CSM, and that collateral response is recorded as before. This simulation corresponds to an empirical approach that builds a CSM from a meta–analysis of other datasets or through clinical observations of multiple patients. In this case the total number of possible CSMs grows to approximately

$$\left(1^3 \times 2^6 \times 3^4 \times 4^2\right)^{15} \approx 6.05 \times 10^{73}.$$

Figure 3.7 shows the best, worst, most likely and average case instances for the CSMs. A comparison of the best and worst case CSMs highlights the pitfalls of *in vitro* collateral sensitivity analysis in designing sequential drug protocols — collateral sensitivity may be reported where collateral resistance can sometimes occur. Indeed, by considering all possible ordered drug pairs, we find that there exist 15 pairs in which the first has no effect on sensitivity to the second drug (this is because the landscape for Cefepime (FEP) has a local optimum of fitness at $g = 0000$), 28 pairs in which the first is guaranteed to induce sensitivity to the second, 94 in pairs in which the first in guaranteed to induce resistance and 88 pairs in which the first can induce both sensitivity and resistance depending on the stochastic outcome of the evolutionary trajectory under the first drug. This result undermines the use of empirically

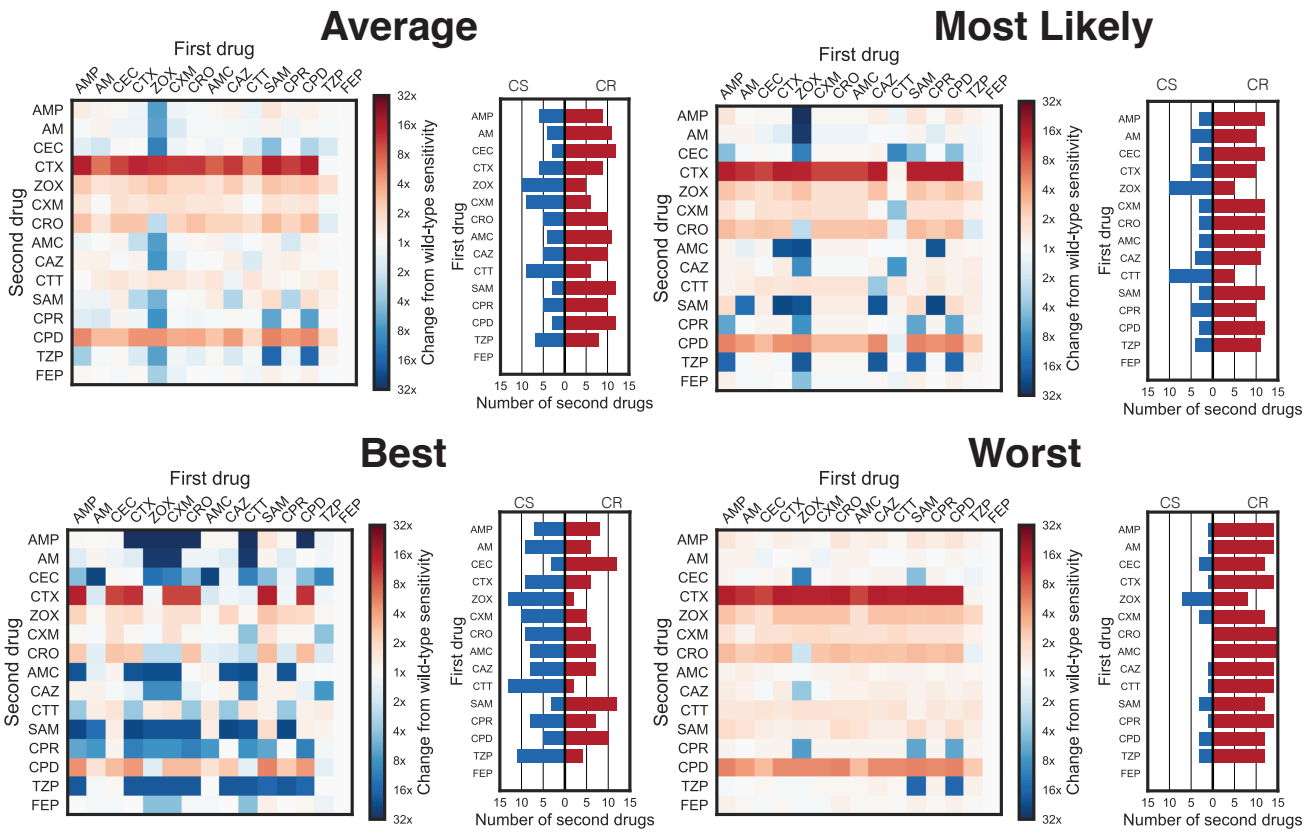

Fig. 3.7 **Analysis of the possible collateral sensitivity matrices (CSMs).** The best, worst, average case and most likely collateral sensitivity matrices are shown. The bar charts indicate the number of drugs that induce sensitivity (blue) or resistance (red) to the first–applied drugs listed on the y-axis.

determined collateral sensitivity in determining antibiotic prescription protocols, as evolution is not repeatable and thus the CSM reported from a small number of replicates need not be the only possible outcome. Indeed, amongst the 15 antibiotics we study here, an ordered pair taken at random is predicted to induce collateral resistance with probability 0.58. If we instead first take a CSM generated at random by simulating evolution on the fitness landscapes, and then choose a pair in which collateral sensitivity is reported by that experiment, then the probability of collateral resistance is 0.52. Thus, reporting a single CSM offers only a marginal improvement over choosing drugs at random.

3.3.2 Collateral Sensitivity Likelihood and Worst Case Response

Given that collateral response is not repeatable, we propose two alternative metrics to describe the efficacy of drug sequences. The first of these metrics is the *collateral sensitivity likelihood*, which we define as the probability that the first drug induces collateral sensitivity in the second. Figure 3.8 shows the matrices of collateral sensitivity likelihood for the 15 empirically derived drug landscapes. This figure is generated using the Markov chain model with $r = 0$. Unlike reporting a matrix of collateral response, the matrix of likelihoods explicitly highlights the risk that cross-resistance may occur. By allowing more risk (for example, $p > 0.75$ in Figure 3.8),

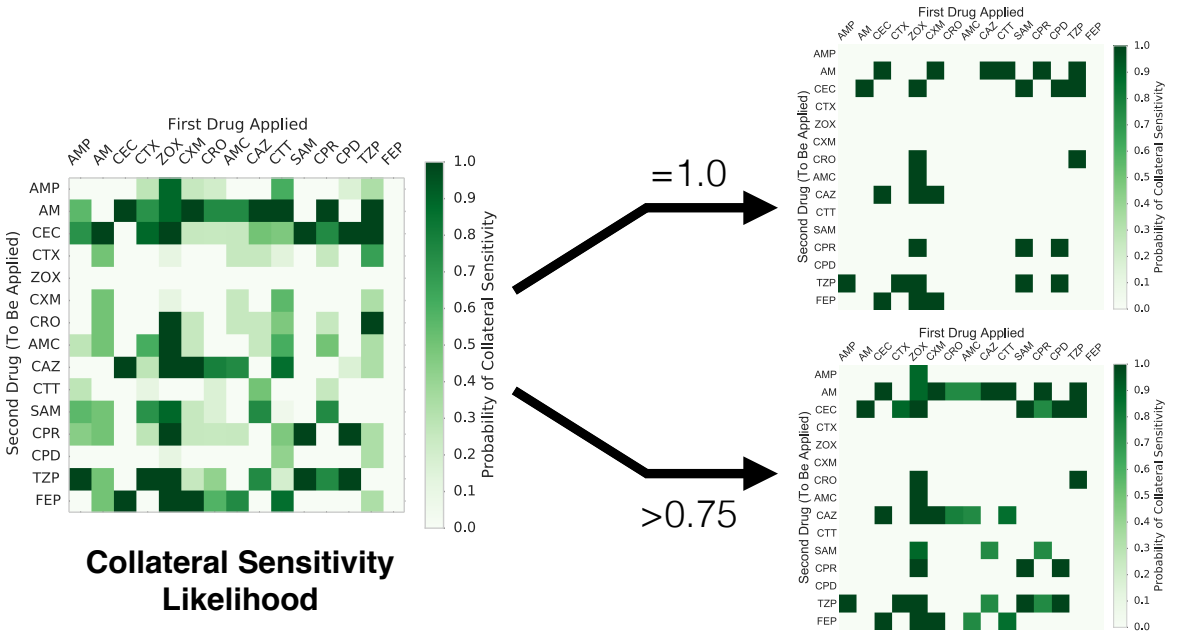

Fig. 3.8 **Collateral sensitivity likelihood.** (Left) The likelihood that a drug pair induces collateral sensitivity is shown for each drug landscape derived by Mira et al. [228]. Two additional matrices on the right show the CSL matrix thresholded for pairs with a likelihood, p, of inducing collateral sensitivity satisfing: $p = 1$, $p > 0.75$.

we see that a greater number of drug pairs become viable. The derivation of a collateral sensitivity likelihood matrix will likely require a large number of observations of evolution under a drug, and thus may be intractable through experimental evolution techniques. This problem may be mitigated by the collection of clinical observations from a number of clinics. The details of this approach are outlined in the discussion below.

The collateral sensitivity likelihood approach does not fully capture the risk associated with a drug pair. For example, a drug pair with an intermediate likelihood of collateral sensitivity, but only a small degree of cross–resistance where it does occur, is certainly lower risk than one with a high likelihood of collateral sensitivity but the potential to induce a highly multi–drug resistant phenotype where it fails. For this reason, we propose matrices of *worst case collateral response* as a companion metric to the collateral sensitivity likelihood. An example of such a matrix is shown in Figure 3.7(Labelled "Worst"). By thresholding this matrix, drug pairs can be identified with a limited worst case outcome (at least, amongst previously observed outcomes). Figure 3.9 shows thresholded matrices of worst case response derived from the 15 empirical fitness landscapes.

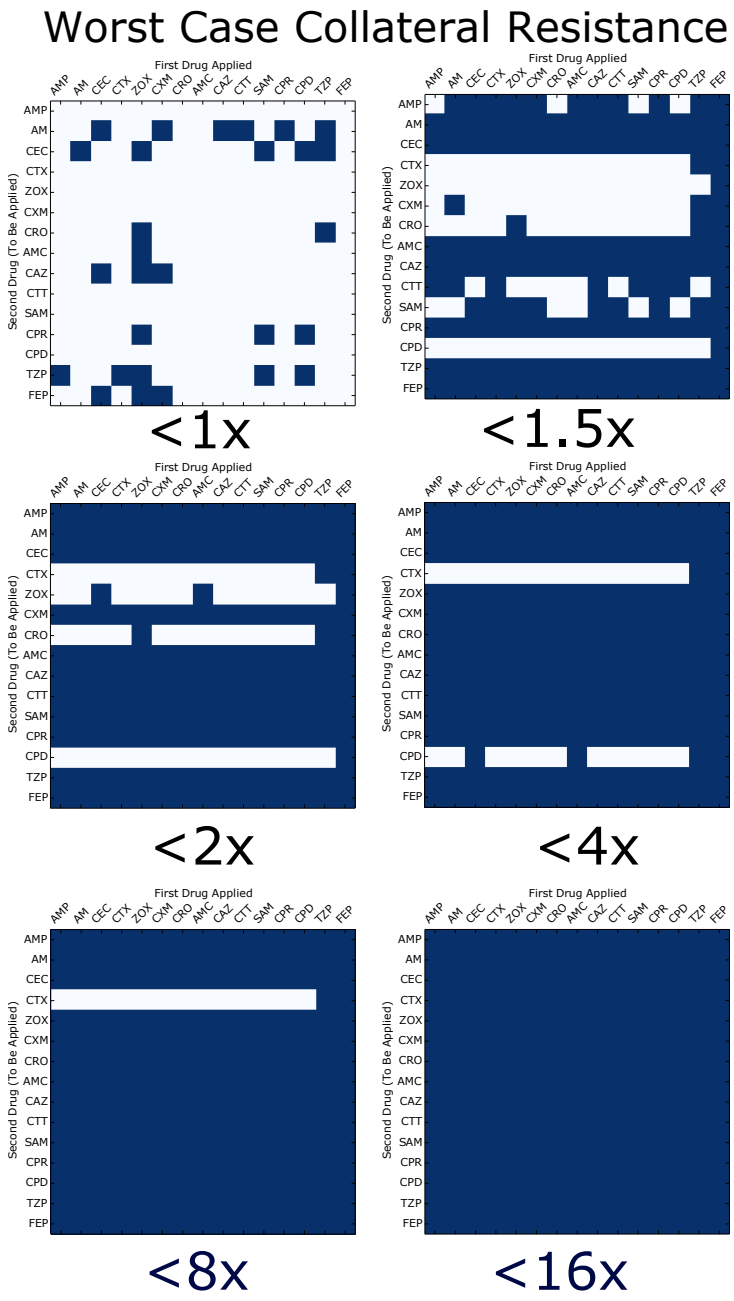

Fig. 3.9 **Thresholded matrices of worst case collateral response.** Blue squares indicate worst case collateral resistance lower than the stated threshold (1×, 2×, 4×, 8× and 16×). Note that the threshold of 1× corresponds precisely to the matrix of guaranteed collateral sensitivity (Figure 3.8).

3.4 Deriving Optimal Sequential Therapies from Empirical Landscapes

We have demonstrated that, owing to the existence of multiple accessible peaks in most fitness landscapes, pairs of drugs or sequences of drugs that ensure collateral sensitivity cannot be

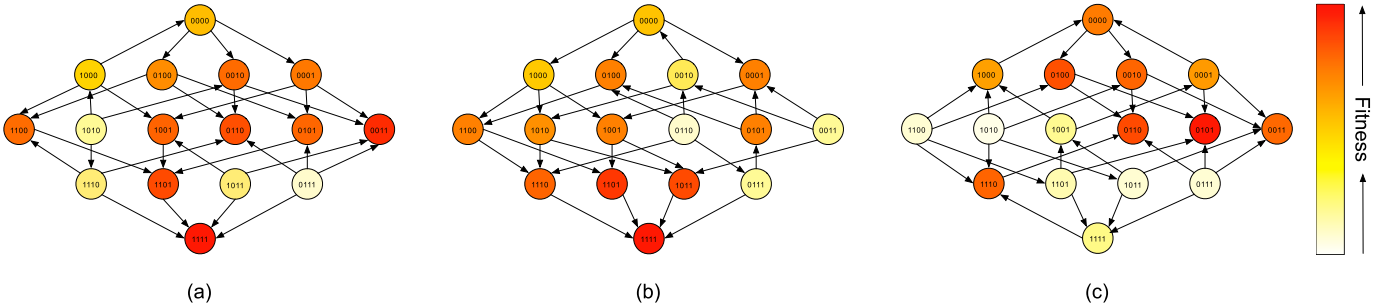

Fig. 3.10 **Graphs of fixable mutations in genotype space.** The graphs of fixable mutations (arrows) on the space of genotypes for (a) Ampicillin (AMP), (b) Ampicillin+Sulbactam (SAM) and (c) Cefprozil (CPR). The probability of fixation for these mutations is determined by the equations defined in Section 3.3.1 and are dependent on the value of r. Where $r = 0$, as in our analysis below, the probability of fixation for a specific mutation at a given genotype vertex is the reciprocal of the number of outbound edges for that vertex. For clarity, these probability values are omitted from the figure.

determined through *in vitro* experimental evolution alone: even where experiments indicate collateral sensitivity occurs, resistance may still be possible. We have suggested a partial solution to this problem: wide–scale data collection to indicate the likelihood of collateral resistance/sensitivity and the worst case outcomes of treatments. However, this data collection fails to account for two important factors of sequential drug therapies. The first is that collateral response analysis tells us nothing about how evolution will proceed once the second drug is applied. Thus, even where sensitivity is induced, there is no guarantee that the evolution of drug resistance to that second drug will not occur.

The second major problem is that it is not possible to predict the effects of longer sequences of drugs from analysis of the collateral response of the individual pairs. Specifically, consider three drugs X, Y and Z. It is not possible to determine the sensitivity to the drug Z after evolution first under X and then under Y from the collateral response pattern amongst each of the pairs. This is because evolution under X from an initial genotype g_0 yields new genotypes. The peaks of Y that are accessible from these new genotypes may differ from those accessible from g_0. Thus, the collateral response of Z to Y determined from experiments performed on an initial population with genotype g_0 may not be correct when evolution under X has taken place. This observation calls into question the second part of the methodology of Imamovic and Sommer [161], in which collateral sensitivity *networks* are constructed to identify longer drug sequences from collateral response amongst pairs.

In this section we introduce a framework through which we can use empirically derived fitness landscapes, coupled with our previous mathematical model, to overcome the above difficulties of collateral response analysis and to design sequential drug therapies. Specifically, we test the following hypothesis: As many evolutionary trajectories through a landscape are inaccessible [346] and evolutionary trajectories that occur in one landscapes are often irreversible in a second [318], can we irreversibly (as in [318]) *steer* evolution with one drug

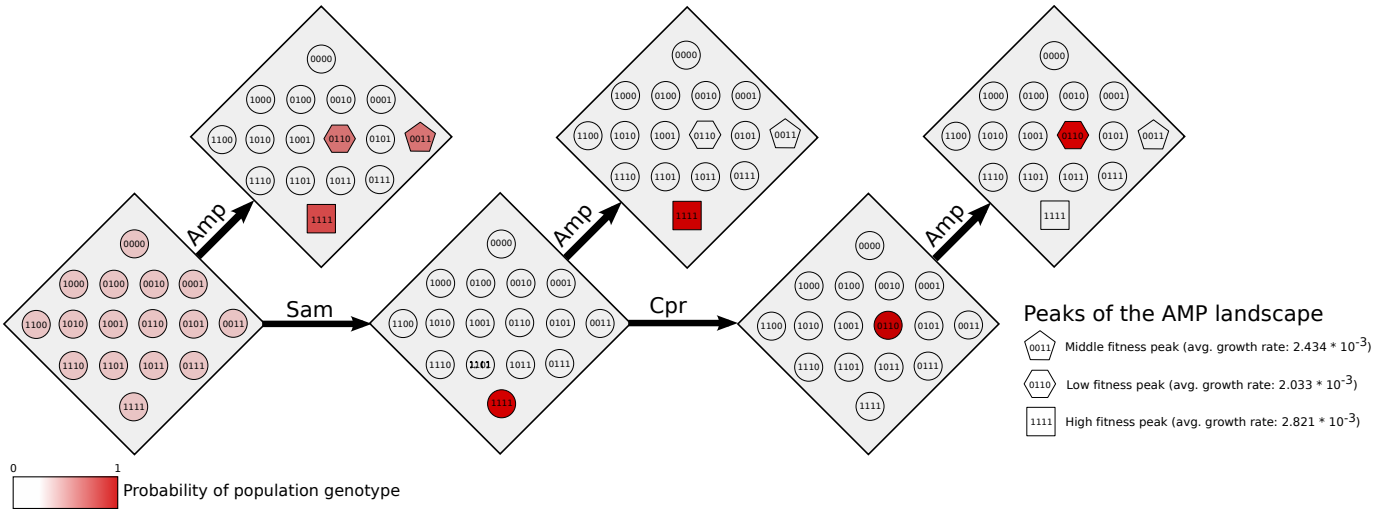

Fig. 3.11 **The probability distributions for accessibility of the peaks of the AMP landscape for different steering regimes.** The starting distribution is $u = [1/2^L, \ldots, 1/2^L]$. When AMP is given first any of the three peaks of the landscape are accessible with most resistant genotype 1111 being most likely. If SAM is given first to steer the population to its sole peak 1111, then resistance to AMP will be guaranteed when it is applied. Alternatively, if SAM is given followed by CPR then the population evolves to the local optima genotype 0110 of the CPR landscape. If AMP is applied to this primed population the global optimum is inaccessible as the population genotype, $g = 0110$, is already constrained to the lowest fitness local optimum genotype in the AMP landscape.

to a state from which the evolutionary trajectory to high resistance is inaccessible (as in [346]) for a second drug?

3.4.1 Evolutionary Steering to Avoid Resistance

We begin by presenting an example of this hypothesis in action — a sequence of drugs that *steers* evolutionary trajectories such that when a final drug is prescribed the trajectory to the highest fitness genotype is inaccessible. This example serves as proof of concept of how, if we have knowledge of the underlying fitness landscapes, we can design sequential drug therapies that satisfy our steering hypothesis.

Consider the landscapes for Ampicillin (AMP), Cefprozil (CPR) and Ampicillin+Sulbactam (SAM) presented in Figure 3.10. Suppose that we have an unknown disease population $u = [\frac{1}{2^4} \ldots \frac{1}{2^4}]$ as we might in a clinical setting where genomic sequencing of infections is, at present, not common practice. If we apply the drug AMP to this population distribution, we find that in the expected distribution $u^* = uP^*_{Amp}$ (shown in the first diamond in the top row of Figure 3.11) each of the fitness peaks could be found. In particular, our model predicts that the most highly resistant genotype 1111 will arise in the population with probability 0.62. Suppose instead we apply SAM first. In this case as the landscape is single peaked the population will converge to the global optimum genotype 1111 under our model. This genotype is also the global optimum of the AMP landscape and hence if we apply AMP after

Final drug	# of Peaks	Fittest peak genotype	Probability of ending at highest fitness genotype (no steering)	Best single steering drug	Probability of ending at highest fitness genotype (with best single steering drug)	Best ordered pair of steering drugs	Probability of ending at highest fitness genotype (with best steering pair)	Best ordered triple of steering drugs	Probability of ending at highest fitness genotype (with best steering triple)
Ampicillin (AMP)	3	1111	0.62	CPR	0.25	AMP→CPR	0.0	-	-
Amoxicillin (AMX)	2	1101	0.75	CTX	0.62	SAM→CPR	0.51	CRO→SAM→CPR	0.51
Cefaclor (CEC)	3	0011	0.19	SAM	0.0	-	-	-	-
Cefotaxime (CTX)	4	1111	0.17	CPR	0.04	AMP→CPR	0.0	-	-
Ceftizoxime (ZOX)	2	0111	0.86	AMX	0.81	SAM→AMX	0.75	CEC→CPR→AMC	0.75
Cefuroxime (CXM)	2	0111	0.56	AMC	0.33	CXM→AMC	0.19	AMC→CXM→AMC	0.11
Ceftriaxone (CRO)	4	1111	0.28	CEC	0.20	TZP→CEC	0.05	AMC→TZP→CEC	0.0
Amoxicillin +Clav (AMC)	2	1101	0.67	CXM	0.39	AMC→CXM	0.23	CXM→AMC→CXM	0.13
Ceftazidime (CAZ)	3	0110	0.39	TZP	0.08	AMC→TZP	0.0	-	-
Cefotetan (CTT)	5	0111	0.18	AMC	0.0	-	-	-	-
Ampicillin +Sulbactam (SAM)	1	1111	1.0	$\star$	1.0	$\star$	1.0	$\star$	1.0
Cefprozil (CPR)	3	0101	0.25	AMP	0.0	-	-	-	-
Cefpodoxime (CPD)	2	1111	0.81	CRO	0.74	CTX→FEP	0.57	SAM→CEC→CRO	0.50
Piperacillin +Tazobactam (TZP)	2	0101	0.80	CTT	0.72	CTT→CTT	0.72	CTX→FEP→CTT	0.68
Cefepime (FEP)	4	1111	0.48	CTX	0.39	CAZ→CEC	0.36	CAZ→CEC→CTX	0.33

Table 3.3 **Summary of evolutionary steering to avoid highly resistant genotypes.** For each of the 15 antibiotics we have derived the (ordered) sets of one, two and three steering drugs which reduce, to the lowest possible, the probability of evolution to the maximal fitness genotype. In the case that an ordered set of steering drugs reduced the probability to 0 we have not considered ordered sets of greater length (marked as - in the table). * – as the landscape for SAM is single peaked there can be no combination of steering drugs which reduce the probability of finding the global optimum. In all experiments the initial population distribution is taken as $u = [1/2^L, \ldots, 1/2^L]$ and $r = 0$.

SAM we will encounter high resistance. We have steered the population with one drug to a configuration which increases the probability of resistance to a second. Next suppose instead we apply CPR after SAM, in this case our model predicts that the population is guaranteed to evolve to a local optimum 0110 of the CPR landscape. 0110 is the least fit local optimum of the AMP landscape. Thus if we apply AMP to the population primed by SAM and CPR then evolution to the global optimum 1111 is not possible. This example demonstrates the steering hypothesis, that evolution can be shepherded through careful orderings of multiple drugs to increase or decrease the likelihood of resistance emerging (Figure 3.11).

To test how common such steering strategies could be, we performed an *in silico* test of steering using combinations of one, two or three preceding drugs taken from Figure 3.3. Table 3.3 shows for each of the 15 antibiotics the combinations which reduce the probability of evolution to a peak fitness genotype to the lowest possible when applied in order followed by the final drug. We found that for 3 of the 15 drugs there exists another which steers an initial population u to a configuration which prevents evolution to the global fitness optimum of the landscape entirely. This number rose to 6 when pairs of drugs applied sequentially are used to steer the population and to 7 when triples were applied in sequence.

Final drug	# of Peaks	Least fit peak genotype	Probability of ending at lowest fitness genotype (no steering)	Best single steering drug	Probability of ending at lowest fitness genotype (with best single steering drug)	Best ordered pair of steering drugs	Probability of ending at lowest fitness genotype (with best steering pair)	Best ordered triple of steering drugs	Probability of ending at lowest fitness genotype (with best steering triple)
Ampicillin (AMP)	3	0110	0.19	CPR	0.51	SAM→CPR	1.0	-	-
Amoxicillin (AMX)	2	0010	0.25	CTX	0.38	SAM→CPR	0.49	AMP→SAM→CPR	0.49
Cefaclor (CEC)	3	0100	0.41	TZP	0.70	CXM→AMC	0.78	AMC→CXM→AMC	0.87
Cefotaxime (CTX)	4	1010	0.21	FEP	0.27	CTX→FEP	0.43	SAM→CEC→CRO	0.50
Ceftizoxime (ZOX)	2	1001	0.14	AMX	0.19	SAM→AMX	0.25	AMP→SAM→AMX	0.25
Cefuroxime (CXM)	2	0100	0.44	AMC	0.67	CXM→AMC	0.81	AMC→CXM→AMC	0.89
Ceftriaxone (CRO)	4	0100	0.30	CXM	0.59	AMC→CXM	0.76	CXM→AMC→CXM	0.86
Amoxicillin +Clav (AMC)	2	0100	0.33	CXM	0.61	AMC→CXM	0.77	CXM→AMC→CXM	0.87
Ceftazidime (CAZ)	3	0011	0.26	FEP	0.28	CAZ→CEC	0.40	CAZ→AMX→FEP	0.45
Cefotetan (CTT)	5	1101	0.19	AMX	0.75	SAM→AMX	1.0	-	-
Ampicillin +Sulbactam (SAM)	1	1111	1.0	$\star$	1.0	$\star$	1.0	$\star$	1.0
Cefprozil (CPR)	3	0011	0.25	ZOX	0.26	CAZ→CEC	0.34	CAZ→CEC→FEP	0.43
Cefpodoxime (CPD)	2	1010	0.19	CRO	0.26	CTX→FEP	0.43	SAM→CEC→CRO	0.50
Piperacillin +Tazobactam (TZP)	2	1000	0.20	CTT	0.28	CTT→CTT	0.28	CTX→FEP→CTT	0.32
Cefepime (FEP)	4	0000	0.14	FEP	0.14	TZP→CEC	0.17	CXM→AMC→CTX	0.20

Table 3.4 **Summary of evolutionary steering towards low fitness genotypes.** For each of the 15 antibiotics we have derived the (ordered) sets of one, two and three steering drugs which increase the probability of evolution to the least fit optimum genotype to the highest possible. In the case that an ordered set of steering drugs increases the probability to 1 we have not considered ordered sets of greater length (marked as - in the table). * – as the landscape for SAM is single peaked there can be no combination of steering drugs which change the probability of finding the global optimum. In all experiments the initial population distribution is taken as $u = [1/2^L, \ldots, 1/2^L]$ and $r = 0$.

As an alternative metric of steerability we also search for a stronger steering condition – identifying combinations of steering antibiotics that maximize the probability that evolution proceeds to the least fit of the local optima of a final antibiotic. Table 3.4 show the results of an exhaustive search of drug sequences of up to four drugs. We found that, excluding the single peaked landscape for Ampicillin with Sulbactam, there exist 0 drugs for which a single other drug is able to steer the population to a configuration from which evolution to only the least fit optimum is possible. If pairs of drugs are used to steer there are 3 such drugs (including the example presented in the above demonstration) and if triples of steering drugs are considered there remains only 3. These findings suggest that through careful choice of steering drugs we may be able to prevent the emergence of resistance. During these experiments we found that 14/15 of the antibiotics in our experiment (Cefpodoxime (CPD) being excluded) appeared in an optimal steering combination of some length. This finding suggests that there is no "archetypal" steering drug and hence that many drugs which have fallen out of favour owing to their poor efficacy as first–line therapies could find use as effective steering drugs.

3.4.2 The Dangers of Arbitrary Drug Ordering

Whilst careful selection of drugs for steering can prevent the emergence of resistance, arbitrary drug orderings can also promote it. We performed an exhaustive *in silico* search of all singles, pairs of, and triples of steering drugs applied sequentially to prime the initial population u for a final application of each of the 15 antibiotics (See Table 3.5). We found that steering with a single drug increased the likelihood of the most resistant genotype emerging in 57.3% of cases and decreased the likelihood in 29.8% of cases. Steering with pairs of drugs increased the likelihood in 64.1% of cases and decreased it in 28.4% of cases and steering with triples increased the likelihood in 65.6% of cases and decreased it in 27.5%. We tested the robustness of these results to changes in the value of r in Equation 3.3 and found that each of these values are changed by less than 2% for $r = 1$. For $r \to \infty$ we found that 56.0%, 68.1% and 71.2% of singles, doubles and triples (respectively) of steering drugs increased the likelihood of the most highly resistant genotype being found whereas only 22.2%, 20.0% and 19.2% of singles, pairs and triples (respectively) decreased it.

For each of the antibiotics, except Cefaclor, Cefprozil and Ampicillin+Sulbactam (which is single peaked making steering irrelevant), we found that a random steering combination of length one, two or three is more likely to increase the chances of resistance than to reduce it. Indeed, for Piperacillin+Tazobactam and Ceftizoxime we found that a random steering combination will increase the probability of the most highly resistant genotype occurring in more than 80% of cases, suggesting that sequential multidrug treatments which use these antibiotics should proceed with caution. These findings suggest that the present system of determining sequential drug orderings without quantitative–optimisation based guidelines could in fact be promoting drug resistance and that to avoid resistance we must carefully consider the order in which drugs are applied.

3.5 The Impact of Variance in Fitness Measurements

In the previous sections we have derived results regarding collateral sensitivity and evolutionary steering under the assumption that the fitness values for the different genotypes are precisely known. In reality, this is likely not true owing both to intrinsic variation in fitness, arising from the stochastic nature of intracellular dynamics or drug action, and extrinsic variation, arising from experimental bias introduced between replicates. Here, we explore the impact of this variance on the robustness of our results.

We begin by noting that the fitness landscapes shown in Figure 3.3 are reported as an average growth rate derived from 12 replicates. However, neither the raw values for the 12 replicate experiments, nor a measure of variance, is reported. As such, it is not possible to directly assess the robustness of our results to variation in fitness values. As we take $r = 0$ in our initial analysis, fluctuations in fitness will not change our results provided that the variance

Final drug	# of single steering drugs better (*/15)	# of single steering drugs worse (*/15)	# of steering pairs better(*/225)	# of steering pairs worse (*/225)	# of steering triples better (*/3375)	# of steering triples worse (*/3375)
Ampicillin (AMP)	7 (46.7%)	7 (46.7%)	97 (43.1%)	126 (56.0%)	1373 (40.7%)	1998 (59.2%)
Amoxicillin (AMX)	5 (33.3%)	9 (60.0%)	60 (26.7%)	164 (72.9%)	787 (23.3%)	2578 (76.4%)
Cefaclor (CEC)	9 (60.0%)	5 (33.3%)	132 (58.7%)	88 (39.1%)	2087 (61.8%)	1264 (37.5%)
Cefotaxime (CTX)	4 (26.7%)	10 (66.7%)	67 (29.8%)	155 (68.9%)	989 (29.3%)	2335 (69.2%)
Ceftizoxime (ZOX)	2 (13.3%)	12 (80.0 %)	31 (13.8%)	193 (85.8%)	444 (13.2%)	2930 (86.8%)
Cefuroxime (CXM)	5 (33.3%)	9 (60.0%)	95 (42.2%)	128 (56.9%)	1486 (43.5%)	1885 (55.9%)
Ceftriaxone (CRO)	5 (33.3%)	9 (60.0%)	61 (27.1%)	163 (72.4%)	810 (24.0%)	2564 (76.0%)
Amoxicillin +Clav (AMC)	7 (46.7%)	7 (46.7%)	99 (44.0%)	124 (55.1%)	1428 (42.3%)	1943 (57.6%)
Ceftazidime (CAZ)	4 (26.7%)	10 (66.7%)	76 (33.8%)	147 (65.3%)	1218 (36.1%)	2153 (63.8%)
Cefotetan (CTT)	4 (26.7%)	10 (66.7%)	58 (25.8%)	166 (73.8%)	843 (25.0%)	2531 (75.0%)
Ampicillin +Sulbactam (SAM)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)
Cefprozil (CPR)	7 (46.7%)	7 (46.7%)	113 (50.2%)	110 (48.9%)	1703 (50.5%)	1668 (49.4%)
Cefpodoxime (CPD)	3 (20.0%)	11 (73.3%)	26 (11.6%)	195 (86.7%)	282 (8.4%)	3077(91.2%)
Piperacillin +Tazobactam (TZP)	2 (13.3%)	12 (80.0%)	11 (5.9%)	213 (94.7%)	81 (2.4%)	3293 (97.6%)
Cefepime (FEP)	3 (20.0%)	11 (73.3%)	34 (15.1%)	190 (84.4%)	402 (11.9%)	2972 (88.1%)
Overall	67 (29.8%)	129 (57.3%)	960 (28.4%)	2162 (64.1%)	13933 (27.5%)	33200 (65.6%)
Overall (Without steering for SAM)	67 (31.9%)	129 (61.4%)	960 (30.5%)	2162 (68.6%)	13933 (29.5%)	33200 (70.3%)

Table 3.5 **Exhaustive summary of all short steering strategies.** For each of the 15 antibiotics we determined how many of the possible single, ordered double and ordered triple combinations of steering drugs (allowing drugs to appear multiple times in the combination and also allowing the target drug to appear as a steering drug) improved or worsened the probability of the most resistant genotype being found after the steering drugs are applied in order, followed by the final target drug. In each case the initial population was given by $u = [1/2^L, \ldots, 1/2^L]$ and $r = 0$. As the SAM landscape is single peaked no combination of steering drugs will improve or worsen the outcome. As such, we have computed the overall numbers both with and without the contribution of the SAM row.

is sufficiently small that the ordering amongst genotypes remains unchanged. However, if the variance is high, then the fitness ordering amongst genotypes can change and the efficacy of steering strategy may be greatly reduced. In the absence of reported variance, we explore the impact of variation in fitness by re–sampling each fitness value, f, uniformly from the interval $[0.7f, 1.3f]$.

We begin by noting that it is the existence of multiple peaks within the fitness landscape that ultimately drives our results. Indeed, without multiple peaks evolutionary steering would not be possible and evolution would be entirely repeatable. As such, we began by verifying that the landscapes remain multi–peaked under resampling of the fitness values. Figure 3.12 shows a histogram of the number of peaks for each drug over 1000 resampled landscapes. We

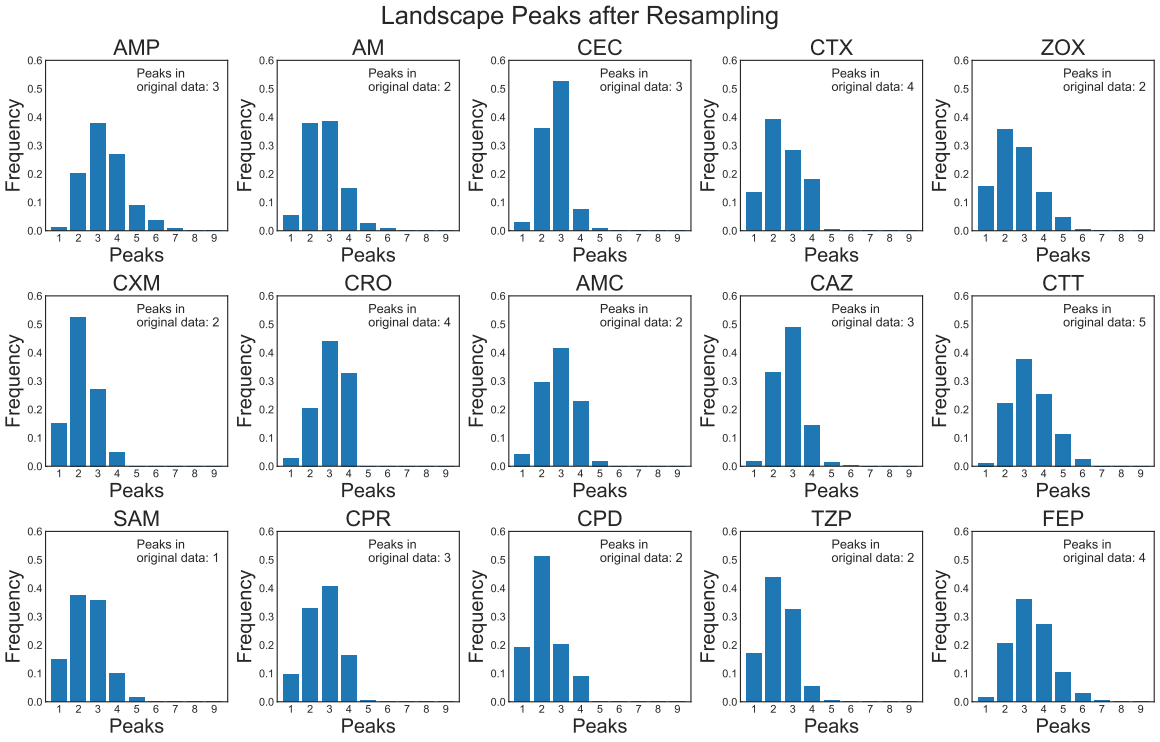

Fig. 3.12 **Histogram of peak numbers over 1000 resampled landscapes.** Each of the fitness landscapes derived by Mira et al. [228] was resampled 1000 times by resampling each fitness value f uniformly from an interval $[0.7f, 1.3f]$. For each sample the total number of peaks in the landscape was counted. Histograms depict the distribution of peak number over the 1000 samples for each of the original 15 landscapes. Note that the landscapes are very often multi-peaked for each drug, indicating the the potential for evolutionary divergence, and hence for non-repeatable collateral resposne, is robust to variation in the fitness values.

find that for all of the landscapes, including that for Ampicillin/Sulbactam which has only one peak in the original data, the most common number of peaks is two or three. This indicates the non–repeatability of evolution arising from 'branching' points in the fitness landscape is not the result of precise fitness values reported and is robust to variance in the fitness values.

To estimate the role of variance in driving the non-repeatability of evolution, we performed an *in silico* experiment wherein all of the fitness landscapes were resampled and a single table of collateral response generated. We then picked a pair of drugs exhibiting collateral sensitivity in this matrix, resampled their respective landscapes and calculated the probability of cross-resistance in these new resampled landscapes. Over 10000 simulations, we found that a random drug pair exhibits cross-resistance with probability 0.538 and a pair previously identified as collaterally sensitive exhibits cross-resistance with probability 0.531. This result indicates that where fitness values are subject to variation, for example due to experiment-specific (or patient-specific) factors, low-replicate-number experimental evolution studies may suggest drug pairs that exhibit collateral sensitivity with no better likelihood than guessing.

We next considered the effect of variance on the prediction of collateral sensitivity between drug pairs. Figure 3.13(A) shows a matrix of collateral sensitivity likelihoods generated by

resampling the landscapes for each drug pair 1000 times and calculating the average likelihood of collateral sensitivity. We call these values the *resampled collateral sensitivity likelihoods (rCSLs)*. Figure 3.13(B) shows the difference between the CSL calculated from the original dataset and the rCSL values, calculated as CSL − rCSL for each drug pair. Here, red entries denote drug pairs with high CSL but low rCSL, indicating that collateral sensitivity in these pairs is not robust to variation in the measured fitness values. Although few, these low-CSL-robustness pairs indicate that fitness-landscape-derived CSLs may be wildly inaccurate if the fitness landscapes themselves are incorrect, and that we should attempt to account for potential variance when predicting collaterally sensitive drug pairs from fitness landscapes. Where the variance in fitness values arises due to external noise arising from the experimental process, these results indicate that a greater number of replicates should be performed to estimate fitness. However, if the variance is intrinsic (or in the case of clinically prescribed antibiotics, patient specific) it may be unavoidable, and estimates of robustness such as the one we have performed here should be performed as part of any model-driven prediction of a drug sequencing strategy.

Figure 3.13(C) shows the rCSL matrix thresholded for entries with likelihood of collateral sensitivity, p, satisfying $p > 0.9$, $p > 0.8$ and $p > 0.7$. Here, we find that no pair of drugs exhibit guaranteed collateral sensitivity but that several pairs exhibit a high ($p > 0.8$) likelihood of collateral sensitivity. This result indicates that even where fitness is highly variable, the collateral response need not be entirely unpredictable, and drug pairs exhibiting high likelihood of collateral sensitivity may be identified.

To conclude our exploration of variance in fitness values, we explored the robustness of the steering strategies derived above. These strategies comprise a sequence of at most three drugs that, when prescribed in order, are predicted to steer evolution through the genotype space to a genotype from which the likelihood of evolution to the highest peak of a final drug landscape is minimised (in the case of Table 3.3), or are predicted to steer evolution to a genotype from which the likelihood of evolution to the lowest peak of a final drug landscape is maximised (in the case of Table 3.4). For each final drug, we resampled the fitness landscape and the landscapes of the drugs comprising the steering strategies outlined in Tables 3.3 and 3.4 before evaluating the likelihood that evolution under that steering strategy, followed by the final drug, ended at the highest peak of the final drug landscape (in the case of Table 3.3) or lowest peak of the final drug landscapes (in the case of Table 3.4). We performed this *in silico* experiment $S = 10000$ times. The results are presented as box-whisker charts in Figure 3.14.

Here, we see that the efficacy of the steering strategies is not robust to resampling the fitness values. Indeed, when attempting to avoid the highest peak of the final drug landscape, the median likelihood of ending at that peak is 0.25 or higher for every drug except CTX (median 0.17). In attempting to steer evolution to the lowest fitness peak the result is even worse, with no strategy achieving a median likelihood greater than 0.5. These results clearly indicate that in order to derive efficacious steering strategies from fitness landscapes, we will need to determine fitness values within a small confidence interval. If the uncertainty or

variance in fitness cannot be reduced, for example where it arises in a patient-specific manner, it may be that 'perfect' evolutionary steering cannot be achieved. In this case, we may need to rely on steering strategies that are *most likely* to steer evolution away from highly resistant genotypes, just as we explored those drugs that were most likely to induce collateral sensitivity. Designing experimental protocols that permit the precise and reproducible measurement of organismal fitness, or the measurement of patient specific fitness values, remains a critical open problem that must be solved in order to move model-derived evolutionary steering strategies into the clinic whilst maintaining confidence in their efficacy.

3.6 Implementation Details and Computational Considerations

The implementation of the matrix model, calculation of likelihoods for the evolutionary trajectories and plotting of all of the figures was performed using the Python programming language. A full annotated implementation of all models, the derivation of all results, and the code used to produce the figures is presented as a Jupyter notebook in Appendix 3. The full code listing is available from github (See Appendix 3 for details).

3.6.1 Performance

Our results are derived by repeated multiplication of small (16×16) matrices. As such, all results can be evaluated almost instantaneously on a personal computer. The full notebook is presented in Appendix 3 and can be evaluated within less than five minutes. The majority of this execution time is spent in the exploration of random *LK* landscapes (Section 3.2.1) and in the plotting methods.

3.6.2 Testing

The models and results presented in this chapter are highly dependent on matrix multiplication. In our code, we performed this matrix multiplication by use of the NumPy package for numerical computing in Python [167]. We did not attempt to explicitly verify the correctness of this implementation of matrix multiplication. To test the correctness of our model implementation and predictions, we compared the predictions of our model to two hand-worked examples. Further, we performed 1000 runs of generating a random *LK* landscape, constructing the associated matrix model and testing that: 1) the matrix was indeed a stochastic matrix defining a Markov chain, 2) the final genotype distribution predicted by the model contained only

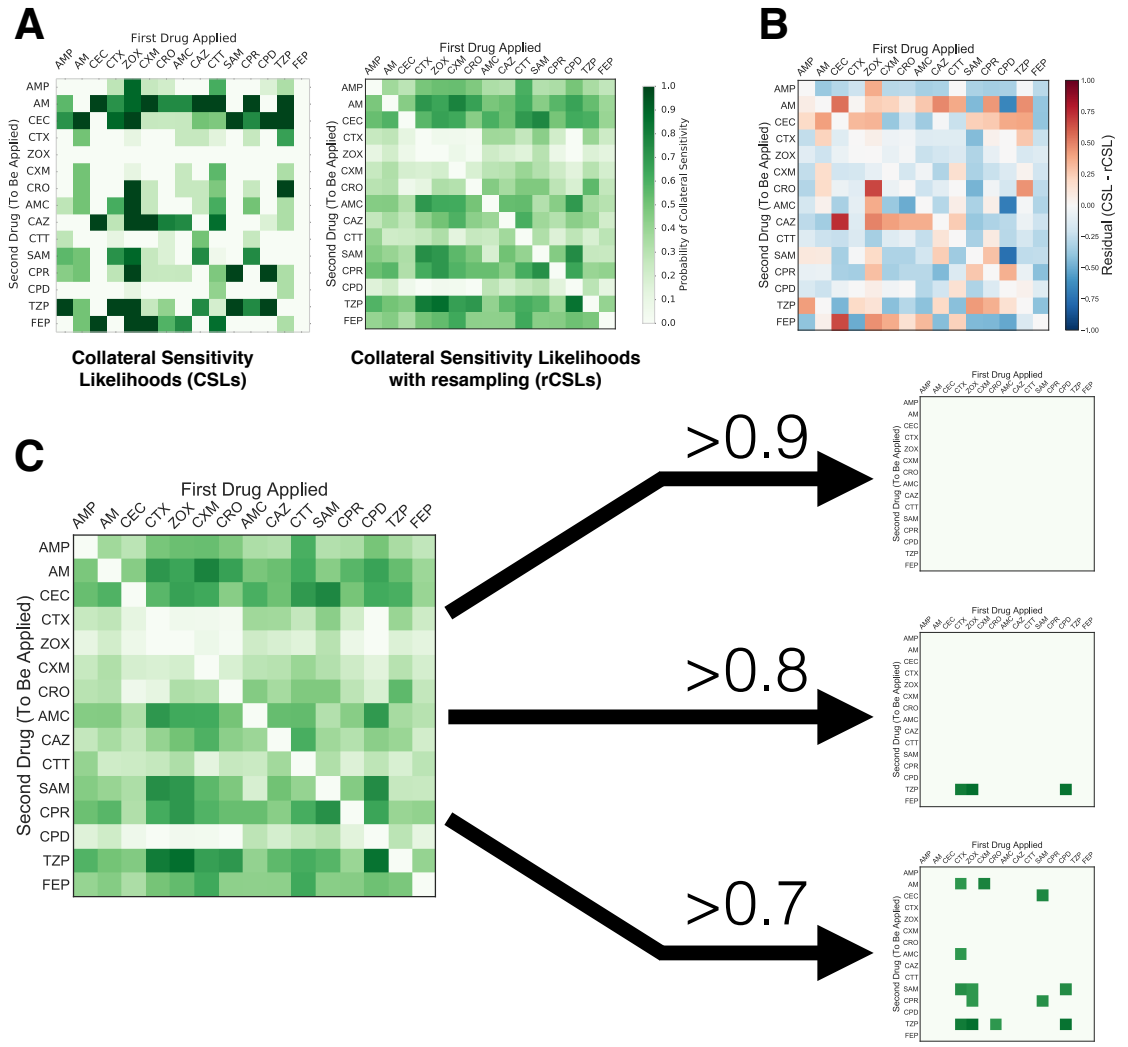

Fig. 3.13 **Collateral Sensitivity Likelihoods between landscapes with resampling (rCSLs).** **A** Matrices showing the likelihood of collateral sensitivity between pairs of drugs with the original landscapes (left) and calculated from 1000 resampled landscape pairs (right). **B** Matrix showing the difference between the original reported CSL and the resampled CSL value. Red entries indicate drug pairs that exhibited a high likelihood of collateral sensitivity in the original analysis but a low likelihood when the landscapes were resampled. These drug pairs can be considered as not robust to noise/variation in the measurement of the fitness values. Blue entries indicate drug pairs for which resampling increased the likelihood of collateral sensitivity. **C** The rCSL table thresholded for entries greater than 0.9, 0.8 and 0.7. There are no longer any guaranteed collaterally sensitive drug pairs, however despite resampling, there are drug pairs that exhibit a high likelihood of collateral sensitivity.

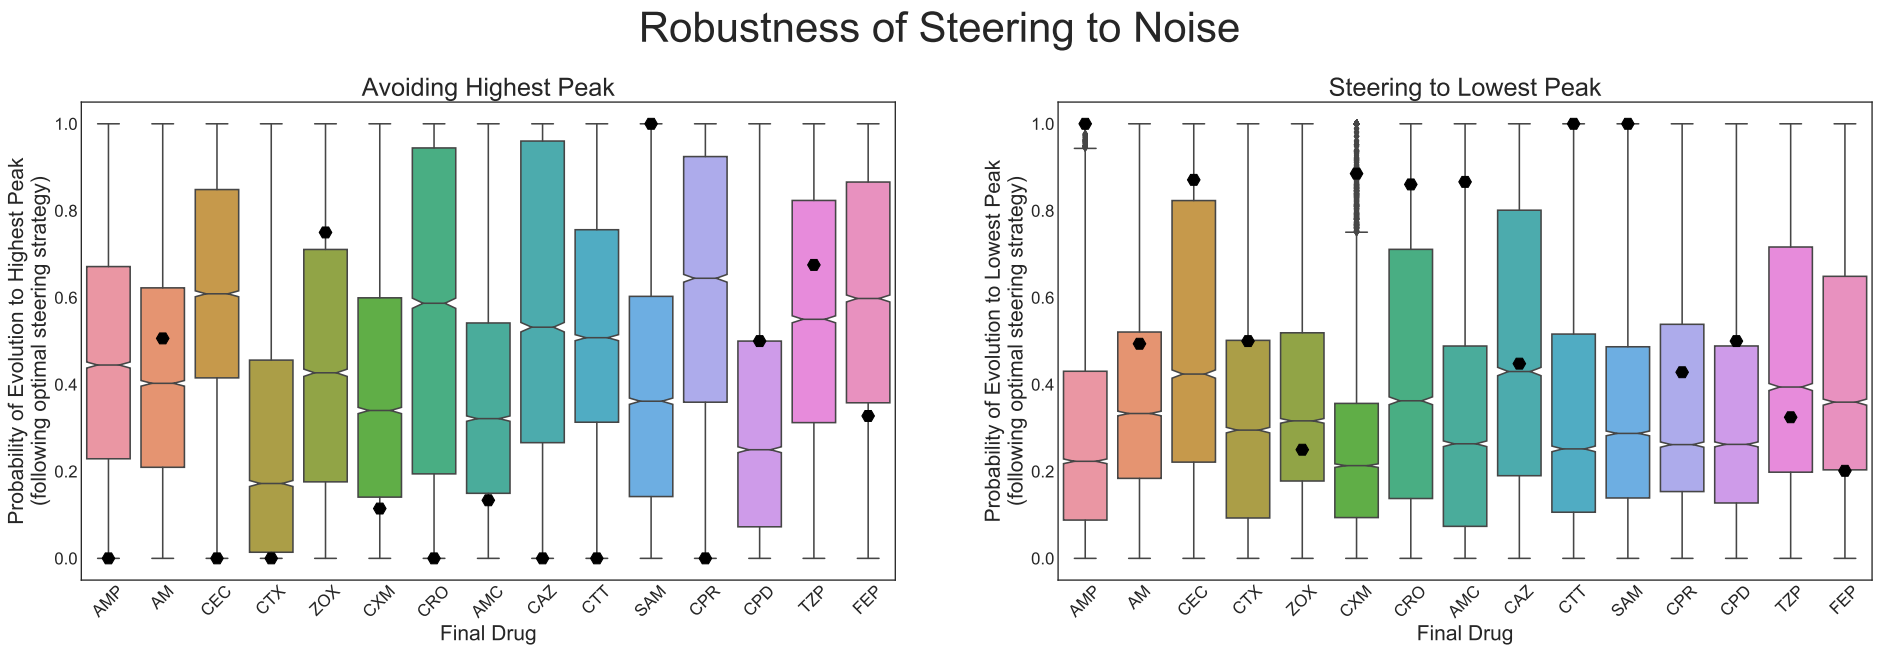

Fig. 3.14 **Robustness of evolutionary steering to variance in fitness.** The robustness of the evolutionary steering strategies outlined in Tables 3.3 and 3.4 was tested by resampling the fitness landscapes. For each final drug and its associated steering strategies, we performed $S = 10000$ replicates of resampling the landscapes before calculating the probability that the highest (left) or lowest (right) peak of the final landscape was reached following the associated steering strategy. These probabilities are presented as box-whisker plots. Black hexagons indicate the probability associated with the original landscape reported by [228].

genotypes that are peaks of the associated landscape, and 3) evolutionary convergence occurs within a number of matrix multiplications that is smaller than the total number of genotypes.

3.7 Discussion

We have introduced a mathematical model of evolution using the fitness landscape model for the GP–mapping. This model takes the SSWM assumptions on population dynamics to reduce evolutionary trajectories to a random walk on the directed graph of the genotype space. Using a parametrisation from empirically derived fitness landscapes, we used this model to test the evolutionary steering hypothesis: that drugs can be given in order to 'steer' evolution to a genetic configuration from which trajectories to high resistance are inaccessible.

One strength of the model presented here is that the SSWM mutations allow us to study evolution for a broad range of population sizes and mutation rates. Thus, these parameters do not need to be explicitly chosen and our results need not be restricted to the dynamics within a single patient. Goulart et al. [122] used the fitness landscape model to explore whole–hospital–scale antibiotic treatment strategies and the model we present here can also be used to make predictions at this scale. A major assumption in our modelling is that evolutionary trajectories converge to a local optima of fitness before the drug is switched. If evolutionary convergence is determined to be prohibitively slow for steering as a treatment strategy within a single patient, our results will still hold in scenarios which admit longer time-scales. For example, evolutionary steering could provide an effective means to avoid the emergence of drug resistance at the hospital scale or in life–long diseases such as HIV.

We conclude this chapter with an overview of the key results derived from the fitness landscape model and interpret these results in the context of current clinical and experimental approaches to preventing or reversing drug resistance. Finally, we will explore the limitations of our model, including the breakdown of the SSWM assumptions, and the potential impacts on evolutionary steering as a viable therapeutic strategy.

3.7.1 The Pitfalls of Collateral Response Analysis

Using mutliple drugs in sequence has emerged as a potential therapeutic strategy in treating highly resistant disease [366, 275, 233, 368, 161, 122, 255, 34]. As the discovery of novel antibiotics has stalled, sequential drug therapies offer a strategy to extend the life of a limited pool of antimicrobials. For cancer therapy, multiple drugs often cannot be prescribed together, owing to high toxicity, so sequential therapies are the only viable option for the use of multiple drugs. Experimentally identifying patterns of collateral sensitivity, wherein the evolution of resistance to one drug induces sensitivity to a second, has been explored as a potential method in designing these therapies. Recent *in vitro* studies have attempted to determine collateral sensitivity profiles and predict optimal drug sequences or combinations from them [366, 275, 233, 196, 161, 197]. Here, we have explored the shortcomings of

empirically–determining collateral response through low–replicate–number experimental evolution.

Owing to complex epistatic interactions within the genotype–phenotype map, the fitness landscapes induced by drugs are often multipeaked [261, 260], with many peaks accessible by evolutionary trajectories from each specified "wild–type" genotype. Thus, it is difficult to predict how drug resistance will evolve during the course of therapy and how the sensitivity to a second drug will change as this evolution proceeds. Evolution is inherently stochastic and can proceed along a number of distinct trajectories to differing genetic configurations. These distinct trajectories have important implications for the efficacy of collateral response analysis.

Firstly, we demonstrated that the differences in fitness between genotypes are likely not preserved when the landscape is changed. Thus, where two genotypes have high fitness in one landscape, they need not have correlated fitness in the second. Critically, this means that for two genotypes that are local optima of one landscape, one may be sensitive in the second whilst the other remains resistant. Hence, owing to the existence of multiple trajectories to distinct genotypes with high resistance, one drug may stochastically induce collateral sensitivity or collateral resistance in a second. Using the small combinatorically–complete fitness landscapes derived by Mira et al. [228] for 15 antibiotics, we identified 94 (out of a possible 225) pairs wherein collateral response is stochastically determined. As collateral response is often derived from a small number of experimental replicates, evolutionary trajectories can be missed and collateral sensitivity reported where collateral resistance can in fact occur. By exploring the space of *in silico* derived collateral sensitivity matrices, we found that drug pairings determined by selecting a pair showing collateral sensitivity in a randomly chosen CSM are only marginally better (probability 0.52 of collateral resistance) than those chosen at random (probability 0.58). This result undermines the current approaches that use experimental evolution as a tool in designing drug pairs to be given sequentially.

Second, from the mathematical construction of our model, we determined that evolution is likely non–commutative, in the sense that the order in which drugs are prescribed will have a (potentially large) impact on the resultant population genotype. Critically, this means that the effects of longer drug sequences cannot be predicted from information about pairs of drugs; a strategy employed by Imamovic and Sommer [161]. Thus, even if collateral response were repeatable, experimentally derived response amongst pairs is likely insufficient for the design of longer sequential therapies. To design such therapies, all distinct sequences of a given length must be treated as unique. As such, the number of experiments required (even without the increased number of replicates required to account for the non–repeatability of evolution) grows intractable even for modest length sequential therapies.

To improve collateral sensitivity analysis as a potential tool to aid clinical decision making, we propose alternative clinical metrics. To account for the inherent non–repeatability of evolution we suggest that a *likelihood* of collateral sensitivity should be reported for a drug pair. Through this metric the risk of resistance is explicitly stated and can be fine–tuned on a patient specific basis, for example dependent on the severity of infection, drug allergies or

therapies that have previously failed. However, likelihood measures do not account for the severity of resistance associated with failure. For this reason we propose that the worst case collateral response should also be reported for each drug pair. Indeed, in our analysis we find that some drugs have similar collateral sensitivity likelihoods, specifically Cefotaxime, Cefotetan and Cefuroxime, but drastically different ($2\times$ vs $16\times$) worst case responses. Thus, in order to distinguish between them as potential first–line therapies it is important to know both the *likelihood* of collateral sensitivity/resistance and the *severity* of resistance when it does occur. Whilst these alternative clinical metrics may offer a significant improvement over the current practice of arbitrary drug ordering, or the current experimental approach of deriving collateral response from *in vitro* evolution, there are still substantial drawbacks. For instance, this strategy likely requires considerably larger datasets or an intractably large number of experiments to be performed. Below, we discuss potential solutions to these problems.

3.7.2 Model–Driven Evolutionary Steering

The failure of empirically–derived collateral response estimates emphasizes the need to understand the complex, rugged, drug fitness landscapes, and thus to take into account additional features of the genotype–phenotype mapping, in designing sequential therapies. We have demonstrated the potential benefits of using a sequence of drugs (or even treatment breaks, which themselves impose a selective pressure [14], and which we explore further in Chapter 4) to steer, in an evolutionary sense, the disease population to a configuration from which resistance cannot evolve. Specifically, our model predicts that for 7 of the 15 β–lactam antibiotics for which empirical landscapes have been derived, there exists a sequence drugs of length at most three that can be used to steer evolution and restrict evolutionary trajectories to high resistance. However, for only three of the drugs can this steering be performed with a single prior drug, suggesting that restricting our search for sequential drug therapies solely to pairs, as in empirical derivation collateral response, may prevent us from identifying potentially effective therapies. Crucially, amongst the sequential drug therapies we derive, 14/15 of the drugs appear at least once, suggesting that there is no archetypal 'good' steering drug. From this observation we propose that antibiotics previously shelved due to their low curative efficacy could find a new lease of life as steering drugs in a multi–drug treatment therapy.

At present there exist no clinical guidelines regarding the order in which drugs should be applied. Through an exhaustive analysis of the evolutionary trajectories of *E. coli* under short drug sequences (two to four drugs) we found that the majority, approximately 70%, of sequential treatments with 2–4 drugs increase the likelihood of resistance emerging. Thus, we may be inadvertently selecting for highly resistant disease populations through arbitrary drug ordering in the same way that highly resistant disease can emerge through irresponsible drug dosing [279]. Perhaps the most striking example of this phenomenon is that of Piperacillin+Tazobactam, a drug combination often prescribed after others have failed, which has an increased likelihood of resistance when prescribed after a pair or triple of others

drugs in over 90% of cases. Coupled with the observation that drug–resistance may in fact be irreversible once it occurs [318], this highlights the critical need for improved antibiotic stewardship. Indeed, if we are to avoid resistance to our most effective drugs we must carefully consider how they are used together, both in combination and in sequence, and take appropriate steps to reduce the risks.

3.7.3 Variance in Fitness Measurements

We found that the evolutionary steering strategies predicted by our mathematical modelling were not robust to variance in the fitness values comprising the fitness landscapes. This result highlights two considerations that are critical to the model-driven design of therapeutic strategies. The first is that experimentally determined fitness measures are inherently noisy, and that this variance should be explicitly quantified and reported. A critical future experimental question will be how to derive fitness measurements with confidence, both in the laboratory and in the clinic. However, we should expect that uncertainty in fitness measurements will always be a factor and should approach computational work appropriately, by testing the robustness of our model-derived treatment strategies to variance in the reported values. Indeed, we found that a number of drug pairs exhibiting collateral sensitivity in the reported data were extremely fragile, often exhibiting cross-resistance when the landscapes were resampled. By identifying this fragility, we were able to identify drug pairs which exhibited a high likelihood of collateral sensitivity, even when the fitness landscapes were changed substantially by resampling.

We further found that the steering strategies we derived were not robust to resampling of the fitness values. Of course, this does not mean that the steering hypothesis is wrong, but rather that the design of effective steering strategies will be strongly reliant on the accurate and precise measurement of fitness landscapes. Whether these measurements can be made sufficiently precise, especially in the a patient-specific context, remains an open question and one that must be answered in order to translate evolutionarily-driven steering strategies into clinical practice. Even if fitness cannot be measured with sufficient accuracy, it is likely that mathematically-derived steering strategies (particularly those that explicitly account for this variance) will outperform current, entirely random, drug ordering.

3.7.4 Future Directions

Our model represents a major simplification of the evolutionary process. As such, a necessary first step in moving towards clinically–relevant predictions is to empirically verify our model predictions. For evolutionary steering, this verification has not yet been attempted owing to the relatively small subspace of the fitness landscape from which the predictions have been made. In fact, both the collateral sensitivity/worst case outcome metrics and the evolutionary steering predictions require considerably better datasets than exist at present. The derivation of such datasets should allow us to design better sequential therapies and also provide additional

information about the topology of the rugged drug fitness landscapes. Empirically deriving larger fitness landscape datasets is thus a crucial avenue of future work. For cancers this is especially true, as at present these exist no empirically measured landscapes for solid tumours at all.

Distributed Data Collection

Whilst collateral sensitivity matrices are often determined from a small number of empirical replicates, many more are needed to accurately determine the distribution of collateral response. Likewise for evolutionary steering, De Visser and Krug [71] find that there exist less than 20 systematic studies of fitness landscapes and that these studies consider between 3 and 9 possible mutations. For steering to be fully effective we must account for all likely fitness conferring mutations and their effects on fitness under many drugs. Thus, many of the studies reviewed by de Visser and Krug are insufficient for determining clinically–actionable steering strategies for certain diseases. Fortunately, for a number of highly resistant infectious diseases [358, 165] and cancers [207, 206] there are only a small number of mutations which seem to contribute to resistance. A potential solution to both of these problems lies in distributed data collection. For example, owing to a concerted data collection effort tracking genotypes and resistance assays to multiple common drugs, a database of over 70,000 HIV samples was created and work by Hinkley et al. [146] used this data set to introduce a method to approximate large fitness landscapes from relatively fewer data points through ridge regression. A similar approach of concentrated genotyping and drug resistance assays could help in deriving bacterial, or even cancer, fitness landscapes. In the case of collateral sensitivity likelihood and worst case response, no genomic sequencing need be performed. Only before and after treatment drug assays are required, meaning that the relevant datasets are surprisingly inexpensive to derive. In fact, for some bacterial infections these measurements are already made and deriving datasets is a problem of proper centralised data storage rather than encouraging new or expensive clinical practice.

Model Improvement

The mathematical model of evolution used here to predict collateral response from empirically–derived fitness landscapes is a simplification of the complex dynamic process that drives drug resistance. However in spite of our simplifying assumptions, we have shown that evolution is not repeatable even in this simple model and it is unlikely to become so once more complexity is introduced. It follows that the criticisms of empirical collateral sensitivity analysis in its current form are well founded and that additional, more nuanced techniques must be developed. Here, we have taken the first steps towards such techniques.

Two major assumptions within our modelling are that drugs are administered for sufficiently long that evolution can converge to a local fitness optimum and that this convergence is guaranteed to occur. This assumption poses two potential problems in converting our model

predictions of evolutionary steering into evolutionarily–informed clinical practice. The first is that if selection is strong and mutations are rare then there is a possibility of the population being driven to extinction before an adaptive mutation occurs. We have chosen to ignore this possibility within our modelling as, in the context of treating bacterial infections or cancers, this would constitute a success. We have not attempted to estimate the likelihood of extinction as to do so would require us to explicitly fix a population size and mutation rate, which can take a large range of values under the SSWM model. The second is that the time to convergence could be prohibitively long for steering to constitute a realistic treatment strategy. We believe that the assumption of reasonable convergence times could be valid as adaptive walks in rugged landscapes are often short [248]. However, it has been shown that for certain landscapes there can exist adaptive walks of length exponential in the number of loci [176]. Since we get to choose those drugs with which to steer, we can avoid landscapes for which the convergence time is prohibitively long. A full theoretical and experimental consideration of the impact of evolutionary convergence times and genetically heterogeneous populations constitutes the next step in translating our model predictions to clinical practice.

The Strong Selection Weak Mutation model we have used here is a highly simplified, yet well studied model of evolution. The model ignores much of the complexity of the evolutionary process, specifically simplifying the genotype–phenotype map, ignoring the disease microenvironment and making the assumption of a monomorphic disease population in which deleterious and neutral mutations cannot fix. Under certain regimes of population size and mutation rates these simplifying assumptions break down [351]. For example, if the population is sufficiently large then stochastic tunnelling [163] – the situation where double mutations can occur allowing the crossing of fitness valleys – can arise causing a breakdown of the Strong Selection assumption. Similarly, if the mutation rate is sufficiently high then the population ceases to be monomorphic and forms a quasispecies [240, 41]. In both of these cases the population may become heterogeneous. Evolutionary steering, or potentially 'herding' to a more genetically homogeneous population, may still be possible in this case, but a more explicit model of population dynamics will likely be required. For example, an explicit agent based model, or an approach similar to that used to study quasispecies [79] may be more appropriate. However, such models will likely require explicitly defined population sizes and mutation rates which may render them less broadly applicable than the SSWM model. One potentially interesting avenue of further study will be to determine how effective the steering strategies derived from the SSWM model are as the underlying assumptions begin to break down.

If the population is sufficiently small then it becomes possible for deleterious mutations to fix [230, 359, 89]. Further, in our modelling we have ignored the possibility of neutral spaces in the fitness landscape which have been shown to have significant impact on whether non-optimal genotypes can fix in the population, as well as the time taken for evolution to find a locally optimal genotype [286, 285]. If both deleterious and neutral mutations are permitted to fix in the population, but the population is still assumed to be isogenic, the previously

published Markov chain model of [295] can be used to study the population dynamics. In this case, the Markov chain is irreducible and the stationary distribution over genotypes is unique. As such, evolutionary steering in the sense we have explored here is impossible and the key question becomes how to order drugs to maximise the first hitting time of the set of highly resistant genotypes.

We believe that each of the possible breakdowns of the SSWM model will have important consequences for the possibility and efficacy of steering, but a full treatment of their implications is beyond the scope of this thesis. In our future work we aim to undertake a comprehensive study of the implications of population size, mutation rate, neutral drift and evolutionary convergence times on the steering hypothesis.

Finally, we note that a key phenomenon in the evolution of drug resistance, and one which was highlighted as under–represented in studies of the genotype–phenotype mapping, is the existence of mechanisms of non–genetic heterogeneity in phenotypes. The fitness landscape model entirely omits the potential for drug resistant phenotypes to arise in a stochastic or environmentally mediated manner, even though these mechanisms of resistance appear to be common in both bacterial infections and cancers [145, 357, 134, 204]. In the next two chapters, we shift our focus towards mechanisms of non–genetic heterogeneity and their impact on the emergence of drug resistance. As with the work presented here, we focus on the potential for evolutionarily–informed treatment strategies to outperform the clinical standard of care.

Chapter 4

Modelling Non–Genetic Heterogeneity

4.1 Introduction

Treatment resistance in many diseases is driven by the pre–existence of resistant phenotypes within the population. Why such phenotypes co-exist (with sensitive phenotypes) and persist in environments never exposed to drug treatment remains a significant unanswered question. Phenotypic heterogeneity has been observed within isogenic populations of a number of organisms and at many scales [21], from the unicellular – bacteria [327], fungi [203] or cancer cells [134] – through insects [66, 151], plants [48] and even aspects of human development [322]. Importantly, this inter–cellular variation has been observed even in homogeneous and constant environments, suggesting that aspects of organismal phenotype may be stochastically determined.

In environments that fluctuate unpredictably this phenomenon can serve as a survival mechanism by increasing the likelihood that at least some offspring are well–adapted to future environments. Thus, non–genetic, non–environmentally–driven variation in phenotypes has been termed *bet–hedging*, as a species diversifies the phenotypes within the population in order to "hedge its bets" against environmental change (see [292] for justification of this naming, and [70] for a discussion of what evolutionary phenomena can be considered bet–hedging). Oscillatory environments are common in a range of ecological settings including fluctuating climates, immune–pathogen interactions or cyclic hypoxia within tumours, and the range of phenotypic traits which are thought to display stochastic determination is just as broad.

Bet–hedging can offer a survival mechanism in the event of rare catastrophic environmental change. An important clinical example is that of persister cells which arise stochastically within isogenic populations of infectious bacteria such as *Escherichia coli* [20, 204, 327]. These cells, which constitute a small fraction of the population ($< 1\%$, [204]), have reduced metabolism and shut down non–essential cellular functions. In this dormant state, the persister cells are tolerant to the cytotoxic effects of a number of antibiotic agents. Although dormant, these cells can retain the ability to proliferate (although at a drastically reduced rate) and when antibiotic treatment ceases persisters will proliferate, producing non–persisters and driving the re–emergence of the bacterial population. Hence bet–hedging, by creating a small sub–

population impervious to those therapies that act on proliferating cells, proves to be an effective survival mechanism against antibiotic treatment. Indeed, bacterial persisters are thought to be a contributing factor to multidrug resistance in a number of diseases [178, 204, 237] and are implicated in the dormancy of chronic diseases, such as Tuberculosis, which can be suppressed but not eradicated [365]. Novel treatment strategies capable of effectively killing persister cells are desperately needed and this need will continue to grow with the increasing incidence of resistance to our presently most effective antibiotics.

In cancer, bet–hedging has been minimally studied; however, a number of aspects of disease course suggest that bet-hedging mechanisms may be important for understanding how tumours evade therapy. Significant regression of tumours post-therapy leads to a period of remission, followed by the regrowth of aggressive, therapy-resistant lesions. These dynamics can be explained by the clonal model of cancer [242, 225, 128, 309, 271], wherein recurring drug–resistant cells are those that have stochastically acquired resistance mechanisms through genetic mutation. However, the high frequency of tumour recurrence in many cancers suggests that therapeutic escape cannot be solely based on mutational 'luck'. Experimental results have shown evidence of transitory resistance [192, 363] indicative of the existence of a small drug–resistant subpopulation that re–establishes a drug–sensitive cancer cell population. Recent experiments have identified the existence of such populations of 'cancer persister cells' in a cell line of EGFR+ non–small cell lung cancer [296] indicating that bet–hedging may play a role in the emergence of cancer drug resistance [265]. Further, bet–hedging may also drive tumour dormancy, the clinical phenomenon wherein residual disease lies dormant for a period before driving recurrence [4]. Thus, an understanding of bet-hedging in normal and abnormal (e.g. cancer) cell function may help us understand why certain types of therapies fail while others succeed.

4.1.1 The Causes of Bet–Hedging

A number of causes of bet–hedging have been identified across different species, but in many cases the cause remains an open question. The difficulty in identifying the precise drivers lies in distinguishing between variability (or *noise*) introduced at different biological scales. For example, gene promoter, transcription and translation dynamics are driven by inherently stochastic molecular interactions that result in the expression of gene products that vary both temporally and between isogenic individuals [264, 81, 169, 33]. These gene products interact in non–linear molecular networks, often forming feedback loops that have the potential to suppress [28] or amplify noise [144], to precisely tune the concentration of a molecular output, or to induce oscillations in the concentrations of molecules [43, 42]. This intra–cellular system is further modulated by variability in environmental factors and inter–cellular signalling.

Current biological thought is that noise in the levels of specific intra–cellular proteins may drive non–genetic phenotype differentiation. Indeed, under certain regimes of promoter switching, transcription and translation, protein abundance can follow a bimodal distribution [169] – inducing two distinct phenotypes in a population. Theoretical modelling coupled

with experimental validation has highlighted how bistable auto–regulatory genetic motifs can induce bimodal protein distributions [162, 144, 27].

The evolutionary origin of bet–hedging is unclear. It is not known whether bet–hedging emerged as an adaptation to unpredictable environments or as a spandrel [125], arising from the inherently stochastic nature of the biochemical reactions governing cellular behaviour, and later co-opted as a survival mechanism. What is clear is that bet–hedging strategies, manifested as phenotypically heterogeneous populations, are subject to natural selection. Beaumont et al. [26] demonstrated the *de novo* evolution of bet–hedging in the phenotypic trait of colony morphology, of the bacterium *Pseudomonas fluorescens*, by imposing stochastically fluctuating environments through replating. The genetic driver underpinning this switching behaviour was partially elucidated by Gallie et al. [99] who identified a single nucleotide change in the gene *carB* as responsible for the emergence of phenotype switching, however, identifying the precise molecular pathways through which this mutation acts to produce bet–hedging remains an open problem. Following the recent development of persister isolation techniques, a number of genetic drivers thought to contribute to bacterial persistence in *E. coli* have been identified [204]. However, whilst over–expression or deletion of these genes were shown to impact the *proportion* of bacterial persisters within a population, none was found to completely inhibit the persister phenotype, suggesting that bet–hedging may arise from the interactions of multiple gene products.

To address the difficulty in identifying genetic drivers whose mutation can inhibit bet–hedging, we introduce a new theoretical model of the GP–mapping wherein phenotypes emerge through the stochastic interactions of proteins in intra–cellular molecular pathways. We explore bet-hedging, which arises naturally within this model, from the perspective of *network intrinsic noise* as opposed to the *gene intrinsic* perspective of previous modelling (Figure 4.1). We argue that, owing to robustness in the network, there need not be single genetic drivers of bet–hedging and thus the reductionist gene–knock-out approach may be doomed to fail. We explore the impact of the structure of the GP–mapping on the rate of evolutionary loss of bet–hedging in environments that select against it. Finally, we highlight how understanding of the GP–mapping is essential in predicting the efficacy of the adaptive therapy strategy of breaks from treatment in treating chronic illnesses.

4.2 Previous Mathematical Modelling of Bet–Hedging and Non–Genetic Heterogeneity

Müller et al. [232] as well as others [356, 320, 193] have used mathematical techniques to demonstrate that bet–hedging offers a fitness advantage in stochastically fluctuating environments, offering strictly greater expected fitness than a deterministic one–phenotype strategy, provided that the fluctuations are neither sufficiently slow that genetic mutation and fixation (of a one–phenotype strategy) can occur, or sufficiently fast that no phenotype can survive to

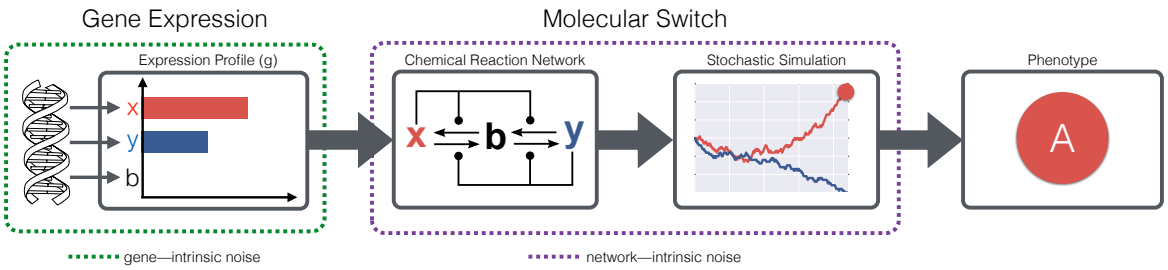

Fig. 4.1 **Schematic representation of the chemical reaction network model for determining phenotypes from genotypes.** The gene expression profiles (g) are assumed to be fixed for each genotype and the dynamics of expression from the biological genotype (green dashed leftmost box) are ignored. These modelling assumptions allow us to explore the implications of network–instrinsic noise (purple dashed central box) independently of gene–intrinsic noise.

reproduce. These models highlight that the average population fitness is maximised when the likelihood of phenotypic switching matches the likelihood of environmental switching. Further theoretical work by Botero et al. [37] considers when bet–hedging can offer a greater fitness advantage than *phenotypic plasticity*, where phenotypes are modulated via the individual's ability to sense environmental changes [328]. This work derives constraints on the cost of sensing, predictability of environmental fluctuations and the fitness effects of environmental change to determine when bet–hedging, plasticity or determinism offers a selective advantage. Jablonka et al. [164] also considered the trade–offs between genetic, stochastic and plastic phenotype determination; highlighting that where a plastic change in phenotype takes time of the order of expected environment change, a bet–hedging strategy is superior.

These mathematical results are in agreement with the results of a number of empirical studies. Beaumont et al. [26] demonstrated the *de novo* evolution of bet–hedging in the phenotypic trait of colony morphology in *Pseudomonas fluorescens* under artificially fluctuated selection. Acar et al. [1] engineered a strain of *Saccharomyces cerevisiae* to stochastically switch phenotypes as the result of stochastic gene expression and empirically demonstrated that fitness is maximised when the phenotype switching rate matches that of the environment – a direct corroboration of the findings of mathematical modelling.

It has been suggested that drug insensitive cells that arise stochastically in an isogenic population can facilitate the emergence of genetically–driven resistance by providing a window of opportunity in which resistance conferring mutations can arise [38]. However, this window is not indefinite, as drug insensitive cells will revert to a sensitive state and likely die in the presence of a drug. Charlebois et al. [46] explored this phenomenon through a mathematical model that incorporates switching from drug–insensitive to drug–sensitive phenotype as the stochastic relaxation from a state of high to low gene expression. This study demonstrated that the time scale of relaxation necessary to facilitate a high likelihood of genetic resistance is comparable to time scales measured for certain genes implicated in human cancers. In further work, Charlebois et al. [47] introduced a model of a feed–forward transcriptional regulatory network to demonstrate that the network architecture can extend the time that drug–insensitive

cells maintain their phenotype, and thus, can increase the likelihood of therapeutic escape occurring through genetically–driven mechanisms. This work highlights that to understand bet–hedging–driven drug resistance it may be necessary to look beyond the genetic scale and toward the gene–gene interactions that comprise the genotype–phenotype mapping.

The mathematical results of Müller et al. [232] and others [320, 356, 193] suggest a paradox when compared to clinical observations of bet–hedging, for example bacterial persistence, as a survival mechanism against rare catastrophic events. Specifically, in hospitable environments bacterial persisters reproduce more slowly than cells with a proliferative phenotype, reducing population fitness. It follows that natural selection will act to minimize the number of persisters that arise in the population, potentially resulting in the disappearance of this trait entirely. Where catastrophic events are rare, we should expect bet–hedging strategies to be lost before the event occurs, resulting in extinction of the population when it eventually does. Despite this selective pressure, persisters remain present in bacterial populations and provide a survival mechanism when antibiotics are introduced into the environment. We propose below that this disconnect can be explained by the structure of the GP–map.

4.3 A Stochastic Genotype–Phenotype Mapping Model

To investigate the impact of the GP–map on the evolution of bet–hedging, we implemented a model GP–mapping, in the sense of an evolutionary system $\mathfrak{E}$ introduced in Chapter 2, in which multiple phenotypes can emerge stochastically from a given genotype. At the heart of this GP–map is an abstract model of chemical reaction networks (CRNs) that are simulated computationally. A CRN is defined as a pair (S,R) specifying a set of chemical species S, which we represent as symbols (x, y, z etc.), and a collection of reactions R. Each reaction $r \in R$ is specified by a (not necessarily distinct) collection of reactants $s_1, \ldots s_k \in S$, a collection of products $s'_1, \ldots, s'_m \in S$ and a reaction rate ρ. We write reactions in short hand as

$$r := s_1 + \cdots + s_k \xrightarrow{\rho} s'_1, \ldots, s'_m. \tag{4.1}$$

As CRNs explicitly represent the biological mechanisms at work within the living cell, they provide a powerful tool for studying the GP–map, with the potential to be closely tied to empirical datasets. Owing to the complexity in explicitly simulating the progression of CRNs, simplified mathematical models have been derived. For example, continuous models of the mass action kinetics [152] provide insight into the average case behaviour and can be extended to account for noise through the use of stochastic differential equations [114]. As our GP–mapping will be pre–computed, we opt to forgo these simplified models in favour of explicit simulation. The chemical reaction networks are stochastically simulated using the Gillespie algorithm with tau–leaping [113, 115]. This algorithm determines a stochastic

progression of the network and returns the sequence of reactions that occur and the times taken between them.

It is entirely infeasible to explicitly model the full array of chemical interactions comprising the environmentally–mediated translation from genes to phenotypes. However following the paradigm of abstraction, we can investigate smaller CRNs to gain insight into the properties that the full mapping may possess. This approach was taken by Cardelli [42] who studied emulation between CRNs – the phenomenon where one network is capable of reproducing the exact mass–action kinetics of another. Cardelli derives algebraic criteria for the existence of emulations from structural properties alone, providing a method to extend studies of simple CRN motifs to larger networks. A well studied class of chemical reaction networks are those which encode bistable (or multistable) switches [43] in which the series of reactions must eventually terminate in a stable configuration from which no further reactions are possible. The different final configurations of a bistable CRN can be considered different states of a stochastic switch. The probability that the CRN progresses to a specific switch state is dependent on the initial conditions for the network. We implement an evolutionary system in which the underlying mapping from genotypes to phenotypes is determined by a bistable molecular switch as follows.

4.3.1 The Model Genotype–Phenotype Mapping

We assume that cells take one of two phenotypes, $\mathscr{P} = \{A, B\}$, that will correspond to the two stable configurations of a bistable CRN (S, R). Depending on the biological system in question, these phenotypes could, for example, represent normal/persister cells in *E. coli* or stem/non–stem cells in a population of cancer cells. The genotypes, G are modelled by varying the initial numbers of the chemical species within the CRN. This approach forgoes modelling the mechanisms of transcription and translation that have formed the basis of previous studies [169] and instead considers genotypes as represented by stable gene–expression profiles. The rationale for this approach is that we focus on the implications of *network–intrinsic noise* and network architecture without the additional complexity of *gene–intrinsic noise*[1].

Genotypes are taken as the initial and instantaneous expression of a collection of proteins that form the initial conditions of a CRN. Thus, $G = \mathbb{N}^{|S|}$. For simplicity, we assume that there exist only two chemical species comprising the genotype, x and y, whose expression is non–zero in the initial conditions and that genotypes take the form $g = (x_0, y_0)$. Further, we assume that there exists a maximum expression $g_{\max}$ that neither x_0 or y_0 can exceed. Mutations are modelled as changes to the initial abundance of x_0 and y_0. For an individual

[1] Although we do not make explicit use of this fact, we note that these genotypic stable expression profiles in part correspond to the *phenotypes* ($p = \mathbf{S}_\infty$) of the gene regulatory network (GRN) model (Section 2.2.4). Thus, our model can be viewed as an extension of the GRN model that bridges between gene products and cellular phenotype.

with genotype $g = (x_0, y_0)$ the possible mutants are given by

$$\mu((x_0, y_0)) = \{(x_0 \pm 1, y_0 \pm 1) \,|\, x_0 \text{ and } y_0 \text{ remain in } \{0, \ldots, g_{\max}\}\ \} \tag{4.2}$$

We assume that mutations occur at replication with some probability q and that where more than one mutation is possible each is equally likely.

We define a GP–mapping, $\mathscr{R}_{\text{GP}}$, from $g \in G$ to a phenotype $p \in \mathscr{P}$ by stochastically simulating a bistable CRN (S, R) from an initial configuration determined by the genotype g. Thus, $\mathscr{R}_{\text{GP}}$ is defined by

$$g \overset{\text{GP}}{\to} \begin{cases} A, & \text{if the simulation starting config. } g \text{ progresses to stable state 1} \\ B, & \text{if the simulation starting config. } g \text{ progresses to stable state 2.} \end{cases} \tag{4.3}$$

A schematic representation of this model is shown in Figure 4.1. In this model we assume that convergence time to an equilibrium state in the CRN is sufficiently fast (in comparison to the cell cycle time-scale) that it may be taken to be instantaneous. Thus, the phenotype is modelled as being determined instantaneously at birth. This assumption permits us to ignore the timing information provided by the Gillespie algorithm. As such, the abscissas of all figures showing stochastic simulations of chemical reaction networks presented in this work measure time discretely, in terms of the number of reactions that have occurred, instead of continuously. The time between successive reactions in the Gillespie algorithm is dependent on a propensity function that accounts for reaction rates and the volume of the container (i.e. cell cytoplasm). As we are only interested in the probability of finding each stable configuration, and not the precise time taken to reach this configuration, we may set the volume to be an arbitrary constant, say 1.0 mm^3, and also normalize the reaction rates such that at least one reaction has rate 1.0 sec^{-1}.

Note that the genotype g corresponding to a stable gene–expression profile is the sole heritable determinant of phenotype in this model. The state of the molecular switch, i.e. the phenotype, of a parent individual has no influence on the phenotype of the offspring. This assumption can be justified by observing that the bistable stable switch is likely to represent only a small sub–network of the complex and dynamic interaction network governing the GP–mapping. Thus, we can expect the chemical species comprising the molecular switch, x and y, to be further transformed or consumed in additional unmodelled reactions that determine phenotype. Further, even when this is not the case, we can expect the x and y molecules to decay naturally over time. Thus, the omission of epigenetic inheritance of switch state can be considered an assumption that the time scale of molecular decay is much faster than that of cellular division. Weakening this assumption, and permitting epigenetic inheritance, represents a potential extension of the model that is explored in Chapter 5.

This model also assumes a single genetic burst of x and y molecules occurs at birth, and that no further molecules are produced throughout the life of an individual. This is a simplifying assumption. Under these assumptions, our model can be considered to be a

stochastic, irreversible developmental program in the case of differentiating cancer cells or higher organisms. In the context of bacteria, the developmental perspective is less appropriate and the model can be justified by considering gene–expression bursts that are cell–cycle dependent (specifically, bursts that occur at the start of the G1 phase). We note that the bistable switches we consider in our analysis display a degree of hysteresis. Thus, the production of a small number of molecules during the life of an individual is unlikely to result in a switch of phenotype.

We note that, in this model, stable configurations of the CRN are analogous to local minima in the epigenetic landscape metaphor introduced by Huang [154, 155] to study phenotypic heterogeneity (Section 2.2.4). Where our GP–mapping differs is that we explicitly simulate the intra–cellular stochasticity introduced by thermal fluctuations through the Gillespie τ–leaping algorithm. As such, our model explicitly distinguishes between intra– and extra–cellular noise and, as a result, between phenotypic heterogeneity driven by bet–hedging and phenotypic plasticity, whereas these two process take an identical form in Huang's work.

4.3.2 Examples of CRN Genotype–Phenotype Mappings

There are numerous examples of bistable switches that can serve as genotype–phenotype mappings within this evolutionary framework. We consider four different bistable switches that can serve as model GP–mappings. These molecular switches, along with examples of their dynamics, are presented in Figure 4.2 where normal arrows indicate the transformation of a chemical species and and dot–headed arrows indicate catalysis. The *Direct Competition* (DC) switch along with two biased versions of this switch (DCx, DCy) consist of a pair of autocatalytic reactions. The *Approximate Majority* switch studied by Angluin et al. [17] and later Cardelli and Csikász-Nagy [43] consists of two catalytic and two autocatalytic reactions. (A biological implementation of the AM switch is presented by Dodd et al. [75] as a potential mechanism for epigenetic cell memory).

The second row of Figure 4.2 shows multiple stochastic simulations of the bistable chemical reactions networks generated by the Gillespie algorithm [113, 115]. This algorithm determines a stochastic progression of the network and returns the sequence of reactions that occur and the times taken between them. The third row of Figure 4.2 shows heatmaps associating with each genotype $g = (x_0, y_0) \in G$ the probability that an individual with that genotype will have phenotype A. For the DC, DCx and DCy switch the stochastic simulation is equivalent to the classical drunkards walk with two absorbing boundaries. Thus, the probabilities can be

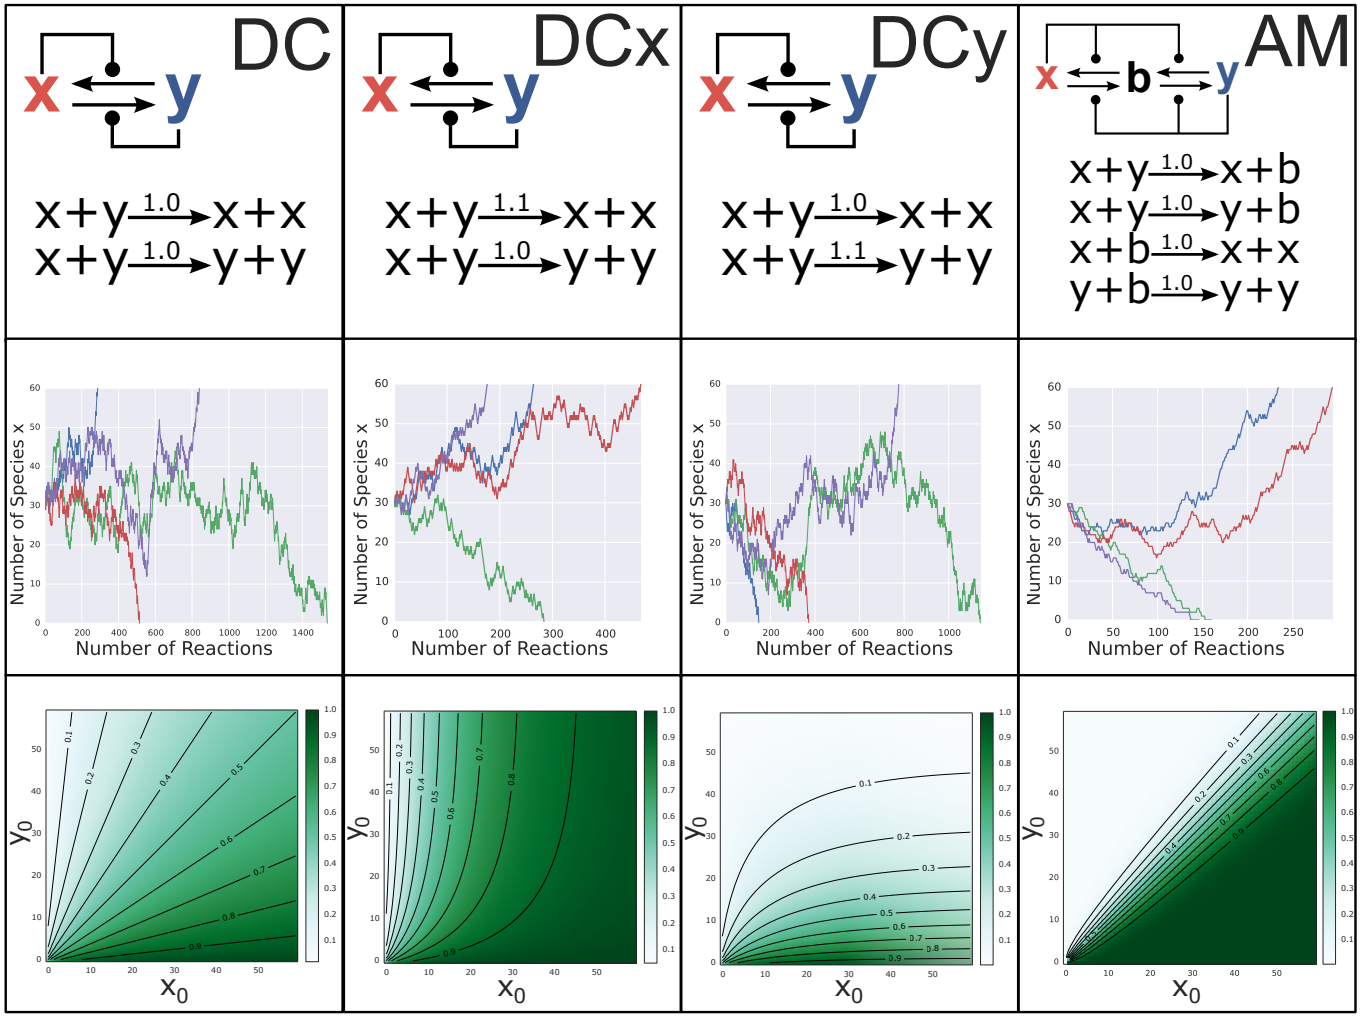

Fig. 4.2 **Example molecular switches as GP–mappings.** Each column shows the characteristics of one of the four switches (Direct Competition, x–Biased Direct Competition, y–Biased Direct Competition and Approximate Majority) introduced in the main text. The first row shows the name, chemical reaction network structure and precise definition of each switch. The second row shows stochastic trajectories of the number of molecules x in the system for four different simulations of each switch. The starting condition in all simulations is $x = y = 30$, ($b = 0$ for the AM network). Note that all of the switches are able to resolve to either of the stable conditions, $x = 60$ or $x = 0$, which correspond to the phenotypes A and B respectively. Row three shows contour plots displaying the probability of switching to phenotype A for each possible initial condition with $0 < x_0, y_0 \leq 60$ and $x_0 + y_0 \neq 0$ ($b = 0$ for the AM switch). Contour lines show subspaces of genotype space of equal hedging probability for hedges equal to $0.1, 0.2, \ldots, 0.9$.

computed explicitly as

$$\mathbb{P}\left(\text{ Phenotype } A \mid \text{ DC switch }, g=(x_0,y_0)\right) = \frac{x_0}{x_0+y_0} \tag{4.4}$$

$$\mathbb{P}\left(\text{ Phenotype } A \mid \text{ DCx switch }, g=(x_0,y_0)\right) = \frac{1-\left(\frac{10}{11}\right)^{x_0}}{1-\left(\frac{10}{11}\right)^{x_0+y_0}} \tag{4.5}$$

$$\mathbb{P}\left(\text{ Phenotype } A \mid \text{ DCy switch }, g=(x_0,y_0)\right) = \frac{1-\left(\frac{11}{10}\right)^{x_0}}{1-\left(\frac{11}{10}\right)^{x_0+y_0}} \tag{4.6}$$

without the need to explicitly simulate the dynamics of the CRN. For the AM switch there is no known closed form solution to calculate the probability of switching to the all x configuration corresponding to phenotype A. We calculate this probability by constructing a Markov chain on the state space of possible configurations. The details are provided in Appendix B.

4.3.3 The CRN Model can Approximate Arbitrary Hedges

The first key observation of this model is that by picking appropriate initial conditions (i.e., genotypes) for the chemical reaction networks, we can instantiate the CRN model to closely approximate any probability of switching to the phenotype A configuration (Figure 4.2 Row 3). Equivalently, we can approximate any proportion of phenotypes arbitrarily closely. However, in order to do so we must choose a sufficiently large value of $g_{\max}$. Figure 4.3 shows for increasing values of $g_{\max}$ the maximum difference between a target switching probability, p, and the closest switching probability, q, attainable for each of the four molecular switches. We find that for $g_{\max} = 100$ each of the molecular switches can attain a switching probability within 0.05 of any target. It follows that bet–hedging can arise from a simple biological mechanism consisting of stochastic interactions between as few as two or three chemical species, and without the need for large numbers of each of these types. This simplicity could help explain the ubiquitous nature of stochastic phenotype determination and why a large number of genetic factors have been implicated in this process across species; as simple molecular switches suffice they are more likely to arise randomly through genetic mutations.

4.4 Robustness and Redundancy

Next, by introducing redundancy, we demonstrate how bet–hedging can be robust to major perturbations to the underlying network. Figure 4.4(A) shows a version of the DC switch from Figure 4.2 in which the species x and y are duplicated. In this network, which we call DCdup, the set of stable configurations are determined by $x+x' = 0$ or $y+y' = 0$. If we associate the phenotypes A and B with these two configurations respectively then the switching

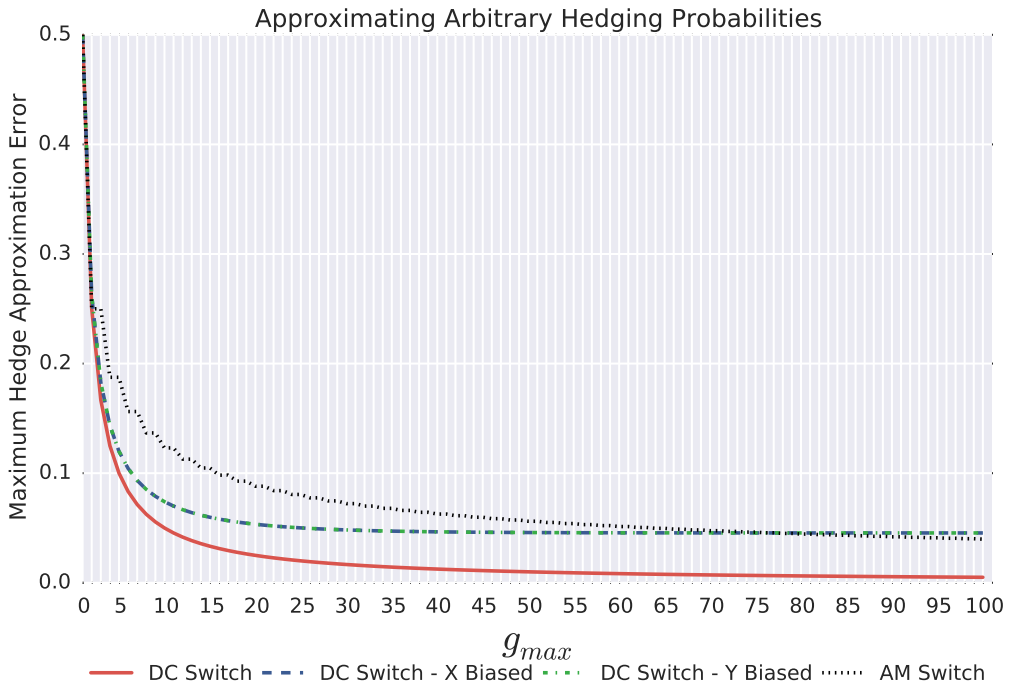

Fig. 4.3 **The CRN model can closely approximate arbitrary switching probabilities.**. For each of the four example bistable CRNs shown in Figure 4.2 and each value of $g_{\max} \in \{1,2,\ldots,100\}$, the maximum error between a target probability, p, of phenotype A and the nearest attainable probability q by the CRN is shown. Here, we see that the CRN molecular switch can closely approximate any bet–hedging strategy with only a small number of molecules (< 100).

probability on initial conditions (x_0, x_0', y_0, y_0') is identical to the switching probability of DC with initial conditions $(x_0 + x_0', y_0 + y_0')$ (a simple mathematical argument to establish this proceeds by symmetry and re–labeling the species). The potential benefit of DCdup is that it maintains its switching properties even if chemical species are removed. Figure 4.4(A) shows numerical solutions for the CRN switching probability when the species x is deleted (middle network) and then when both x and y are deleted (right hand network). These induced CRNs maintain switching behaviour similar to the original network DCdup. The network induced by deleting x (or by symmetry y) behaves precisely as DCdup with initial condition $(0, x_0', y_0, y_0')$ (by symmetry $(x_0, x_0', 0, y_0')$). Further, removing both x and y from DCdup creates a version of the DC switch in the species x' and y' which behaves precisely as the DCdup switch on initial conditions $(0, x_0', 0, y_0')$. It follows that the deletion of chemical species will change the likelihood of switching (and thus the proportion of phenotypes A and B) if initial numbers of all other species remain fixed. An example is shown by the red circles in Figure 4.4(A) where deletion of chemical species shifts the switching probability but does not inhibit bet–hedging entirely. Hence, the DCdup switch is robust to the removal of chemical species — a mutational event that in our model can be interpreted as deletion or down–regulation of a gene.

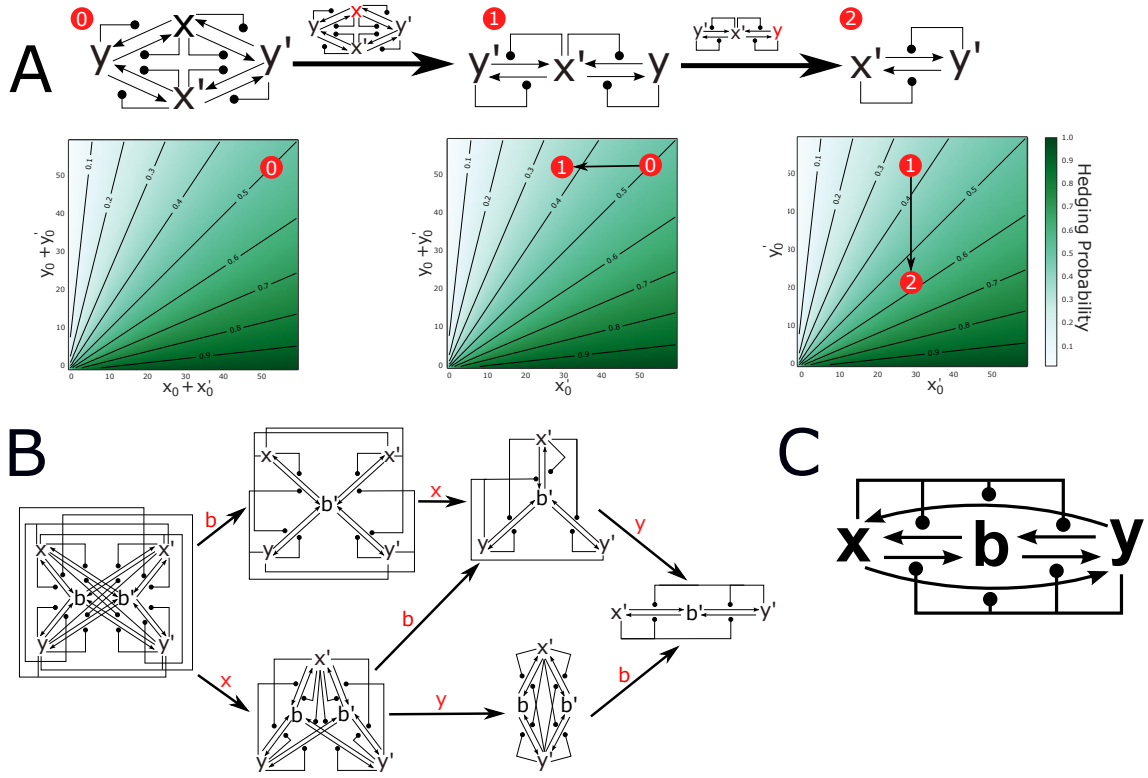

Fig. 4.4 **Redundancy and robustness in the CRN model**. **A)** Redundancy in the chemical reaction network implementing the DC switch maintains molecular switching when chemical species are deleted. Marked in red is the switching probability for initial conditions $(20, 30, 30, 20)$ before deletion (0), after the deletion of x (1) and after the deletion of x and y (2). Contour lines show initial conditions of equal switching behaviour. **B)** Redundancy in the chemical reaction network implementing the AM molecular switch. Switching is maintained if the species x, y and b are removed in any order. We omit the case where y is removed before x due to symmetry. **C)** A molecular switch that can reduce to either AM or DC when specific reactions are inhibited.

Reversing this argument, the switching behaviour of the CRNs in Figure 4.4(A) demonstrate how bet–hedging is robust to gene duplications or upregulating mutations.

A similar implementation of the AM molecular switch is shown in Figure 4.4(B) and is robust to the removal of species, x, y and b in any order. As with the DCdup network, the removal of any chemical species will change the switching probability and shift the proportions of phenotypes in the population but will not inhibit one phenotype entirely. Finally, we note that molecular switches can also be robust to the alteration of reaction rates, including the entire removal of reactions. Each of the switches presented in Figure 4.2 can be derived from a larger CRN shown in Figure 4.4(C) by inhibition of specific reactions. Removal of the autocatalytic reaction $x + y \to x + x$ and $x + y \to y + y$ yields the AM switch. Alternatively, removing the four reactions involving the chemical species b yields the DC switch.

4.5 Implications of the GP–Map for the Evolution of Bet–Hedging

We now turn our attention to the fate of bet–hedging over evolutionary time-scales. Specifically, we highlight how the structure of the CRN determining the GP–mapping determines the likelihood that mutations will fix in a bet–hedging population. We began by presenting a model of population dynamics without mutation, before introducing a long–term evolutionary simulation derived from these population dynamics. Finally, we compare how each of the example CRNs, DC, DCx, DCy and AM, result in different evolutionary fates for bet–hedging where natural selection acts to remove it.

4.5.1 Population Dynamics

Assume a fixed GP–map that determines a probability p of an individual with genotype g having phenotype A at birth. This phenotype is fixed throughout the life of the individual. As we wish to study evolutionary loss of bet–hedging we may assume $p \in (0,1)$, as when $p = 1$ we will end our simulations. We also assume a fixed environment. A discrete time model is used to simulate the population dynamics as follows. Denote by $\mathbf{x}(t) = (x_A(t), x_B(t))^\intercal$ the number of individuals of phenotype A and B at discrete timestep t. We assume Wright–Fisher sampling wherein each individual in the population at time t can contribute any number of individuals to the population at time $t+1$. Denote by w_A (respectively w_B) the expected number of individuals in the population at time $t+1$ that are descended from a single individual of phenotype A (resp. B) present in the population at t. We assume $w_A, w_B > 0$. Each new offspring takes phenotype A with probability p (and B with probability $1-p$) where p is determined by the genotype, g, and the GP–mapping.

For a fixed population size, the dynamics of the population can be modelled using mathematics from the theory of quasispecies that describe a population in mutation–selection balance [353]. Specifically, if the average number of individuals at time $t+1$ produced by a single individual at time t is given by $\langle w \rangle$ the population dynamics are governed by the projection matrix,

$$\mathbf{P} = \frac{1}{\langle w \rangle} \begin{bmatrix} w_A p & w_B p \\ w_A(1-p) & w_B(1-p) \end{bmatrix}. \tag{4.7}$$

The population distribution after one discrete time step is given by $\mathbf{x}(t+1) = \mathbf{P}\mathbf{x}(t)$. As $p \in (0,1)$, the matrix $\mathbf{P}$ is positive and the Perron-Frobenius theorem [63, 205] tells us that the normalized eigenvector corresponding to the dominant eigenvalue of $\mathbf{P}$ gives the long–term stationary distribution of the two phenotypes A and B. As we have only two phenotypes, we can easily determine this dominant eigenvalue as

$$
\begin{aligned}
0 = |\mathbf{P} - I\lambda| &= \begin{vmatrix} \frac{w_A p}{\langle w \rangle} - \lambda & \frac{w_B p}{\langle w \rangle} \\ \frac{w_A(1-p)}{\langle w \rangle} & \frac{w_B(1-p)}{\langle w \rangle} - \lambda \end{vmatrix} \\
&= \left(\frac{w_A p}{\langle w \rangle} - \lambda \right) \left(\frac{w_B(1-p)}{\langle w \rangle} - \lambda \right) - \frac{w_A w_B p(1-p)}{\langle w \rangle^2} \\
&= \lambda \left(\lambda - \frac{w_A p}{\langle w \rangle} - \frac{w_B(1-p)}{\langle w \rangle} \right).
\end{aligned}
$$

Thus, $\lambda = \frac{1}{\langle w \rangle}(pw_A + (1-p)w_B)$. It is easy to verify then that $\mathbf{x}^* = (p, 1-p)$ is the normalized eigenvector corresponding to this eigenvalue as

$$
\mathbf{P} \begin{pmatrix} p \\ 1-p \end{pmatrix} = \frac{1}{\langle w \rangle} \begin{bmatrix} p^2 w_A + p(1-p)w_B \\ p(1-p)w_A + (1-p)^2 w_B \end{bmatrix} \tag{4.8}
$$

$$
= \frac{1}{\langle w \rangle}(pw_A + (1-p)w_B) \begin{pmatrix} p \\ 1-p \end{pmatrix} \tag{4.9}
$$

$$
= \lambda \begin{pmatrix} p \\ 1-p \end{pmatrix} \tag{4.10}
$$

Note that this is also the distribution of phenotypes in the case where the population grows without bound (when the factor of $1/\langle w \rangle$ is omitted). From this stationary distribution we can derive $\langle w \rangle = pw_A + (1-p)w_B$. This phenotype equilibrium will be used in the next section to determine the invasion probability of a different (i.e. mutant) bet–hedging population into an existing one.

In reality, the values w_A and w_B are dependent on a number of stochastic processes, most importantly reproduction and death. Later we will consider the effect of an increased death rate associated with drug treatment, so it is informative to consider the relationship between parameters governing the explicit processes of birth and death and the offspring numbers w_A and w_B. Suppose individuals of type A (respectively B) die with probability d_A (resp. d_B) over each time step. Further, suppose that individuals of phenotype A (respectively B) that survive reproduce with probability f_A (resp. f_B) over the timestep. In this case, the expected number of offspring of an individual with phenotype i over a single time-step is determined by the likelihood of survival multiplied by the expected number of offspring for a surviving individual,

$$
w_i = (1-d_i)(1+f_i) \text{ for } i \in \{A, B\}. \tag{4.11}
$$

Note that multiple pairs of values for f_i and d_i can yield the same w_i. The population dynamics we present here, and the invasion dynamics presented below, are identical for all such pairs.

It is worth noting that, for a genotype g corresponding to a probability $p(g)$, the population dynamics are determined entirely by $p(g)$ and are independent of the shape of the GP–mapping.

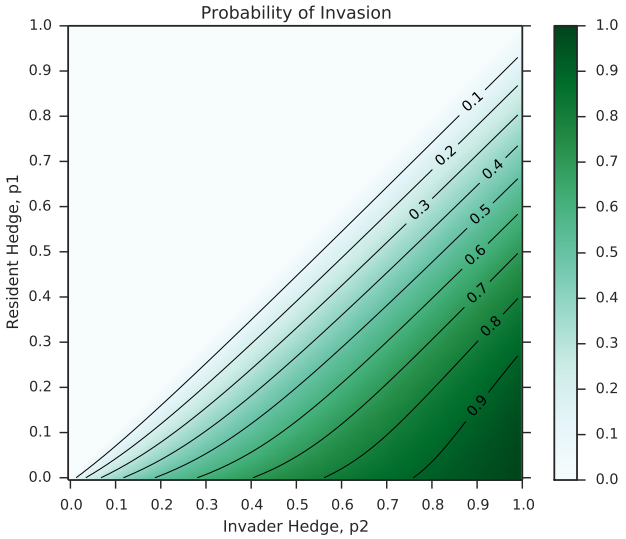

Fig. 4.5 **Invasion probabilities for a mutant with bet–hedging probability p_1 into a resident population of phenotype p_2.** The parameters are those presented in Figure 4.2. Note that deleterious and neutral mutations cannot fix under our model of invasion dynamics, hence invasion in the case $p_1 \leq p_2$ (above the antidiagonal of the plot) is impossible.

It is only when we consider the difference in dynamics between populations of two different genotypes that the form of the GP–map begins to reveal itself. This scenario bears a striking resemblence to the property of GRNs (Section 2.2.4), which we motivated by considering the Spemann Organiser in embryonic development [51], wherein both the genotype and phenotype of two organisms can be identical but the GP–maps can differ.

4.5.2 Invasion Dynamics

Our aim is to determine the long–term evolutionary trajectories of a population of cells endowed with different GP–maps. It is intractable to determine these trajectories through explicit simulation of the entire population, so instead we derive an analytic solution for the probability of a mutant genotype invading an existing isogenic population. Consider a large fixed–size population and assume that mutation is sufficiently rare (explicitly that the mutation rate u and population size N satisfy $Nu\log N \ll 1$) that we may consider strong–selection weak mutation (SSWM) evolutionary dynamics (Section 3.1.1). Under these assumptions we can assume that the population is isogenic and that each time a new mutant appears in the population that mutant either fixes as the new population genotype or becomes extinct.

Suppose a single mutant of genotype g' arises in an isogenic population of genotype g and denote by $\pi(g')$ the probability that this mutant reaches fixation as the population genotype. This probability is dependent on the phenotype of this initial mutant and is given by,

$$\pi(g') = p(g')\pi(g'|A) + (1 - p(g'))\pi(g'|B). \tag{4.12}$$

Denote $\pi = (\pi(g'|A), \pi(g'|B))$ and suppose that the population size, N, is sufficiently large that we may approximate it by the limit $N \to \infty$. As we assume Wright–Fisher sampling for reproduction, the value of π can be determined from the theory of branching processes. In particular, π can be calculated numerically as the solution to the equation

$$1 - \pi = e^{-\mathbf{P}\pi} \tag{4.13}$$

where $\mathbf{P}$ is the matrix governing the population dynamics defined above. A proof of this identity, modified from the theory of viral quasispecies [352], proceeds as follows.

Under Wright–Fisher sampling the probability that a randomly chosen individual in the next generation is the offspring of a given individual of phenotype $i \in \{A, B\}$ in the current generation is $\frac{w_i}{\langle w \rangle N}$. Thus, the probability that an individual is the offspring of a particular parent of phenotype i and genotype g and also has phenotype r is

$$\xi_r = \frac{w_i \mathbb{P}(g \xrightarrow{\text{GP}} r)}{\langle w \rangle N} \tag{4.14}$$

as not all offspring will have phenotype r. It follows that the probability that a given individual of phenotype i has precisely k_r offspring of phenotype r in the next generation is given by

$$p(k_r|i) = \binom{N}{k_r} \xi_r^{k_r} (1 - \xi_r)^{N - k_r}. \tag{4.15}$$

We can extend this argument to both phenotypes. The probability that an individual of phenotype i has k_A offspring of phenotype A and k_B offspring of phenotype B is

$$p(k_A, k_B|i) = \frac{N!}{k_A! k_B! (N - k_A - k_B)!} \xi_A^{k_A} \xi_B^{k_B} (1 - \xi_A - \xi_B)^{N - k_A - k_B} \tag{4.16}$$

$$= \frac{N!}{k_A! k_B! (N - k_A - k_B)!} \frac{\mathbf{P}_{i1}}{N}^{k_A} \frac{\mathbf{P}_{i2}}{N}^{k_B} \times \left(1 - \frac{\mathbf{P}_{i1}}{N} - \frac{\mathbf{P}_{i2}}{N}\right)^{N - k_A - k_B} \tag{4.17}$$

where the second equality follows from the definition of $\mathbf{P}$. Assume now that the population size, N, is sufficiently large that we may approximate $p(k_A, k_B|i)$ by taking the limit as $N \to \infty$. This limit is a bivariate Poisson distribution

$$p(k_A, k_B|i) = \frac{\mathbf{P}_{i1}^{k_A}}{k_A!} \times \frac{\mathbf{P}_{i2}^{k_B}}{k_B!} \times e^{-\mathbf{P}_{i1} - \mathbf{P}_{i2}}. \tag{4.18}$$

In the following derivation we will use the theory of branching processes (Harris [140]). This method, along with the assumption of infinite population size, restrict the resulting theory to fixation events of mutations which increase the average fitness of the population. Denote by χ_i the probability that the lineage of a single mutant individual with phenotype i in a population of average fitness $\langle w \rangle$ becomes extinct after finitely many generations. From the theory of branching processes we know that the vector of extinction probabilities $\chi = (\chi_A, \chi_B)$

satisfies $\chi = \mathbf{f}(\chi)$ where $\mathbf{f}(\mathbf{z}) = (f_A(\mathbf{z}), f_B(\mathbf{z}))$ is the probability generating function of the offspring probabilities $p(k_A, k_B|i)$ given by,

$$f_i(\mathbf{z}) = \sum_{k_A}\sum_{k_B} p(k_A, k_B|i) z_1^{k_A} z_2^{k_B}. \tag{4.19}$$

Substituting equation 4.18 into equation 4.19 we have

$$f_i(\chi) = \sum_{k_A}\sum_{k_B}\left(\frac{\mathbf{P}_{i1}^{k_A}}{k_A!} \times \frac{\mathbf{P}_{i2}^{k_B}}{k_B!} \times e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}}\right)\chi_1^{k_A}\chi_2^{k_B} \tag{4.20}$$

$$= e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}} \sum_{k_A}\sum_{k_B}\left(\frac{\mathbf{P}_{i1}^{k_A}}{k_A!}\chi_1^{k_A}\right)\left(\frac{\mathbf{P}_{i2}^{k_B}}{k_B!}\chi_2^{k_B}\right) \tag{4.21}$$

$$= e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}} \sum_{k_A}\left(\frac{(\mathbf{P}_{i1}\chi_1)^{k_A}}{k_A!}\right)\sum_{k_B}\left(\frac{(\mathbf{P}_{i2}\chi_2)^{k_B}}{k_B!}\right) \tag{4.22}$$

$$= e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}} \times e^{\mathbf{P}_{i1}\chi_1} \times e^{\mathbf{P}_{i2}\chi_2} \tag{4.23}$$

$$= e^{\mathbf{P}_{i1}(\chi_1-1)+\mathbf{P}_{i2}(\chi_2-1)}. \tag{4.24}$$

Taking the convention that $e^{\mathbf{x}} = (e^{x_1}, \ldots, e^{x_n})$ this gives

$$\mathbf{f}(\chi) = e^{\mathbf{P}(\chi-\mathbf{1})}. \tag{4.25}$$

Now note that the probability of fixation of the mutant individual with phenotype i is precisely the probability that its lineage does not go extinct in finitely many generations, $\pi = \mathbf{1} - \chi$. Hence

$$1 - \pi = \mathbf{f}(1-\pi) = e^{-\mathbf{P}\pi}. \tag{4.26}$$

As the value of $\langle w \rangle$ can be calculated from the population dynamics described above, this equation can be solved numerically. Figure 4.5 shows an example heatmap of invasion probability for an invader with genotype corresponding to probability p_2 of phenotype A into a resident population with probability p_1. The parameters are $w_A = 2.0$, $w_B = 1.01$. Heatmaps of invasion probability in terms of population genotypes are presented in Figure 4.6(C).

4.5.3 The Genotype–Phenotype Map Preserves Bet–Hedging

Each molecular switch we study has the ability to (approximately) produce any arbitrary switching probability between phenotype A and B, as shown in Section 4.3.3. For fixed gene expression, the specific switching mechanism responsible for producing multiple phenotypes is irrelevant as the proportion of phenotypes remains fixed. However, over longer time scales the structure of the molecular switch has a significant impact on the evolution of bet–hedging. We now consider the rate at which evolutionary loss of the bet–hedging occurs in a fixed

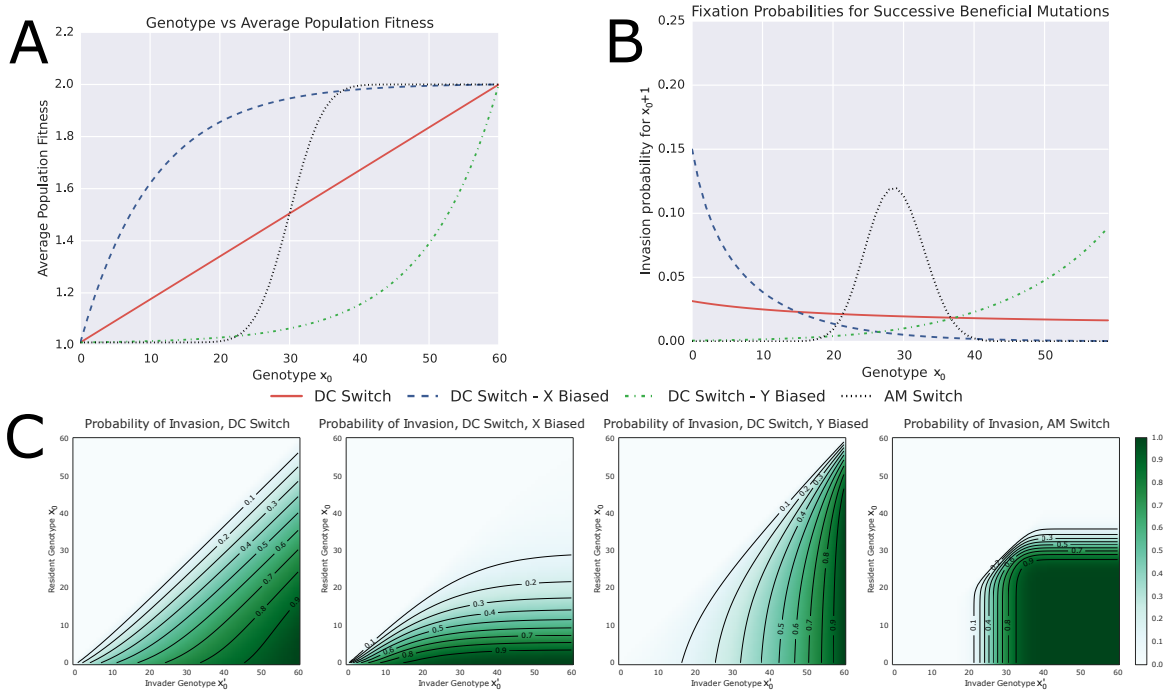

Fig. 4.6 **The switch–dependent relationship between genotype and fitness/invasion probability. A)** The relationship between population genotype, x_0, and average population fitness for each molecular switch. These values are determined by solving for λ in Equation 4.10 using switching probability p derived from the initial genotype x_0 and the associated molecular switch as outlined in Figure 4.2. **B)** Invasion probabilities for a new genotype x_0+1 into a resident population of genotype x_0. The values are determined by solving Equation 4.26 wherein the matrix **P** takes switching probability p defined by x_0 and the molecular switch. **C)** Invasion probabilities for resident and invader genotypes. The values are determined by solving Equation 4.26 wherein the matrix **P** takes switching probability p defined by x_0 and the molecular switch.

environment strongly favouring one phenotype (A) over a second (B) to demonstrate the importance of the GP–mapping on predicting the evolutionary fate of bet–hedging.

We simulate evolution by assuming a large asexually reproducing bet–hedging population exists under Strong Selection Weak Mutation (SSWM) dynamics (Section 2.1.2). The population is assumed to be isogenic and the population genotype is periodically replaced by a fitter mutant (Figure 4.7 shows a schematics of this model, Algorithm 1 shows the implementation in pseudocode). The invasion probability is calculated by numerical solution of Equation 4.26. To avoid issues regarding the phenotypic/fitness implications of switching on fewer molecules, we enforce that x_0+y_0 remains constant ($x_0+y_0 = g_{\max} = 60$ for computational efficiency) throughout our simulations. Thus, the mutation relation μ now specifies only two possible mutations: simultaneously incrementing x (or y) and decrementing y (resp. x). Under this assumption the genotype is then entirely determined by the value of $x_0 \in \{0,\ldots,60\}$.

Throughout the remainder of this work we take the two phenotypes A and B in our model to correspond to a high fitness, proliferative phenotype and a low fitness, slow proliferating phenotype respectively, to mirror the phenomenon of bacterial persistence. The invasion

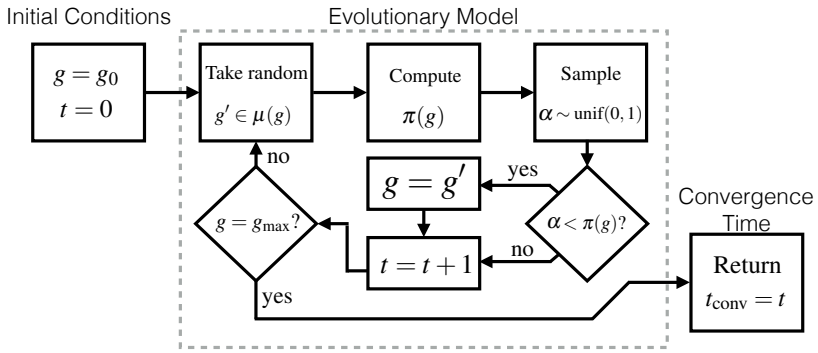

Fig. 4.7 **Schematic of the evolutionary model.** This model is used to stochastically determine a time (measured in mutational events) until bet–hedging is completely lost by the population. Pseudocode is presented in Algorithm 1.

probability is computed independently of the (assumed to be large) population size and is dependent only on the (stable) distribution of phenotypes in the population at equilibrium. For this reason we need only consider the relative fitnesses of the phenotypes. This distribution can be computed from the values w_A and w_B. We parametrise the model determined as follows. The discrete timesteps are taken to be $t = 1$ h (hours), mirroring an expected division time of *E. coli*. The time from birth until reproduction for a phenotype A individual is exponentially distributed with rate parameter $\lambda_A = 1.0\ \text{h}^{-1}$. The number of reproductions of a single individual per timestep is then Poisson distributed with parameter λ_A. Hence on average, each phenotype A individual present at time t reproduces once per timestep and produces two individuals at time $t+1$ (itself and its offspring). It follows, $w_A = 2.0\ \text{h}^{-1}$. The reproductive rate of persister–type cells is unknown. To match their behaviour qualitatively, we take our persister–like phenotype B cells to have an expected reproduction time 100 times slower than phenotype A individuals, 100 h. Thus, reproduction time for phenotype B cells is distributed exponentially with rate parameter $\lambda_B = 0.01\ \text{h}^{-1}$. The expected number of reproductions over the one hour time step is then 0.01 and $w_B = 1.01\ \text{h}^{-1}$. The environment remains fixed (w_A, w_B are unchanged) throughout the evolutionary simulation. Although this parametrisation is only an approximation it is sufficient as an illustrative model demonstrating the importance of the structural properties of the genotype–phenotype mapping.

Figure 4.6(A) shows how changes in the population genotype manifest themselves as changes in the average population fitness. The expected population fitness increase associated with a mutation from x_0 to $x_0 + 1$ is not constant and instead is dependent on the underlying molecular switch. As a result, invasion probabilities are dependent on both the population genotype, invading genotype, and the form of the molecular network. Figure 4.6(C) shows the probability of a single mutant genotype x_0' invading a resident population of genotype x_0. Note that only mutations that increase the proportion of phenotype A are beneficial and hence, as our invasion probabilities are determined from the theory of branching processes, are the only mutations that can fix. Figure 4.6(B) shows the probability of successive beneficial mutations, $x_0 + 1$, invading a resident population of genotype x_0. In this figure we see the impact of the

Algorithm 1 Stochastic loss of bet–hedging - Pseudocode for the algorithm presented in Figure 4.7. †Numerical solution can be performed through a third party package. In this case, SciPy [167].

procedure SIMULATEBHLOSS($g_0, g_max, \mathscr{R}_{\text{GP}}$)
    $t \leftarrow 0$
    $g \leftarrow g_0$
    g_list $\leftarrow$ []
    p_list $\leftarrow$ []
    while $g \neq g_{\max}$ **do**
        g_list.append(g)
        $p \leftarrow \mathbb{P}\left((g, g_{\max} - g) \overset{\text{GP}}{\rightarrow} A\right)$ (Equations 4.4-4.6 or Appendix B)
        p_list.append(p)
        Sample $\alpha \sim \text{Uniform}[0,1]$
        $g' \leftarrow g+1$ if $\alpha < 0.5$ else $g \leftarrow g-1$ (Equation 4.2)
        $q \leftarrow \mathbb{P}\left((g', g_{\max} - g') \overset{\text{GP}}{\rightarrow} A\right)$ (Equations 4.4-4.6 or Appendix B)
        Construct $\mathbf{P}(p)$ as Equation 4.7
        Numerically solve $1 - \pi = e^{-\mathbf{P}\pi}$ for $\pi = (\pi(g'|A), \pi(g'|B))$ (Equation 4.26)†
        Sample $\alpha \sim \text{Uniform}[0,1]$
        if $\alpha < q\pi(g'|A) + (1-q)\pi(g'|B)$ **then**
            $g \leftarrow g'$
        end if
        $t \leftarrow t+1$
    end while
    return g_list, p_list, t
end procedure

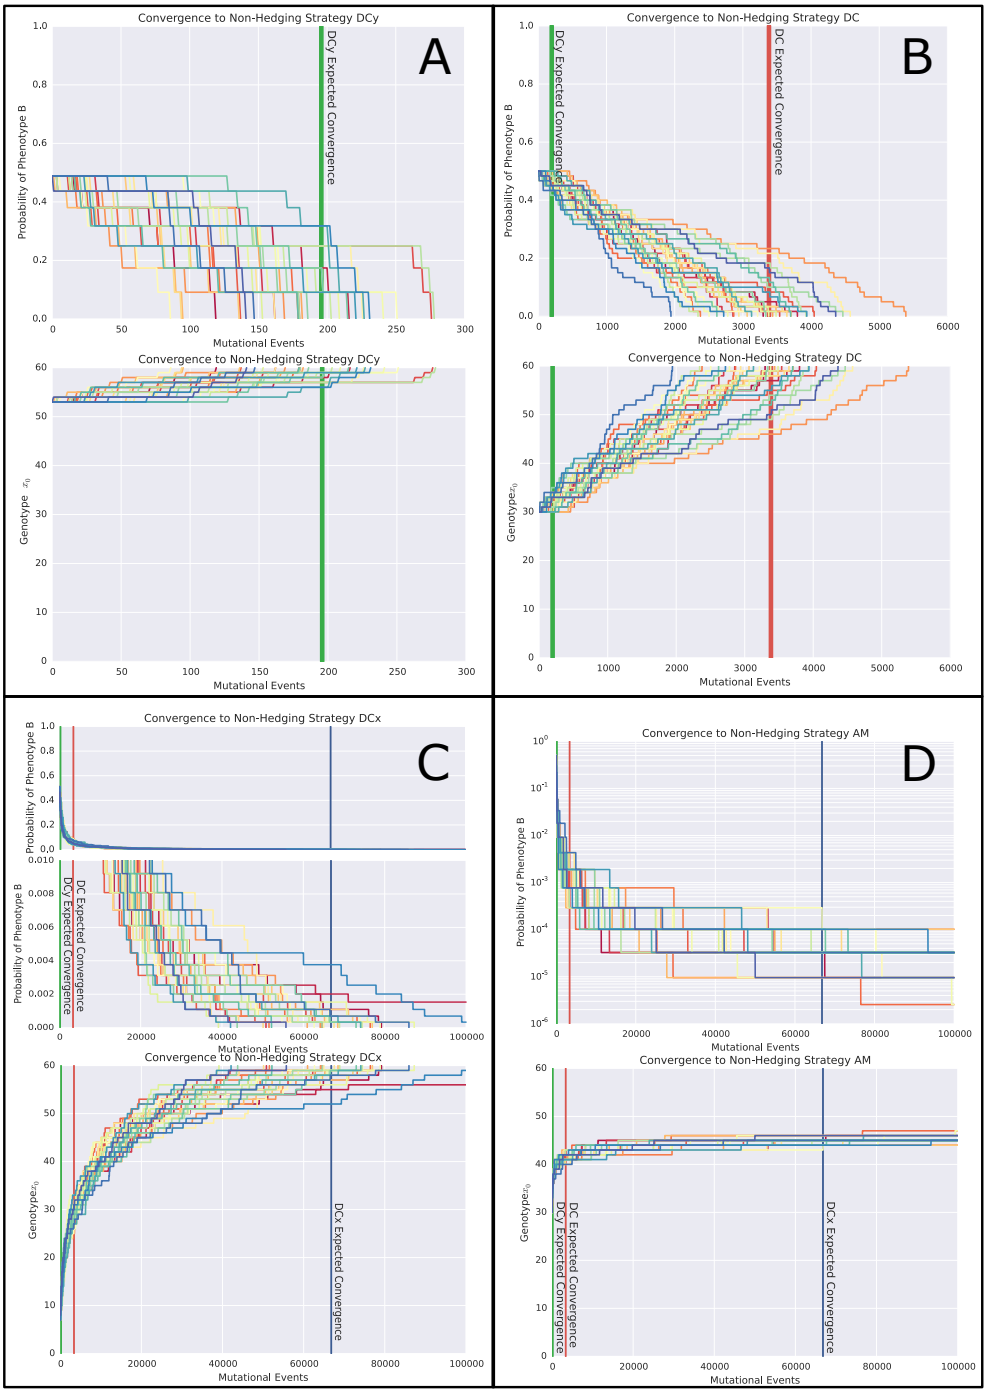

Fig. 4.8 **Convergence dynamics through genotype and probability space for the GP–maps defined by DC, DCx, DCy and AM.** 30 stochastic realisations of the evolutionary simulation through both genotype and probability space are shown for **A)** The DCy switch. **B)** The DC switch. **C)** The DCx switch. **D)** The AM switch. Due to the rapid initial change in hedging probability for the DCx switch, the convergence dynamics are also shown on a restricted scale. As the probability of phenotype *B* rapidly approaches 0 in the AM switch simulation but never converges, the dynamics are shown on a logarithmic scale. The expected convergence time for the DCy switch is marked in green, for the DC switch in marked in red and for the DCx switch is marked in blue.

genotype–phenotype map on the evolutionary dynamics. For the DC, DCx and AM switches the probability of the next beneficial mutation fixing reduces for each successive mutation. The magnitude of this decrease is dependent on the switch and in the case of DCx and AM approaches zero. Conversely, for the DCy switch each successive mutation is more likely to fix.

Consider the evolutionary trajectories of phenotype proportions determined by each of the molecular switches from an initial gene expression profile corresponding closely to a population consisting of approximately 50% of each phenotype. As the DC and AM switches are symmetric the genotype corresponding to a 0.5 hedging probability is $x_0 = 30$. For DCx the closest genotype to a 50% hedge is $x_0 = 7$ which corresponds to a probability of 0.49. For DCy the closest genotype to a 0.5 hedging probability is $x_0 = 53$ and corresponds to a probability of 0.51. As deleterious and neutral mutations cannot fix under our model of population dynamics, the population genotype will be periodically incremented until $x_0 = 60$ at which point bet–hedge is lost. Figure 4.8 shows the evolutionary trajectories towards this loss of bet–hedging, highlighting considerably different convergence dynamics. Here, time is measured in discrete units representing a mutational event. By taking this unit of time we need not state an explicit mutation rate.

For the DC, DCx and DCy switches the expected convergence times can be determined as the expectation of a sum of non–identical independent geometric distributions. We find that the expected number of mutational events required for a complete loss of bet–hedging, as highlighted in Figure 4.8, are given by

$$\begin{aligned}
&\mathbb{E}[\text{time to loss of bet–hedge} \,|\, \text{DCy switch}, x_0 = 53] = 195.8 \\
&\mathbb{E}[\text{time to loss of bet–hedge} \,|\, \text{DC switch}, x_0 = 30] = 3382.8 \\
&\mathbb{E}[\text{time to loss of bet–hedge} \,|\, \text{DCx switch}, x_0 = 7] = 68000.1 \\
&\mathbb{E}[\text{time to loss of bet–hedge} \,|\, \text{AM switch}, x_0 = 30] = \infty
\end{aligned}$$

where in each case time is measured in units of mutational events.

In the case of the AM network, each subsequent mutation provides a diminishing increase in fitness until mutations are approximately neutral. The probability of neutral mutations fixing within our model of invasion dynamics, which models the population size as tending to infinity, is zero. In reality, the actual convergence times in the AM will depend on the population size. For large populations, as is our assumption, the time scales will be sufficiently long that we take it as equivalent to the evolutionary trajectory never converging [2].

[2]This assumption can be justified by the observation that over these time scales either unmodelled mutations (such as mutations to the GP–map itself, to other genes governing the phenotypes A and B, or to other aspects of the phenotype) or unmodelled changes in the environment or ecosystem will occur, rendering our model unsuited to the situation.

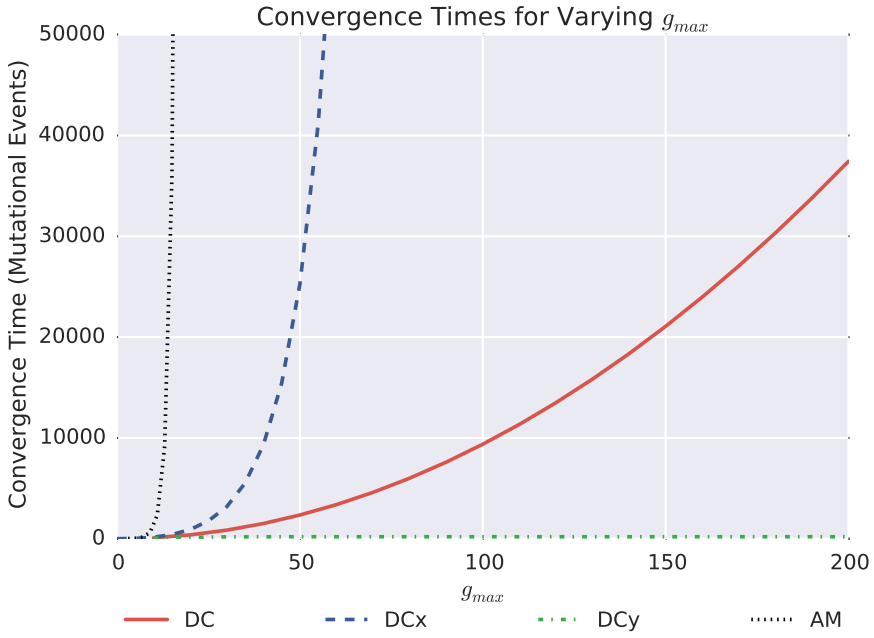

Fig. 4.9 **Comparison of convergence times for increasing $g_{\max}$.** The expected evolutionary convergence time to loss of bet–hedging is shown for each switch and for increasing $g_{\max}$. Here, we see that the large difference in convergence time between switches is maintained as $g_{\max}$ varies, and in fact is amplified as $g_{\max}$ grows large.

4.5.4 Robustness to Parameter Variation

The parameters $g_{\max}$, w_A and w_B used to derive our results take specific values chosen for reasons of computational efficiency (for $g_{\max}$) or to mirror qualitative behaviour (for w_A and w_B). Here, we briefly demonstrate the robustness of our results to variations in these parameters. Figure 4.9 shows the effect of increasing $g_{\max}$ on evolutionary convergence time (in the sense of Figure 4.8). For each switch, the convergence time increases monotonically with $g_{\max}$ and the difference in convergence time between the different molecular switches, on which our results are predicated, holds as $g_{\max}$ increases. Figure 4.10 shows the effect of increasing w_A on convergence time when $w_B = 1.01$ h^{-1} is fixed. Again, the differing dynamics of evolutionary convergence are preserved for increasing w_A, although convergence time is monotonically decreasing as w_A increases. It follows that our results concerning differing rate of evolutionary loss of bet–hedging are robust both to changes in the number of molecules comprising the molecular switch, $g_{\max}$, and the relative fitnesses for the two phenotypes, w_A and w_B.

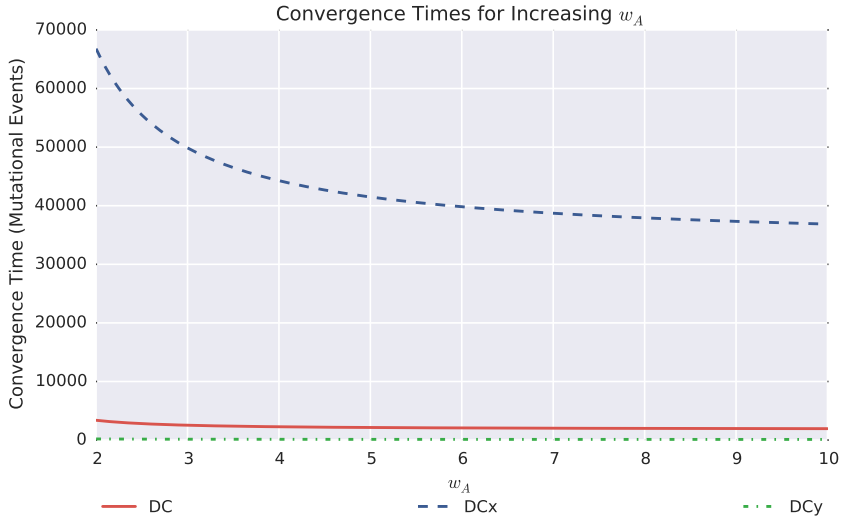

Fig. 4.10 **Comparison of convergence times for increasing w_A, $w_B = 1.0$.** The time to evolutionary loss of bet–hedging is plotted against varying w_A with $g_{\max} = 60$ for each curve. The AM network dynamics are omitted as convergence does not occur. Although the convergence time is monotonically decreasing as w_A increases, the large differences in convergence time are preserved.

4.6 Implications for Sequential Therapies and Treatment Holidays

Recently, breaks from treatment or *treatment holidays* have been suggested as a potential therapeutic strategy to slow the emergence of drug resistance arising due to the ecological principle of competitive release. At present, the efficacy, or potential pitfalls, of treatment holidays in diseases with bet–hedging–driven resistance have not been explored. To demonstrate the importance of the underlying molecular switch on the efficacy of treatment holidays for diseases with bet–hedging–driven resistance, we performed a two part simulation comprising a simulation of the evolutionary loss of bet–hedging followed by an individual based model of extinction under drug therapy.

First, an evolutionary simulation similar to the one described above is performed to determine a population genotype following a treatment holiday (Algorithm 2 – GetExpectedGenotype). From an initial genotype g_0 with $p(g_0) \approx 0.5$ the expected population genotype g following a treatment holiday of length T (measured in mutational events) is computed. The expected time, in mutational events, taken for the genotype $g+1$ to arise by mutation and fix in the population is $\frac{1}{(0.5\pi(g+1))}$. Here, the factor of 0.5 arises as only half of mutations are beneficial ($g \to g-1$ is not). Using this fact, an expected post–holiday genotype can be easily determined by repeatedly incrementing g while keeping a sum of the expected number

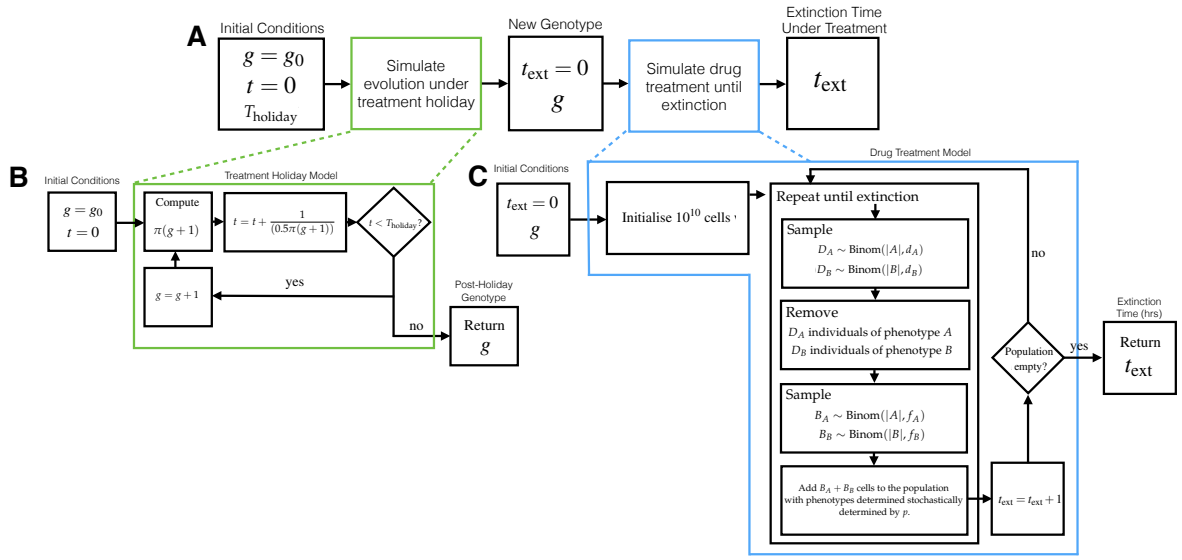

Fig. 4.11 **The model for simulating treatment holidays. A)** The overall model consists of first determining a post–treatment holiday genotype, g, and then simulating drug treatment on a population with bet–hedging determined by that genotype. **B)** The expected post–holiday genotype is determined from the invasion probabilities $\pi(g)$ which are in turn dependent on the molecular switch. **C)** A stochastic death–birth process is used to determined an extinction time t_{ext}. Pseudocode for this simulation is presented in Algorithm 2.

of mutational events required for that new g to fix in the population. This process is terminated when this number of events exceeds T.

Using this post–holiday genotype, and the associated GP–map–dependent hedging probability $p(g)$, a stochastic death–birth process without mutation was performed to determine an approximate time to extinction (Algorithm 2 - AgentBasedModel). A population of size 10^{10} was initialized with a proportion $p(g)$ of phenotype A and $(1-p(g))$ of phenotype B. The simulation proceeds as depicted in Figure 4.11 and is terminated when the population is extinct or 20,000 simulated hours has elapsed. Each time step corresponds to an hour. Unlike the simulation of evolutionary time–scales, the treatment simulation is dependent on the parameters of per–hour probability of death, d_i^{drug}, and reproduction f_i^{drug}, for individuals of each phenotype in the drug–treated environment. The full simulation was performed 2000 times for each treatment holiday length $T \in \{0, 3000, 5000, 50000, 100000\}$ (mutational events) and each molecular switch from Figure 4.2. For each pair of T and molecular switch, the 2000 extinction times are presented as histograms.

The parameter values used in the evolutionary simulation of treatment holidays were chosen to coincide with the values $\mathbf{w} = [w_A^{\text{hosp}}, w_B^{\text{hosp}}] = [2.0\ \text{h}^{-1}, 1.01\ \text{h}^{-1}]$ used in our simulation of evolutionary loss of bet–hedging. For the simulation of treatment, explicit values of f_i^{drug} and d_i^{drug} (for $i \in \{A, B\}$) are required. First, we take the values for the non–drug environment that correspond to those used in the evolutionary simulation. For the hospitable environment, we assume a fixed death probability for both As and Bs of $d_A^{\text{hosp}} = d_B^{\text{hosp}} = 0.005\ \text{h}^{-1}$. The per-hour likelihood of reproduction for the surviving individuals are taken as $f_A^{\text{hosp}} = 1.0\ \text{h}^{-1}$

Algorithm 2 Pseudocode for the simulation of drug treatment following a treatment holiday. A schematic for this algorithm is presented in Figure 4.11.†Numerical solution can be performed through a third party package. In this case, SciPy [167].

procedure GETEXPECTEDGENOTYPE($g_0, g_{\max}, T, \mathscr{R}_{\text{GP}}$)
  $g \leftarrow g_0$
  $p \leftarrow \mathbb{P}\left((g, g_{\max}-g) \overset{\text{GP}}{\to} A\right)$ (Equations 4.4-4.6 or Appendix B)
  $t \leftarrow 0$
  while $t < T$ **do**
    $g' \leftarrow g+1$
    $q \leftarrow \mathbb{P}\left((g', g_{\max}-g') \overset{\text{GP}}{\to} A\right)$ (Equations 4.4-4.6 or Appendix B)
    Construct $\mathbf{P}(p)$ as Equation 4.7
    Numerically solve $1-\pi = e^{-\mathbf{P}\pi}$ for $\pi = (\pi(g'|A), \pi(g'|B))$ (Equation 4.26)†
    $p_{\text{fix}} \leftarrow \frac{1}{2}(q\pi(g'|A) + (1-q)\pi(g'|B))$ (Factor $\frac{1}{2}$ accounts for $g \to g-1$ mutation)
    $t \leftarrow t + \frac{1}{p_{\text{fix}}}$
    $p \leftarrow q$
    $g' \leftarrow g$
  end while
  return g
end procedure
procedure AGENTBASEDMODEL($N_0, g_0, g_{\max}, f_A, f_B, d_A, d_B$)
  $q \leftarrow \mathbb{P}\left((g_0, g_{\max}-g_0) \overset{\text{GP}}{\to} A\right)$ (Equations 4.4-4.6 or Appendix B)
  $t_{\text{ext}} = 0$
  $A \leftarrow \lceil q * N_0 \rceil$
  $B \leftarrow \lfloor (1-q) * N_0 \rfloor$
  while $(A+B) \neq 0$ **do**
    Sample $D_A \sim \text{Binomial}(A, d_A)$
    Sample $D_B \sim \text{Binomial}(B, d_B)$
    $A \leftarrow A - D_A$, $B \leftarrow B - d_B$
    Sample $B_A \sim \text{Binomial}(A, f_A)$
    Sample $B_B \sim \text{Binomial}(B, f_B)$
    $A \leftarrow \lceil qB_A \rceil + \lceil qB_B \rceil$
    $B \leftarrow \lfloor (1-q)B_A \rfloor + \lfloor (1-q)B_B \rfloor$
    $t_{\text{ext}} \leftarrow t_{\text{ext}} + 1$
  end while
  return t_{ext}
end procedure
procedure TREATMENTHOLIDAYSIM($g_0, g_{\max}, \mathscr{R}_{\text{GP}}, T, N_0, f_A, f_B, d_A, d_B$)
  $g \leftarrow$ GetExpectedGenotype$(g_0, g_{\max}, T, \mathscr{R}_{\text{GP}})$
  $t_{\text{ext}} \leftarrow$ AgentBasedModel$(N_0, g_0, g_{\max}, f_A, f_B, d_A, d_B)$
  return t_{ext}
end procedure

and $f_B^{\text{hosp}} = 0.015\ \text{h}^{-1}$. These parameters correspond to overall expected offspring for each phenotype

$$\begin{aligned}[w_A^{\text{hosp}}, w_B^{\text{hosp}}] &= [(1-d_A^{\text{hosp}})(1+f_A^{\text{hosp}}), (1-d_B^{\text{hosp}})(1+f_B^{\text{hosp}})] \\ &\approx [2.0, 1.01].\end{aligned}$$

We consider three different regimes of drug–environment parameter sets corresponding to a purely cytostatic drug, affecting only f_A, a purely cytotoxic drug, affecting only d_A and drug that is a mixture of the two, affecting both f_A and d_A intermediately. In the following discussion the 'drug' superscript is omitted for readability. For each regime we assume that the parameters for phenotype B, mirroring a drug–impervious persister–like phenotype, are unchanged in the presence of the drug.

A standard course of antibiotic treatment lasts between 7 and 10 days. Using this period as a guide, we defined extinction times of less than 240 hrs as corresponding to a potentially successful course of therapy, and thus, a viable treatment holiday strategy. The extinction time histograms lying within this period are shaded green and marked by a $*$ in Figure 4.13. Mutations were not modelled during the simulations of therapy.

The cytostatic drug parameter set was taken to correspond to a substantial reduction in the proliferative rate of phenotype A individuals. We note that extinction will almost surely never occur if

$$\begin{aligned}\langle w^{\text{drug}} \rangle &= p(g) w_A^{\text{drug}} + (1-p(g)) w_B^{\text{drug}} \\ &= p(g)(1-d_A)(1+f_A) + (1-p(g))(1-d_B)(1+f_B) \\ &\geq 1.\end{aligned}$$

As $(1-d_B)(1+f_B) > 1.0\ \text{h}^{-1}$, it follows that a cytostatic drug can only drive extinction in this case if $(1-d_A)(1+f_A) < 1.0\ \text{h}^{-1}$. As $d_A = 0.005\ \text{h}^{-1}$, this occurs only if $f_A < 0.995^{-1} - 1 \approx 0.005\ \text{h}^{-1}$. For our simulation of cytostatic drug treatment we assume that the drug is entirely effective at inhibiting reproduction, giving $f_A = 0\ \text{h}^{-1}$. For this parameter, we find that extinction never occurs within the 240 hrs time period, indicating that regardless of the underlying switching mechanism no length of treatment holiday will result in complete cure with a cytostatic drug. The histograms of extinction times, where extinction does occur, are presented in Figure 4.12.

For the cytotoxic drug parameter set we consider an increased death probability of $d_A = 0.8\ \text{h}^{-1}$, holding all other parameters the same. The histograms of extinction times are presented in Figure 4.13. Here, we see that the length of treatment holiday necessary for a successful (<240hr) follow up course of therapy varies by orders of magnitude depending on the underlying switch. It follows that the efficacy of treatment holidays as a potential therapeutic intervention for a disease with bet–hedging–driven resistance is dependent on the underlying driver of bet–hedging.

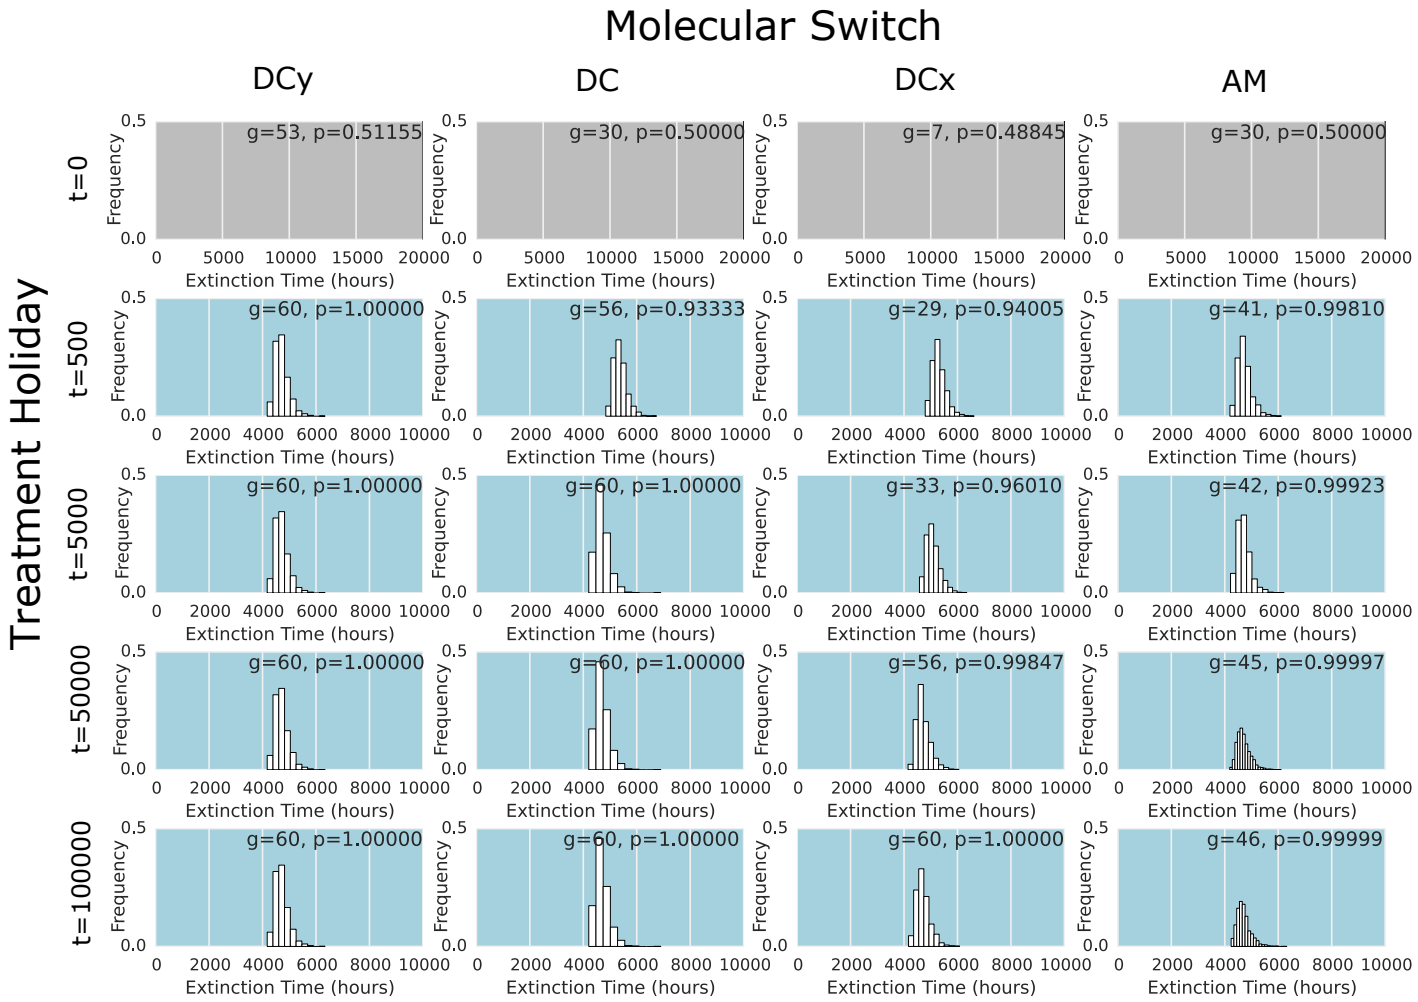

Fig. 4.12 **Extinction dynamics for populations endowed with different molecular switches over different timescales of treatment holidays.** Each histogram shows the distribution of extinction times over 2000 simulations of treatment in an individual–based model with drug parameters corresponding to a cytostatic drug. The molecular switch used as the GP–map is shown as the column heading. The genotype and associated probability of phenotype A (shown inset to each subfigure) are determined by an evolutionary simulation of a treatment holiday for a timescale determined by the row. A blue background indicates extinction times longer than a timeframe viable for an antibiotic treatment, a green background (or inset star) indicates extinction times within this timeframe.

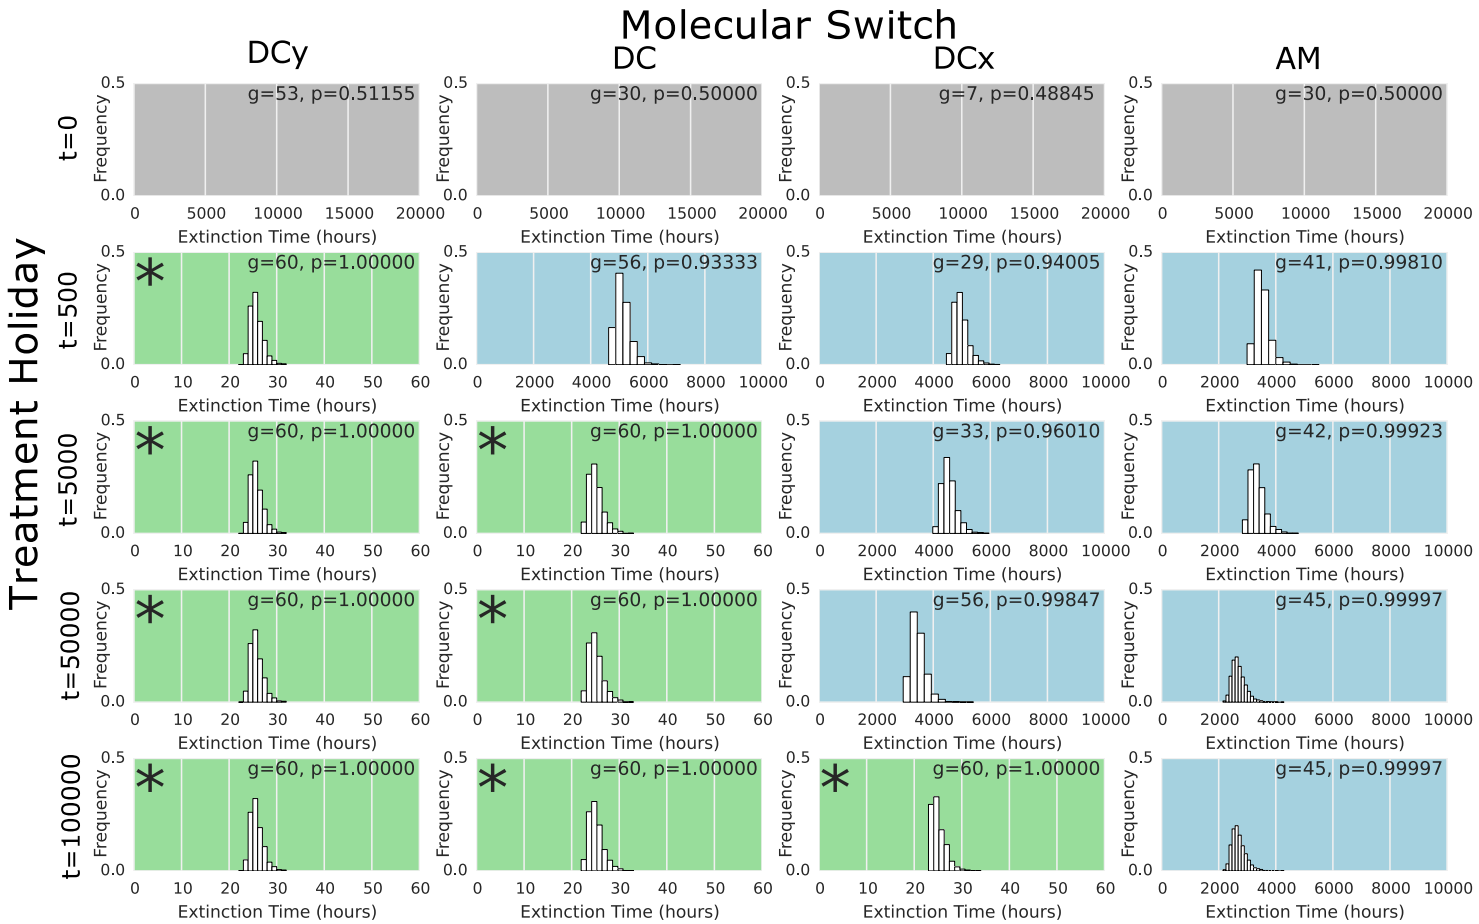

Fig. 4.13 **Extinction dynamics for populations endowed with different molecular switches over different timescales of treatment holidays.** Each histogram shows the distribution of extinction times over 2000 simulations of treatment in an individual–based model with drug parameters corresponding to a cytotoxic drug. The molecular switch used as the GP–map is shown as the column heading. The genotype and associated probability of phenotype A (shown inset to each subfigure) are determined by an evolutionary simulation of a treatment holiday for a timescale determined by the row. A blue background indicates extinction times longer than a timeframe viable for an antibiotic treatment, a green background (or inset star) indicates extinction times within this timeframe.

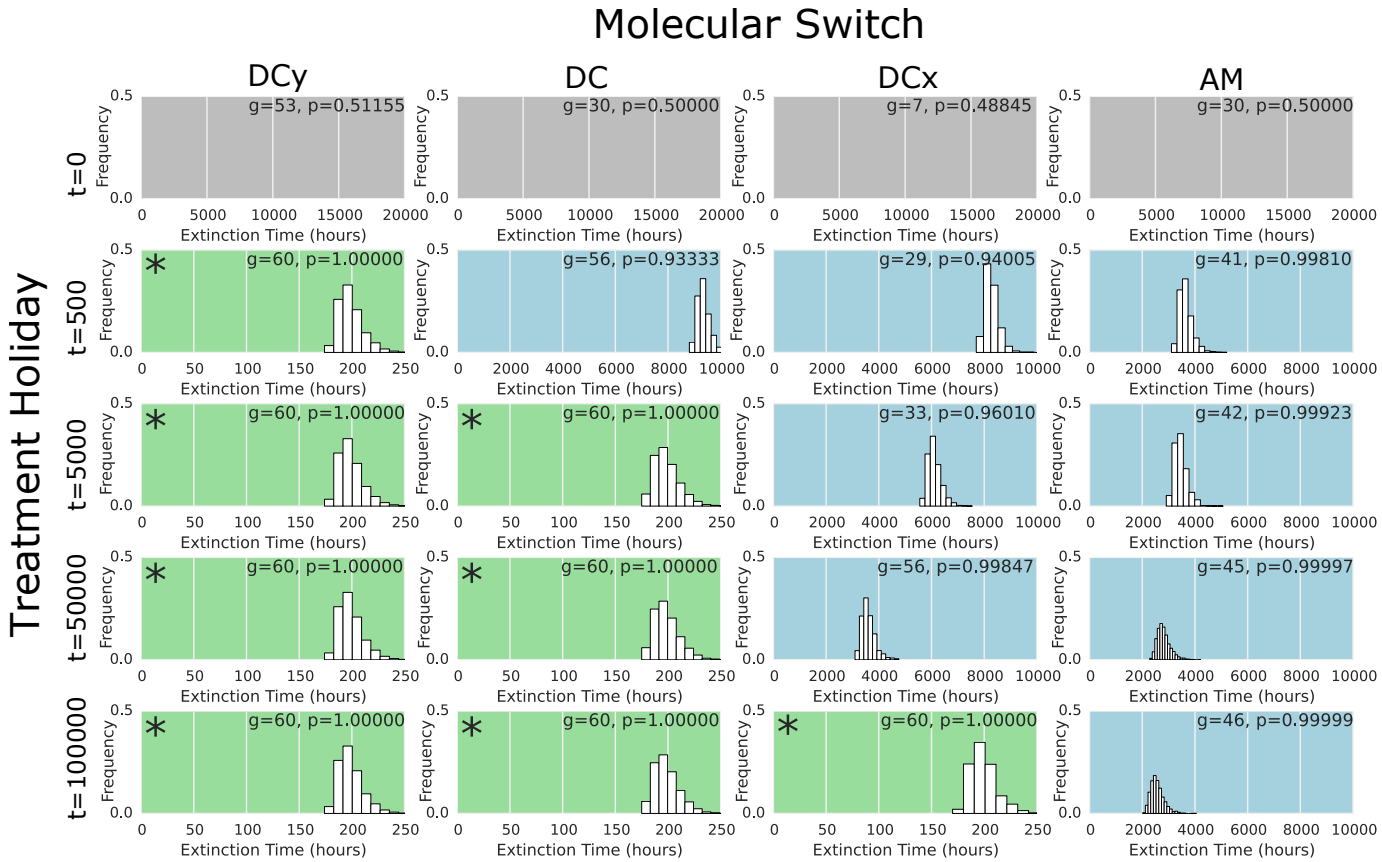

Fig. 4.14 **Extinction dynamics for populations endowed with different molecular switches over different timescales of treatment holidays.** Each histogram shows the distribution of extinction times over 2000 simulations of treatment in an individual–based model with drug parameters corresponding to a mixed cytostatic and cytotoxic drug. The molecular switch used as the GP–map is shown as the column heading. The genotype and associated probability of phenotype *A* (shown inset to each subfigure) are determined by an evolutionary simulation of a treatment holiday for a timescale determined by the row. A blue background indicates extinction times longer than a timeframe viable for an antibiotic treatment, a green background (or inset star) indicates extinction times within this timeframe.

To test the robustness of our results we considered an intermediate parameter set of $d_A = 0.4$ h^{-1} and $f_A = 0.5$ h^{-1} as a trade-off between cytostasis and cytotoxicity. The histograms of extinction times in this simulation are presented in Figure 4.14. We find that, although the specific extinction times change, the qualitative results remain the same. Specifically, those combinations of treatment holiday, T, and switching mechanism that resulted in an effective follow–up treatment with a cytotoxic drug (Green, Figure 4.13) also permit effective treatment under this mixed parameter set. The converse is also true, with the blue (extinction but not within 240hrs) and grey (no extinction with 20,000hrs) histograms from Figure 4.13 being preserved.

These simulations demonstrate the critical influence of the underlying molecular mechanisms that drive phenotypic bet–hedging. Specifically, we see that the length of treatment required to drive a population to extinction is dependent on the proportion of each phenotype in the population, which is in turn determined by the population genotype and the form of the molecular switch. The length of treatment holiday required to steer the evolving population into a more readily treatable configuration (i.e. a sufficiently small proportion of persister–like individuals to permit rapid extinction) is dependent on the underlying molecular mechanism or GP–mapping. To predict the efficacy of treatment holidays as a potential therapeutic intervention for a disease with bet–hedging–driven resistance it is imperative that we account for the influence of the underlying GP–mapping.

4.7 Implementation Details and Computational Considerations

The implementation of the models, derivation of the results and the production of all of the figures was performed using the Python programming language. A full annotated implementation of all models, the derivation of all results, and the code used to produce the figures is available from Github (see Appendix 3). Appendix 3 contains a Jupyter notebook that provides an overview of this code.

4.7.1 Performance

The switching probabilities for the DC, DCx and DCy switches are calculated almost instantaneously from the analytic equations. The derivation of the look up table for the AM switch with $g_{\max} = 60$ takes on the order of minutes in a personal computer but grows intractable for $g_{\max} > 200$. Simulations of the stochastic network dynamics are implemented by a Gillespie algorithm and terminate by reaching a stable configuration within seconds for all switches.

The calculation of invasion probabilities is performed by numerical solution of Equation 4.12 with the SciPy package for scientific computation. This solution is derived almost

instantaneously on a personal computer. 30 simulations of Algorithm 1 are performed to generate Figure 4.8, these simulations take approximately 10 minutes to conclude. Finally Algorithm 2, which is unoptimised, is run 2000 times to generate each of the extinction time histograms in Figures 4.13, 4.12 and 4.14. Each set of 2000 samples takes between 12 and 24 hours on a personal computer. The histogram data presented in this thesis were generated by running Algorithm 2 on a scientific computing cluster.

4.7.2 Testing

Testing of the derivation of switching probabilities was performed as follows. 1) Known values ($x_b = 0$, $x_b = 60$, and $x_b = 30$ for DC/AM) were compared to the values produced by the code. 2) The code-derived switching values were sanity checked, for example by testing monotonicity ($x_{b_1} > x_{b_2} \implies p\left(x_{b_1}\right) > p\left(x_{b_2}\right)$) and correct limits ($0 < p(x_b) < 1$).

The calculation of invasion probability by numerical solution of Equation 4.12 was performed by the SciPy scientific computing package. We did not attempt to validate the correctness of the solver implemented in this package. To test the values derived through this method we performed sanity checks as follows: 1) Verified that deleterious mutations have invasion probability 0 and 2) Verified that the invasion probability is monotonic with respect to the fitness benefit associated with a mutation. Further, the numerical derivation of invasion probability was compared between two solvers (by comparison to the solution derived with Mathematica)

Testing for the simulation of stochastic evolutionary loss of bet–hedging comprised of comparing the distribution of convergence times in the stochastic simulation to the analytically derived mean value. Testing of the stochastic simulation of population dynamics was performed by comparing multiple realisations of the stochastic simulation to the deterministic growth dynamics. All tests confirm the correctness of our approach.

4.8 Discussion

We have introduced a model for the genotype–phenotype mapping that uses stochastic simulation of chemical reaction networks to determine phenotypes. Where this model differs from previous theoretical models of the GP–map (Chapter 2), specifically the neural network model of Gerlee and Anderson [106] or the GRN model studied by Wagner [331] and later Huang [154, 155], is that ours is the first to explicitly model the stochastic aspect of phenotypic determination that results from intra–cellular thermal fluctuations. Using this model we have demonstrated how remarkably simple CRNs can implement complex switching behaviour and produce populations in which different proportions of cells, determined by the initial conditions of the network, take on different phenotypes. Critically, simple network motifs have been demonstrated to implement switching behaviour that can produce populations in which different proportions of cells, determined by the initial conditions of the network, take

on different phenotypes. Whilst abstract in its representation of reactions within a network, the reactions in our model are closely related to the physical mechanisms that govern intracellular regulatory networks, providing valuable insight into the impact of network architecture on the stochastic determination of phenotypes.

4.8.1 Implications for the Drivers of Bet–Hedging

We have demonstrated how introducing redundancy, a common feature of many biological systems, into the CRNs governing the GP–map can create molecular switches that are robust to the addition or removal of chemical species. Assuming that the chemical species of our CRN correspond to expressed gene products, this finding shows how redundancy, which can arise through neutral or nearly–neutral mutation events, can ensure that the phenomenon of bet–hedging is not lost through gene deletions or duplications. Critically, this observation may explain the failure to identify genes responsible for bacterial persistence. For example, in the review by Lewis [204] mutations to the genes hipA, rmf, sulA, and toxin–antitoxin (TA) loci relBE, dinJ and mazEF were highlighted as possible drivers of bacterial persistence. However, deletion of rmf, relBE or mazEF has been demonstrated to have no effect on the phenomenon of persistence, owing possibly to redundancy in TA modules, whilst deletion or over expression of hipA can change the proportion of bacterial persisters but not eradicate them. This is consistent with our predictions, since (when redundancy is present) deleting any single species in the CRN will not prevent bet–hedging or non–genetic heterogeneity, instead it will only alter the switching probability, or at the population scale, the proportions of each phenotype. In the context of bacterial persistence, the conclusion taken from the experiments reviewed in Lewis [204] need not be that none of the potential genetic factors identified are the ones driving bacterial persisters, instead it may be that the search for a single genetic factor responsible for bet–hedging is doomed to fail. Indeed, it may be that bet–hedging emerges from the interactions of a collection of genetic factors in the sense of the epigenetic landscape introduced by Brock et al. [38] and Huang [154]. If this is the case, then to identify the biological mechanisms responsible for hedging we will need to move beyond the gene–centric perspective and begin to identify those interactions responsible for stochastic determination of phenotypes.

Mutations in cancer have often been associated with their direct effect on phenotypes; the concept of a driver mutation being that it induces a novel adaptive phenotype leading to clonal expansion [222]. However, our results suggest another phenomenon, wherein mutations do not induce novel phenotypes, but rather alter frequencies of pre–existing phenotypes within the population. This change in phenotypic ratio can have implications for cancer progression, a phenomenon previously explored by Charlebois et al. [46]. Consider the phenomenon of tumorigenic cells, where it is thought that only cells of a certain phenotype can form a growing tumor mass [251, 268, 223]. Genetic heterogeneity can explain the existence of a tumorigenic subpopulation if certain driver mutations are responsible for the tumorigenic phenotype. However, an alternative mechanism is that stem-like tumour cells give rise to

a population of heterogeneous phenotypes. In the traditional stem-cell model, a hierarchy exists where the stem–cells produce a range of tumour cell phenotypes [266]. Cancer stem cells divide to produce either more cancer stem cells (self–renewal) or cells with non–stem phenotypes down the hierarchy. This cellular decision is often taken to be stochastic (an example of bet–hedging) and our results highlight the potential for mutations to alter the probabilities of self–renewal or differentiation which have been shown to have significant impact on many aspects of tumor progression [82].

An alternative bet–hedging mechanism for tumorigenicitiy is that the tumorigenic phenotype is transient and stochastically determined (potentially with microenvironmental influence). Evidence for this phenomenon was highlighted by Quintana et al. [263] who demonstrated the tumorigenic potential of individual melanoma cells is similar, despite high heterogeneity of many markers in the initializing cell. No driver or stem population was found, and indeed, the heterogeneity of marker expression was recapitulated by most tumorigenic cells, regardless of the starting pattern of expression. Such a mechanism would have different implications from the stem cell model presented above, as potentially sensitive cells would be more difficult to define. However, our predictions remain the same: that genetic mutations could shift the frequency of tumorigenic phenotypes and profoundly impact cancer progression.

Our results further demonstrate that the structure of the network governing phenotypic differentiation also has important implications for the evolutionary loss of bet–hedging. By considering mutations to the expression levels of genes, we find that the time taken for a two–phenotype bet–hedge to be lost, in an environment favourable to only one of the phenotypes, can vary by orders of magnitude depending on the network structure. If bet–hedging serves as a survival mechanism in the event of rare catastrophic environmental change (e.g. drug treatment) then the GP–mapping can prevent loss of this survival mechanism over the long time scales in which catastrophe does not occur. For example, if bet–hedging is driven by a network such as the Approximate Majority (AM) network (Figure 4.2), then each successive mutation towards a one–phenotype strategy induces a diminishing increase in the probability of generating that phenotype and, thus, a diminishing increase in expected fitness. Eventually mutations become (essentially) neutral and unable to fix in a large population. It is the structure of the molecular switch that substantially slows evolutionary convergence. This result provides a possible solution to the apparent paradox of bet–hedging – how can bet–hedging persist in environments where natural selection acts to remove it? The structure of the molecular mechanism, itself subject to natural selection, can slow the loss of bet–hedging strategies to ensure a survival mechanism, even where environmental catastrophes are very rare.

4.8.2 Implications for Therapy

These results have important implications for therapeutic strategies to treat diseases displaying bet–hedging–driven drug resistance. Theoretical strategies to combat bet–hedging–induced drug resistance have focused on identifying novel agents capable of killing persister cells, or identifying genetic mechanisms that can be targeted to prevent the persister phenotype from

emerging. This latter strategy bears a striking resemblance to the targeted therapy revolution in the treatment of many cancers. The identification of molecular targets whose inhibition induces death (or inhibits growth) has led to the discovery of a number of novel therapies for melanoma, non–small cell lung cancer and colorectal cancers. These drugs are, in the short term, remarkably effective, however the effects are rarely durable. Mutations that abrogate the effects of targeted therapies quickly emerge during treatment, driving resistance and ultimately mortality. The results of our chemical reaction model shed light on this Darwinian adaptation and suggest that targeted therapies to prevent bet–hedging may either be impossible, or where they do exist, prone to fail due to the re–emergence of bet–hedging through evolution. More precisely, the discovery of a single "magic bullet" genetic factor [314] which when targeted can switch off multi–drug resistant dormant phenotypes is unlikely, owing to redundancy in the network architecture. However, we should not rule out the potential of targeted therapies entirely. It may be possible to identify multiple targets for which simultaneous inhibition prevents bet–hedging. Alternatively, targets may be identified that shift the proportion of resistant or dormant individuals within a population to a manageable level, either permitting treatment with other cytotoxic agents or driving the disease into a dormant state.

A second theoretical treatment strategy suggested to combat resistance in cancers and resistant infections [32] is the introduction of treatment holidays. The traditional doctrine for therapy is one of *maximal dose–density* – that we should treat diseases using the most potent drug with the highest tolerable dose for the longest possible time until the disease is cured, the therapy ceases to be effective, or the drugs become too toxic. Mathematical models of disease progression assuming genetically–driven resistance indicate that this approach could drive the emergence of resistance through an ecological principle called *competitive release* [9, 2]. Before treatment, cells compete with one another for limited resources within a spatially constrained population. In a non–treated environment, pre–existing resistant cells are often less fit than sensitive ones and thus, owing to competition, do not grow to large numbers within the population. Selective pressures, for example within a growing tumour, are often not sufficiently strong for clonal sweeps (the fixation of a single genetic clone) to occur [309, 271] and the population contains a heterogeneous mix of phenotypes. When this population is exposed to a maximum dose–density therapy the sensitive cells are killed, allowing the rapid outgrowth of the previously small resistant population. This population then drives the recurrence of drug–resistant disease. Treatment holidays have been suggested as a potential therapeutic strategy to avoid drug resistance driven by competitive release [84].

Here, we implemented an individual–based model of the dynamics of a bet–hedging population under treatment to explore the efficacy of treatment holidays in combating bet–hedging–driven resistance. Coupled with a long–term evolutionary simulation, we explored the impact of the mechanism driving bet–hedging on the efficacy of treatment holidays. Our model suggests that it is the GP–map, and in particular how it hinders or promotes the rate of evolution, that determines the efficacy of treatment breaks in reversing drug resistance. If the mechanism driving bet–hedging permits evolutionary loss in a short to medium time scale,

then treatment holidays may drive the loss of bet–hedging and re–establish drug sensitivity. However, if instead the driving mechanism slows the evolutionary loss of bet–hedging, then a treatment break is unlikely to re–establish drug sensitivity in a time frame relevant to disease progression. We note that interfering with the mechanism driving bet–hedging through targeted therapy could alter the switching to allow the fast evolutionary loss of bet–hedging where previously this would not occur.

4.9 Future Directions

We have introduced a model of the GP–mapping in which phenotypes arise stochastically from a heritable gene–expression profile. In Chapter 2, we highlighted that bet–hedging and other mechanisms of non–genetic heterogeneity are under–represented in the study of GP–mappings. We believe that our work serves as a good first step to remedy this of lack modelling. Critically, we have found that the evolutionary loss of bet–hedging is strongly determined by the structure of the GP–mapping, independently of the genotype. Further, we demonstrated how the efficacy of adaptive therapies in the treatment of diseases displaying bet–hedging–driven resistance is strongly dependent on the GP–mapping. Taken together, we believe that these results highlight the power of abstract modelling of the GP–mapping.

The model presented here is predicated on a number of assumptions. For example, we have taken the initial conditions of our reaction networks to be genetically determined and stable, allowing us to explore the implications of the structure of the network itself on the evolution of bet–hedging. The model we present could also be used to study additional aspects of non–genetic heterogeneity, other than those presented here, by introducing extensions. Firstly, the model could be extended to directly account for the dynamics of gene expression (for example using the model presented by Kaern et al. [169] or the GRN model presented by Huang [154]). Secondly, we could take the initial conditions as environmentally determined and the network itself to be genetically determined. In this case the model would give a GP–map similar to the neural network model used by Gerlee and Anderson [106] and Gerlee et al. [110] to study phenotypic plasticity. However, such a model would differ in that the determination of phenotypes would be stochastic, permitting the study of environment–dependent bet–hedging strategies. Thirdly, we have assumed that the chemical species comprising the reaction networks are entirely depleted before reproduction and determination of the next phenotype. Omitting this assumption to permit the epigenetic inheritance of molecular switch state represents an extension of our model to account for phenotypic memory or carry–over. It is this extension of the model that we explore in Chapter 5. Model extensions to account for phenotypic plasticity or for the dynamics of the gene regulatory network are not pursued as they are beyond the scope of this thesis, however phenotypic plasticity is briefly explored in Chapter 6.

Chapter 5

Mechanisms of Phenotypic Memory

5.1 Bet–Hedging, Epigenetic Inheritance and Phenotypic Memory

In the previous chapter we introduced a simple mechanistic model for a GP–mapping in which bet–hedging emerges naturally from the stochastic interaction of intracellular molecules. Under this GP–mapping, each possible genotype, g, was associated with a probability $p(g)$ of an individual having phenotype A (as opposed to an alternative phenotype B) through the simulation of a bistable switch. This probability was fixed and independent of the phenotype of the parent or the environment in which the individual resides. Thus, there was no epigenetic inheritance, no phenotypic memory (Lamarckian inheritance) and no phenotypic plasticity. However, these assumptions are not taken in all models of bet–hedging and weakening them results in evolutionary systems endowed with different kinds of bet–hedging. The assumptions taken in Chapter 4 can be considered a *strong–genetic* form of bet–hedging, in which it is solely the genotype (and GP–mapping) that govern the stochastic determination of phenotypes.

Many mathematical models of bet–hedging do not take this strong–genetic form. Instead, the likelihood of an individual taking a specific phenotype is taken as dependent on the phenotype of the parent. The continuous–time model for a bet–hedging population presented by Müller et al. [232] takes the form

$$\frac{dA}{dt} = A(t)(b_A - d_A - p_1) - K(A(t) + B(t)) + p_2 B(t) \tag{5.1}$$

$$\frac{dB}{dt} = B(t)(b_B - d_B - p_2) - K(A(t) + B(t)) + p_1 A(t). \tag{5.2}$$

Here, b_i and d_i denote the birth and death rates of individuals of phenotype $i \in \{A, B\}$, K is a function governing how competition impacts growth rate, p_1 denotes the likelihood of an individual of phenotype A producing offspring of phenotype B and p_2 denotes the likelihood of an individual of phenotype B producing an offspring of phenotype A. As p_1 and p_2 are permitted to vary independently, the stochastic process governing phenotypic differentiation can be dependent on the parent phenotype. This is a form of *phenotypic memory* in the

bet–hedging strategy and results in a larger space of bet–hedging strategies (than the strong genetic model) as p_1 and p_2 can vary independently. The strong–genetic form of bet–hedging is captured by this more general model when $p_1 = p(g)$ and $p_2 = 1 - p(g)$.

In this chapter, we consider two alternative *modes* of bet–hedging in addition to the strong–genetic form. In the first, the switching values p_1 and p_2 are taken to be determined genetically according to $p_1 = p_2 = p(g)$. Here, the genotype determines the likelihood that an offspring inherits the phenotype of its *parent*, instead of determining the likelihood of producing each phenotype independent of the parent phenotype. Thus, this model exhibits phenotypic memory and is a restricted version of the one presented by Müller et al. [232]. We refer to this model as *(genetically) constrained memory* and refer to the model of Müller et al. [232] as *unconstrained memory*. As the unconstrained memory model has been studied extensively, we choose not to include it in the following analysis, opting instead to focus on restricted models that are less well–studied.

For the constrained (and indeed unconstrained) memory models, no assumptions are made as to the mechanism that drives epigenetic memory[1]. The molecular interactions that result in offspring taking (or not taking) the phenotype of their parent are not modelled. Instead, the resulting stable state of a bistable molecular switch determines $p(g)$, which is then interpreted differently (without mechanistic justification) for the two models.

For the second mode of bet–hedging studied here, we introduce a mechanistic model of epigenetic inheritance for the intracellular molecules comprising the bistable switch. These inherited molecules act to bias the switch, serving as a mechanistic form of phenotypic memory. As outlined by our exploration and review of GP–mappings in Chapter 2, organismal phenotype is determined by the environment, the GP–mapping, and whatever is *inherited* by the offspring. In Chapter 2, we designated all these inherited factors as the genotype, but noted that in biological reality much more than solely genetic factors are inherited. For example, intracellular factors, epigenetic modifications and the cell–wall are all inherited. From this modelling philosophy it follows that phenotypic memory in bet–hedging must derive from the inheritance of some phenotype–dependent substrate.

The explicit distinction between different modes of bet–hedging, as determined by differing contributions of genetic and epigenetic inheritance, is not commonly made in theoretical modelling. Indeed, the majority of models are founded on the assumption of unconstrained memory [232, 164]. This focus on unconstrained memory may stem from mathematical convenience, or alternatively, it may be that the need for such an explicit distinction has not previously been noted. By constructing a general framework of evolution, and elevating the GP–mapping to equal standing with (epi–)genotype, phenotype or environment, the distinction between the modes of bet–hedging becomes much clearer. Indeed, the explicit designation of *what* is inherited, coupled with a precise definition of the GP–mapping, suggests the possibility of different modes of bet–hedging.

[1]Note that throughout this chapter, we refer to memory as *epigenetic* in the sense that it is heritable but not genetic.

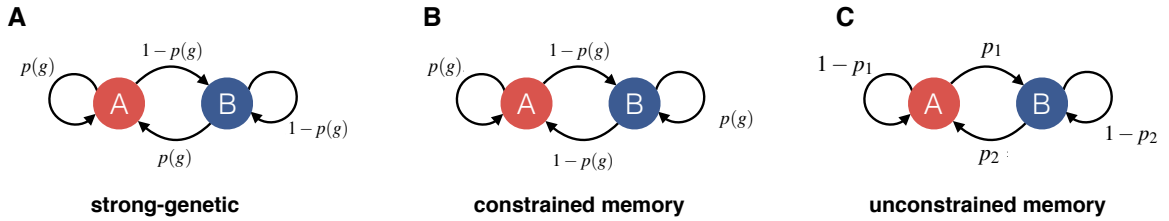

Fig. 5.1 **The different modes of bet–hedging.** (Left) Strong–genetic bet–hedging wherein all offspring take phenotype A with probability $p(g)$, regardless of the parental phenotype. (Centre) Constrained memory bet–hedging wherein offspring take the same phenotype as their parent with likelihood $p(g)$. (Right) Unconstrained memory bet–hedging wherein offspring take a phenotype matching their parent with a probability that is dependent on the phenotype of the parent.

In this chapter, we extend our previous theoretical modelling to account for different mechanisms of bet–hedging. We begin by providing a brief overview of the results of [232], wherein (unconstrained memory) bet–hedging is proven to be an optimal strategy in fluctuating environments. We then introduce two extensions of the molecular switch model described in Chapter 4 that account for phenotypic memory. The first of these extensions is the constrained model described above. In the second, explicit modelling of the production, decay and inheritance of intracellular factors is used to implement a mechanistic model of phenotypic memory, which we call the *epigenetic memory* model. We compare the dynamics of the strong–genetic, constrained memory and epigenetic memory models, in both fixed and fluctuating environments, to determine how they impact on fitness and evolution. Finally, we explore the relationship between these modes of bet–hedging and the efficacy of metronomic therapy, deriving optimal therapeutic strategies, and exploring the potential for adaptive therapies in controlling bet–hedging diseases.

Bet–Hedging with Unconstrained Memory is Optimal in Fluctuating Environments

Müller et al. [232] demonstrated that under certain environmental conditions, unconstrained memory bet–hedging between two phenotypes offers an adaptive benefit over any single phenotype strategy. The mathematical basis of this work, and of other models of bet–hedging [164], is the theory of adaptive dynamics. Consider Equations 5.1 represented in matrix form

$$\begin{pmatrix} \frac{dA}{dt} \\ \frac{dA}{dt} \end{pmatrix} = \mathbf{M}(t) \begin{pmatrix} A(t) \\ B(t) \end{pmatrix} \tag{5.3}$$

where $\mathbf{M}(t)$ encapsulates the (environmentally dependent) per capita rates of birth and death, the effects of competition (K) and the (fixed) rates of transition, p_1 and p_2, at each time t. The expected initial growth rate of an infinitesimal bet-hedging population governed by these

dynamics (also called the *invasion fitness*) is given by the *dominant Lyapunov exponent*

$$\rho = \lim_{t\to\infty} \frac{1}{t} \ln(N(t)) \tag{5.4}$$

that is determined by the matrix sequence $\mathbf{M}(t)$, where $N(t) = A(t) + B(t)$. The existence of this limit is not guaranteed; however, under certain technical conditions (that are usually satisfied by biological models) the limit exists. For example, provided the environmental fluctuations governing the sequence $\mathbf{M}(t)$ are ergodic, that the entries in each $\mathbf{M}(t)$ are uniformly bounded above, and provided that for all initial conditions, $A(0), B(0)$, there exists T such that for $t \geq T$, $A(t), B(t) > 0$, ρ is guaranteed to exist [226].

The study of evolutionary dynamics using Lyapunov exponents is equally valid for continuous models, such as the one defined by Equations 5.1, and discrete time models, as considered in Chapter 4. However for the cases of continuous and discrete models, the computation of ρ will differ [226]. Further, where the dynamics defined by the $\mathbf{M}(t)$ are non–linear, or stochastic, the explicit computation of ρ may not be possible. Instead, the calculation of ρ must be performed using numerical approximations or the derivation of appropriate inequalities.

Where the matrices $\mathbf{M}(t)$ are determined by an environment induced by a resident population, ρ determines whether a new, initially infinitesimal, population can grow in size. Specifically, where ρ is positive, *invasion* – defined as the initial outgrowth of the initially small population – can occur. This criterion does not actually guarantee the invasion of a mutant, only that the initial growth rate is (expected to be) positive. Indeed, the mutant population may experience stochastic extinction through drift. Alternatively, when the invading population becomes large, it may alter the environment, changing selective pressures (the matrices $\mathbf{M}(t)$) and preventing invasion. By estimating the invasion fitnesses of populations exhibiting either bet–hedging or monomorphic strategies, and under different regimes of stochastically fluctuating environments, Müller et al. [232] proved (a stronger form of) the following result.

Theorem (Muller 2013): Consider a stochastically fluctuating environment governed by a Markov process such that the times spent in each environment, e_1 and e_2, are exponentially distributed with means τq_1 and τq_2 (resp.) with $q_1 + q_2 = 1$ and $\tau \in \mathbb{R}$. Suppose that e_1 favours phenotype A and e_2 favours phenotype B. Then there exists a bet–hedging strategy with

$$p_1 = \frac{1}{q_1\tau} + \mathscr{O}(\tau^2) \tag{5.5}$$

$$p_2 = \frac{1}{q_2\tau} + \mathscr{O}(\tau^2) \tag{5.6}$$

that is an Evolutionarily Stable Strategy.

This result requires that the switching probabilities p_1 and p_2 are permitted to vary independently and to be finely tuned. Thus, it is dependent on the assumption of bet–hedging with unconstrained memory. Whether mechanistically restricted modes of bet–hedging, for example the strong–genetic or constrained memory models described earlier, offer a comparable evolutionary advantage remains unanswered. It is for this reason that we study these restricted forms in this chapter.

The result of Müller et al. [232] holds for both stochastic and cyclic environments. However, in the case of stochastic environments, the computation of ρ requires sophisticated mathematics. Conversely, where the environment is cyclic, the value of ρ can be computed easily from the dominant eigenvalues of the matrices $\mathbf{M}(t)$. A numerical scheme for this derivation was provided by Metz [226]. As we are interested in the evolution of drug resistance to pre–determined treatment strategies, we choose to focus only on cyclic environmental regimes, simplifying the calculation of population growth rates.

5.2 Extending the CRN Model to Account for Phenotypic Memory

5.2.1 Intra–Cellular Dynamics

Epigenetic Inheritance

To mechanistically model phenotypic memory we extend the CRN model introduced in Chapter 4 to account for *epigenetic inheritance* of intra–cellular factors, specifically the molecules x and y, whose interactions encode a bistable switch. In our previous (strongly–genetic) model, the initial abundances of the molecules x and y were taken as genetically determined and instantaneously produced at the birth of an individual. Further, no residual x or y molecules are present at the time of reproduction; an assumption justified by the observation that the bistable switch is likely a motif in a much larger interaction network. We now weaken this assumption and extend our model such that the initial conditions for x and y at the birth of an individual are determined from a combination of a *genetic burst* component, $g = (x_b, y_b)$ and an epigenetically inherited component (x_I, y_I). This inherited component will be dependent on the phenotype and age of the parent, and will serve to bias the molecular switch towards producing offspring that match the parent phenotype.

We extend the bistable switch model by endowing each chemical reaction network with additional first order decay reactions for the molecules x and y,

$$\begin{aligned} x &\xrightarrow{\sigma_1} \emptyset \\ y &\xrightarrow{\sigma_2} \emptyset. \end{aligned} \tag{5.7}$$

The decay of these molecules will progress throughout the life of an individual. For example, intracellular factors may be degraded or combined with others in the production or maintenance

of the cellular phenotype. As a simplifying assumption, we maintain that the resolution of the CRN to a stable configuration is instantaneous, and that no decay reactions occur during this process. This assumption is equivalent to supposing a separation of time scales, specifically that $\sigma_i \ll \rho_j$ for all $i \in \{1,2\}$, $j \in \{1,2,3\}$ (for the ρ_j values defined in Equation 4.1). This assumption simplifies our modelling, as the method for computing switching probabilities introduced in Chapter 4 remains valid. For the Approximate Majority switch, all molecules of type b are consumed during the switching process and thus we need not consider their decay dynamics.

Under this extended model, the GP–mapping, intra–cellular dynamics and population dynamics will be closely linked. Feedback from the population scale, specifically the age of an individual at reproduction, determines the intracellular state at the time of reproduction. This intracellular state then directly influences the likelihood of offspring taking each of the phenotypes. Consider an individual with intra–cellular molecular state (x_r, y_r) at time of reproduction. As we assume the chemical species x and y exist as a part of a bistable switch, either $x_r = 0$ or $y_r = 0$. We consider reproduction to be symmetric – producing two new offspring and removing the parent[2]. We assume that each offspring inherit precisely half (subject to ± 1 where the number is odd) of the parent's internal molecules x_r and y_r[3]. Each offspring adds to the number of molecules with a genetic burst to yield an initial internal molecular state of

$$(x_0, y_0) = \left(\frac{x_r}{2}, \frac{y_r}{2}\right) + (x_b, y_b). \tag{5.8}$$

The phenotype of the offspring is then determined, as in Chapter 4, by stochastically simulating a bistable CRN, or equivalently, sampling from a look–up table of switching probabilities for the CRN. At birth, the offspring are assigned an internal configuration dependent on the switch state (i.e. phenotype) according to

$$(x,y) = \begin{cases} (\frac{x_r}{2} + \frac{y_r}{2} + x_b + y_b, 0) & \text{if the bistable switch proceeded to the all } x \text{ config} \\ (0, \frac{x_r}{2} + \frac{y_r}{2} + x_b + y_b) & \text{if the bistable switch proceeded to the all } y \text{ config} \end{cases} \tag{5.9}$$

This internal state will then be subject to the decay of x and y molecules governed by the reactions in Equation 5.7 throughout the life of the individual to determine the values (x_r, y_r) when at reproduction.

[2]This is a simplifying assumption, particularly in a model of evolution as general as we consider here, as in reality the two offspring arising from mitosis can differ. For example, senescence has been observed in cancer cells [276, 65], in asymmetrically dividing bacteria [243], and potentially in bacteria reproducing through morphologically symmetric division [313] although recent studies refute this phenomenon [338].

[3]Again, this is a simplification. A more realistic assumption would be to determine the split between each daughter from a probability distribution. We forgo this approach to maintain a tractable model and to avoid clouding our observations of the relative contributions of stochasticity introduced at different biological scales.

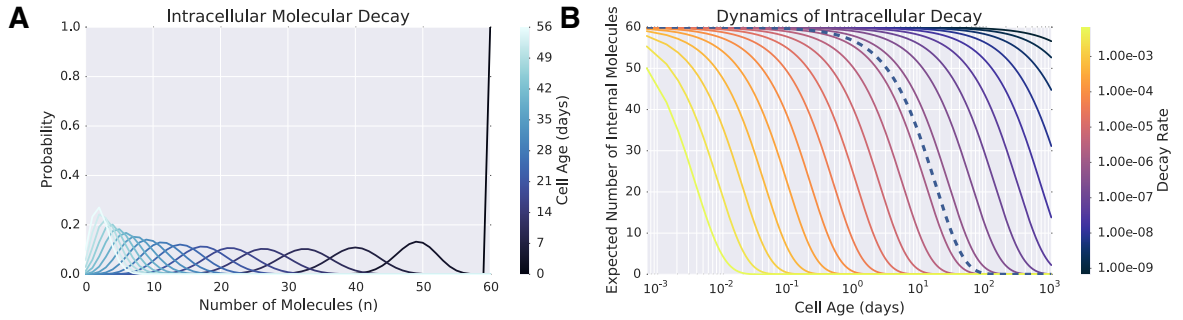

Fig. 5.2 **Internal molecular state varies through time. A)** For an individual with initial internal molecular state $(x_b, y_b) = (60, 0)$, probability distributions $x_b(t)$ are shown varying individual age t (indicated by curve colour). The molecular decay rate is $\sigma = 6.77 \times 10^{-7}$ s^{-1}. **B)** Curves show the expected intracellular state, $\mathbb{E}\left(x_b(t)\right)$, as an individual ages for different σ values (curve colour). The curve highlighted in dashed blue corresponds to $\sigma = 6.77 \times 10^{-7}$ s^{-1} used in (A).

Molecular Decay Dynamics

To complete the model the values of (x_r, y_r) must be calculated. Suppose an individual contains precisely n_0 molecules of x or y at birth and following the determination of the phenotype. These molecules are modelled as subject to decay by the reactions defined in Equation 5.7. Denote by $p(n,t)$ the probability that there are n molecules of the chemical species after t seconds, then

$$p(n,t) = \binom{n_0}{n} e^{-\sigma_1 nt}(1 - e^{-\sigma t})^{n_0 - n}. \tag{5.10}$$

A derivation of this identity by use of a chemical master equation is provided by Erban et al. [86]. This identity provides a method to sample a value for n at a given time t. Figure 5.2(A) shows the evolution of the likelihood function $p(n,t)$ over time, t, with parameters $n_0 = 60$, $\sigma = 6.77 \times 10^{-7}$ s^{-1}. The expected value for the number of molecules n at t seconds is given by

$$\mathbb{E}(n|n_0, t, \sigma) = n_0 e^{-\sigma t}. \tag{5.11}$$

A plot of this expected value over time is shown in Figure 5.2(B) for varying decay rates σ. Here, the curve highlighted in blue corresponds to the σ value used to generate Figure 5.2(A).

Determining an empirical value for decay rate σ from experimental evidence is not possible as empirical factors driving hedging or memory have not yet been clearly identified [204]. We can however estimate σ by considering the qualitative behaviour of memory that is expected. If we define a value T_i^* for the average age at which we expect an individual of phenotype i to produce offspring with an entirely depleted memory component, then the value of σ can be estimated according to

$$\mathbb{E}\left[\text{number of molecules}|n_0, T_i^*, \sigma\right] = n_0 e^{-\sigma T_i^*} < 1. \tag{5.12}$$

Rearranging this equation yields

$$\sigma > -\frac{1}{T_i^*}\log\left(\frac{1}{n_0}\right). \tag{5.13}$$

Throughout this Chapter, we suppose $n_0 = 60$ as in Chapter 4. For a strong memory component, we assume that complete depletion of memory occurs in 10 weeks. Thus, we require

$$\sigma > \frac{-1}{70\times 24\times 60\times 60}\log\left(\frac{1}{60}\right) \approx 6.77\times 10^{-7}\ \mathrm{s}^{-1}. \tag{5.14}$$

We take $\sigma = 6.77\times 10^{-7}\ \mathrm{s}^{-1}$. This time to depletion might be considered very long. However, as the decay occurs exponentially, the number of intracellular molecules will become small (and thus memory weak) long before the 10 week period. In our study of drug resistance, we explore how the value of σ is expected to evolve in response to drug therapy and find that selection can act to both minimise or maximise σ under certain conditions. This observation lends credibility to the assumption of a small σ parameter value, as selection can act to minimise it. Further, it should be noted that the relationship between the genetic component of memory, (x_b, y_b), the average age to loss of memory T_i^*, and the decay rate σ are independent of the underlying molecular network driving phenotypic differentiation.

We note that the number of molecules required to induce a significant phenotypic bias, and thus the impact of changes in the rate σ, will be strongly dependent on the molecular switch that drives bet–hedging. Figure 5.3 shows the effects of the decay dynamics of Figure 5.2 on the likelihood of phenotype A emerging under each of the DC, DCx, DCy and AM switches.

The relationship between cellular age, internal molecular state and phenotypic switching is not a simple one and indicates the potential for epigenetic inheritance to increase fitness. Consider an individual with a high–fitness phenotype that reproduces quickly. Offspring of such an individual will inherit a greater number of additional molecules at birth, and these additional molecules will *bias* the probability of the offspring taking the same phenotype as the parent. This effect will be stronger where the parent is fitter and reproduces more quickly. Thus, epigenetic inheritance may serve as a mechanism for phenotypic memory that increases the likelihood of individuals, adapted to the present environment, producing offspring that are also adapted.

This memory effect would be *environmentally determined* and reversible. When the environment changes, the reproductive rate may slow and the number of molecules inherited epigenetically is lowered. Offspring will then become less likely to take the same phenotype as their parents and the memory will be weakened. Thus, the memory phenomenon will be transient, allowing a population to be better adapted in a number of environments without the need for genetic alteration. This phenomenon represents a straightforward mechanism for *plasticity of bet–hedging*, as environment changes will induce change in the proportion of phenotypes without any alteration to genotype or the GP–mapping itself.

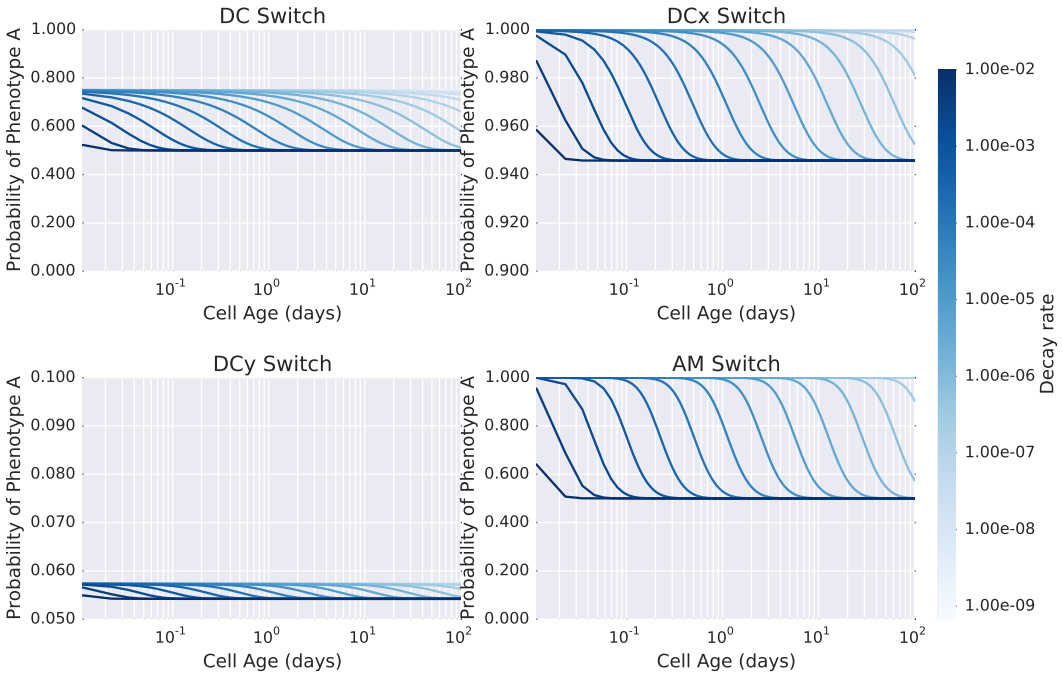

Fig. 5.3 **The decay of molecular memory for phenotype A.** For an individual with phenotype A and initial internal molecular state $(x_I, y_I) = (60, 0)$, the probability of an offspring taking phenotype A is plotted over time (with $(x_b, y_b) = (30, 30)$. A plot is shown for each molecular switching mechanism studied in Chapter 4. Line colours indicate the decay rate σ. Note that the curves in each case tend to a fixed probability regardless of the decay rate, this is the memory–free (strong genetic) bet–hedging studied in Chapter 4. This switching probability is achieved when all of the internal molecules have decayed.

Memory and Multiple Generation Carry–Over

Where reproduction is sufficiently fast, or molecular decay sufficiently slow, offspring will epigenetically inherit more molecules at birth than the parent did. This can occur over multiple generations and result in a lineage in which each individual exhibits a greater degree of memory than its parent. Here, we outline the dynamics of this phenomenon and demonstrate that it is self–limiting.

Consider a bet–hedging population endowed with the memory dynamics described above wherein individuals with phenotype A reproduce on average after T_A seconds. We compute the expected number of molecules such an individual will inherit when descended from a lineage in which every individual has phenotype A (Note that the following is also applicable to a lineage of all B phenotypes). Suppose the founder of a lineage of A-type individuals inherited no molecules epigenetically. This individual will have initial intracellular state $x_0 = x_b + y_b$, $y_0 = 0$. Denote by M_κ the expected number of molecules inherited by the κ^{th} generation of the lineage. Then $M_0 = x_b + y_b$, and we can compute the M_κ by the recurrence relation

$$M_\kappa = \frac{1}{2} M_{\kappa-1} e^{-\sigma T_A} + M_0. \tag{5.15}$$

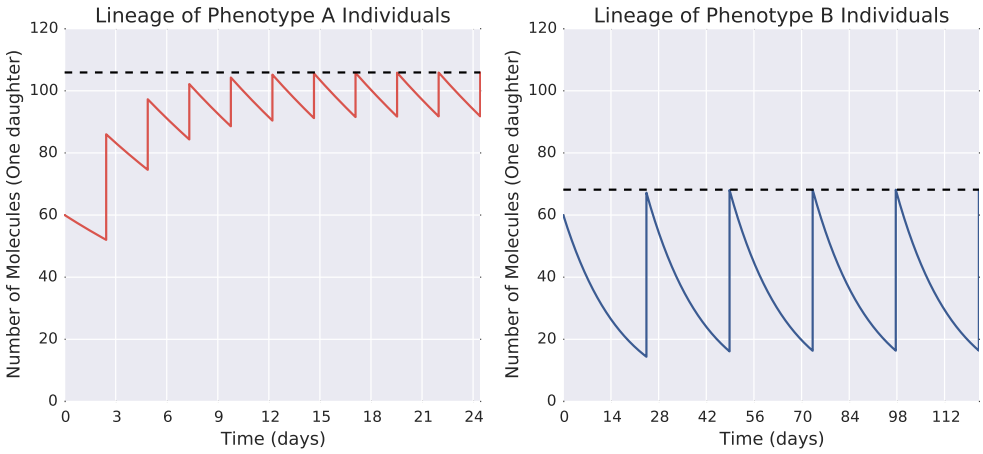

Fig. 5.4 **Dynamics of molecular decay and inheritence.** The expected number of internal molecules for a lineage of phenotype A (left, red) and phenotype B (right, blue) is shown. The intracellular parameters are $x_b = y_b = 30$ and $\sigma = 6.77 \times 10^{-7}$ s^{-1}. Births are simulated every $T_A = 2.44$days for phenotype A and $T_B = 24.4$days for the phenotype B individuals. The dashed lines show the value M^* (Equation 5.18).

Unwrapping this equation yields a geometric series

$$M_\kappa = \frac{1}{2^\kappa} M_0 e^{-\kappa\sigma T_A} + \frac{1}{2^{\kappa-1}} M_0 e^{-(\kappa-1)\sigma T_A} + \cdots + \frac{1}{2} M_0 e^{-\sigma T_A} + M_0 \tag{5.16}$$

$$= M_0 \frac{1 - (\frac{e^{-\sigma T_A}}{2})^{\kappa+1}}{1 - (\frac{e^{-\sigma T_A}}{2})}. \tag{5.17}$$

Taking the limit of the lineage length to infinity, we see that the expected number of inherited molecules approaches a limit

$$M^* = \lim_{\kappa\to\infty} M_0 \frac{1 - (\frac{e^{-\sigma T_A}}{2})^{\kappa+1}}{1 - (\frac{e^{-\sigma T_A}}{2})} \tag{5.18}$$

$$= M_0 \frac{1}{1 - (\frac{e^{-\sigma T_A}}{2})} \tag{5.19}$$

as $T_A, \sigma > 0$ implies $e^{-\sigma T_A} < 1$. Figure 5.4 shows a trace of the expected total number of molecules within the offspring cell over multiple generations for both a proliferative (A) and persister–like (B) phenotype. Here, the dotted lines correspond to the limit derived in Equation 5.18. Note that M^* is not an upper bound on the number of molecules an individual can inherit. As reproduction time is assumed to be exponentially distributed in our model (i.e. $T \sim \text{Exponential}(\frac{1}{T_A})$), reproduction can occur in a time faster than T_A. An upper bound can be computed by assuming instantaneous reproduction at each step in the lineage, $T_A = 0$. In this case, Equation 5.18 yields

$$M_{\max} = 2M_0 = 2(x_b + y_b) \tag{5.20}$$

demonstrating that the epigenetic memory will exhibit a saturating or self–limiting behaviour due to the dilution of intracellular molecules that occurs at reproduction.

5.2.2 Simulating Population Dynamics

We now consider both an explicit stochastic model and a deterministic model of population dynamics with epigenetic memory as described above.

Stochastic Simulation

An explicit simulation of the population dynamics with epigenetic memory can be implemented by a Gillespie algorithm. Algorithm 3 shows a pseudocode implementation of these explicit population dynamics. Here, the number of intracellular molecules within a cell is updated according to Equation 5.10 whenever an individual is chosen for reproduction. This explicit implementation is computationally expensive for an exponentially growing population. Indeed, Figure 5.5 (blue line) shows that the computational time taken to simulate each hour of population growth increases exponentially for this model. Throughout this chapter, we will be aiming to compare the growth rates for exponentially growing populations. For this reason, an explicit stochastic simulation such as the one outlined in Algorithm 3 is infeasible. Instead we implement a deterministic model of epigenetic memory. This model, which is outlined below, is dependent on matrix multiplication and as such is much less computationally expensive (See Figure 5.5).

Deterministic Simulation

The (age–dependent) internal molecular state for each individual must be known at the time of reproduction for each individual in the system. As such, the Wright–Fisher model presented in the previous chapter is unsuitable. Instead, we require a model in which the intracellular state is tracked and the timesteps are sufficiently fine–grained that the intracellular state can be accurately updated. The most conceptually simple approach to modelling such dynamics is through a non–spatial, age–structured individual–based–model (IBM), for example through an explicit simulation of the continuous–time population dynamics using a Gillespie–like algorithm [113, 115] (See Appendix B for details). However, such a model of population dynamics is prohibitively slow to simulate for a large or exponentially growing population. For this reason we develop a deterministic approximation to such an explicit stochastic model.

We begin by decomposing the population of phenotype A and phenotype B individuals into sub-populations A_n, B_n according to the internal number of molecules n. Denote by $A_n(t)$ (respectively $B_n(t)$) the number of individuals with phenotype A (resp. B) and internal state n at time t. By our calculation in Equation 5.20, each individual of phenotype A (respectively B) must exist in a sub–population A_n (resp. B_n) for $0 \leq n \leq M_{\max}$. The decay reactions

Algorithm 3 Continuous time population dynamics for the epigenetic memory model using the Gillespie algorithm

procedure SIMULATE(T_{end}, hedge, x_b, y_b, σ, N_0)
  $n \leftarrow 0$, pop $= []$, $N_A \leftarrow 0, N_B \leftarrow 0$
  while $n < N_0$ **do**
    Sample $p \sim \text{Unif}(0,1)$
    if $p <$ hedge(x_b, y_b) **then**
      pop.append($(0, A, x_b + y_b, 0)$), $N_A \leftarrow N_A + 1$
    else
      pop.append($(0, B, 0, x_b + y_b)$), $N_B \leftarrow N_B + 1$
    end if
    $n \leftarrow n + 1$
  end while
  $t \leftarrow 0$
  while $t < T_{\text{end}}$ **do**
    Compute $p_{\text{tot}} = N_A f_A^{\mathbf{e}} + N_B f_B^{\mathbf{e}} + N_A d_A^{\mathbf{e}} + N_B d_B^{\mathbf{e}}$
    Sample $t_{\text{next}} \sim \text{Exponential}\left(\frac{1}{p_{\text{tot}}}\right)$
    $t \leftarrow t + t_{\text{next}}$
    Sample $p \sim \text{Unif}(0,1)$
    temp $\leftarrow \frac{N_A f_A^{\mathbf{e}}}{p_{\text{tot}}}$
    if $p <$ temp **then**
      Remove a random $(t_0, A, x_n, 0) \in$ pop, $N_A \leftarrow N_A - 1$
      Sample x_I from $p(x_n, t - t_0)$ (Equation 5.10)
      Sample $p_1, p_2 \sim \text{Unif}(0,1)$
      for $i = 1, 2$ **do**
        if $p_i <$ hedge$(\frac{x_I}{2} + x_b, y_b)$ **then**
          pop.append($(t, A, \frac{x_I}{2} + x_b + y_b, 0)$), $N_A \leftarrow N_A + 1$
        else
          pop.append($(t, B, 0, \frac{x_I}{2} + x_b + y_b)$), $N_B \leftarrow N_B + 1$
        end if
      end for
    end if
    temp $\leftarrow$ temp $+ \frac{N_B f_B^{\mathbf{e}}}{p_{\text{tot}}}$
    if $p <$ temp **then**
      Remove a random $(t_0, B, 0, y_n) \in$ pop, $N_B \leftarrow N_B - 1$
      Sample y_I from $p(y_n, t - t_0)$ (Equation 5.10)
      Sample $p_1, p_2 \sim \text{Unif}(0,1)$
      for $i = 1, 2$ **do**
        if $p_i <$ hedge$(x_b, \frac{y_I}{2} + y_b)$ **then**
          pop.append($(t, A, \frac{y_I}{2} + x_b + y_b, 0)$), $N_A \leftarrow N_A + 1$
        else
          pop.append($(t, B, 0 \frac{y_I}{2} + x_b + y_b)$), $N_B \leftarrow N_B + 1$
        end if
      end for
    end if
    temp $\leftarrow$ temp $+ \frac{N_A d_A^{\mathbf{e}}}{p_{\text{tot}}}$
    if $p <$ temp **then**
      Remove a random individual of phenotype A from pop, $N_A \leftarrow N_A - 1$
    end if
    temp $\leftarrow$ temp $+ \frac{N_B d_B^{\mathbf{e}}}{p_{\text{tot}}}$
    if $p <$ temp **then**
      Remove a random individual of phenotype B from pop, $N_B \leftarrow N_B - 1$
    end if
  end while
end procedure

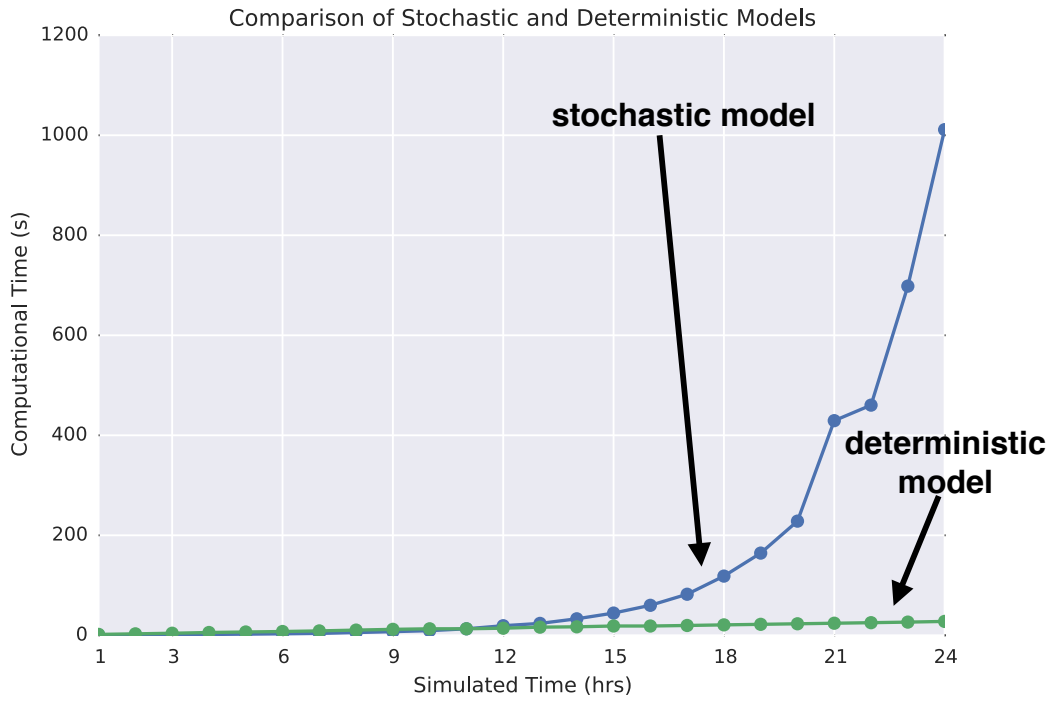

Fig. 5.5 **Timing Comparison for the Stochastic and Deterministic Models.** A comparison in the computational time taken to simulate a number of hours in the stochastic (blue) and deterministic (green) model is shown. Each value is determined as the average of 100 runs. All simulations were performed on a 2.6 GHz Intel Core i7 processor (Macbook Pro (2012)). The explicit stochastic simulation exhibits exponential complexity whilst the deterministic simulation is linear.

(Equation 5.7) determine the probability that an individual of type A_n at time t enters a different subgroup over a small time step Δt (ignoring birth and death) according to

$$\mathbb{P}(\text{The individual remains in group } A_n) = 1 - \sigma n \Delta t + \mathscr{O}(\Delta t^2) \tag{5.21}$$

$$\mathbb{P}(\text{The individual enters group } A_{n-1}) = \sigma n \Delta t + \mathscr{O}(\Delta t^2) \tag{5.22}$$

$$\mathbb{P}(\text{The individual enters a group other than } A_{n-1}) = \mathscr{O}(\Delta t^2). \tag{5.23}$$

For sufficiently small time steps we may assume that at most one event occurs over each time step and that the factor $\mathscr{O}(\Delta t^2)$ may be ignored. Under these assumptions a deterministic model for the populations A_n and B_n can be constructed by repeated population updates over small discrete timesteps Δt.

Suppose individuals of phenotype A, B die with probability d_A, d_B (resp.), and reproduce with probability b_A, b_B (resp.) over the timestep Δt. At timestep t, the contributions to the populations at the next timestep, $A(t+\Delta t)$, $B(t+\Delta t)$ come from the intracellular molecular decay of individuals with state $n+1$ at time t, the reproduction of individuals with intracellular state $n_p = 2(n-x_b-y_b)$ or $n_p = 2(n-x_b-y_b)+1$ and the survival of individuals of state n at time t which experience no decay. The populations $A(t+\Delta t)$, $B(t+\Delta t)$ are also diminished by death events.

Define a length $2M_{\max}+2$ population vector $\mathbf{v}$ wherein the first $M_{\max}$ entries, $\mathbf{v}[0\ldots i\ldots M_{\max})$, correspond to the sub–populations of individuals with phenotype A and intracellular molecule number i. The second $M_{\max}$ entries, $\mathbf{v}[M_{\max}+1\ldots M_{\max}+1+i\ldots M_{\max}+2]$, correspond to the sub–populations of individuals with phenotype B and intracellular molecule number i. Denote by $\mathbf{v}(t)$ the population state at a discrete timestep t, where time proceeds in discrete intervals of $\Delta t = 1$ s. The matrix population dynamics can be encapsulated as a matrix $\mathbf{P}$ wherein $P_{i,j}$ specifies the contribution of sub–population j to sub–population i over the small timestep Δt (for $0 \leq i,j \leq 2M_{\max}+2$). Denote by $p(x_0,y_0)$ the (GP–mapping dependent) probability of switching to the stable state consisting of solely x-molecules (phenotype A), from the initial condition (x_0,y_0).

An explicit equational form for the matrix $\mathbf{P}$ is difficult to represent concisely, because the case of odd and even intracellular numbers must be treated independently. Instead, it is much clearer to present an intuitive algorithmic approach to constructing the matrix from the contributions of birth, death and intracellular decay. The construction is made as follows:

Unchanged individuals, less death and decay: For $0 \leq k < M_{\max}$ set

$$P_{k,k} = (1-k\sigma)(1-d_A)$$
$$P_{k+M_{\max}+1,k+M_{\max}+1} = (1-k\sigma)(1-d_B)$$

Decay from above: For $0 \leq k < M_{\max}$ set

$$P_{k,k+1} \mathrel{+}= (1-d_A)(k+1)\sigma$$
$$P_{k+M_{\max}+1,k+M_{\max}+2} \mathrel{+}= (1-d_B)(k+1)\sigma$$

Births from even state individuals: For $0 \leq k \leq M_{\max}$, if k is even:

$$P_{\frac{k}{2}+x_b+y_b,k} \mathrel{+}= (1-d_A)b_A p\left(\frac{k}{2}+x_b, yb\right)$$
$$P_{\frac{k}{2}+x_b+y_b,k+M_{\max}+1} \mathrel{+}= (1-d_B)b_B p\left(x_b, \frac{k}{2}+y_b\right)$$
$$P_{\frac{k}{2}+x_b+y_b+M_{\max}+1,k} \mathrel{+}= (1-d_A)b_A\left(1-p\left(\frac{k}{2}+x_b, y_b\right)\right)$$

$$P_{\frac{k}{2}+x_b+y_b+M_{\max}+1,k+M_{\max}+1} \mathrel{+}= (1-d_B)b_B\left(1-p\left(x_b, \frac{k}{2}+y_b\right)\right)$$

Births from odd state individuals: For $0 \leq k \leq M_{\max}$, If k is odd:

$$P_{\frac{k}{2}+x_b+y_b,k} \mathrel{+}= \frac{1}{2}(1-d_A)b_A p\left(\frac{k}{2}+x_b, y_b\right) + \frac{1}{2}(1-d_A)b_A p\left(\frac{k}{2}+x_b+1, y_b\right)$$

$$P_{\frac{k}{2}+x_b+y_b,k+M_{\max}+1} \mathrel{+}= \frac{1}{2}(1-d_B)b_B p\left(x_b, \frac{k}{2}+y_b\right) + \frac{1}{2}(1-d_B)b_B p\left(x_b, \frac{k}{2}+y_b+1\right)$$

$$\begin{aligned} P_{\frac{k}{2}+x_b+y_b+M_{\max}+1,k} \mathrel{+}= & \frac{1}{2}(1-d_A)b_A\left(1-p\left(\frac{k}{2}+x_b, y_b\right)\right) \\ & + \frac{1}{2}(1-d_A)b_A\left(1-p\left(\frac{k}{2}+x_b+1, y+b\right)\right) \end{aligned}$$

$$\begin{aligned} P_{\frac{k}{2}+x_b+y_b+M_{\max}+1,k+M_{\max}+1} \mathrel{+}= & \frac{1}{2}(1-d_B)b_B\left(1-p\left(x_b, \frac{k}{2}+y_b\right)\right) \\ & + \frac{1}{2}(1-d_B)b_B\left(1-p\left(x_b, \frac{k}{2}+y_b+1\right)\right) \end{aligned}$$

This model construction resembles the construction of a master equation in the study of statistical mechanics [86], where the dynamics over a small time step δt are first defined and then the limit of $\delta t \to 0$ is taken to derive differential equations. We do not follow this approach as, if we did, the equations would need to be solved through numerical methods (which can be more computationally expensive than matrix multiplication). Further, such numerical methods would discretise time, essentially reconstituting the explicit discrete system derived above. Of course, where the equations governing a dynamic system are complicated, sophisticated numerical techniques are required to avoid numerical instabilities and guarantee precision in a solution. Despite the somewhat dense notation used to define the system **P**, each sub–population is governed by linear equations, and sophisticated numerical techniques are not needed.

The matrix **P** determines an update to a population in state $\mathbf{v}(t)$ according to

$$\mathbf{v}(t+\Delta t) = \mathbf{P}\mathbf{v}(t). \tag{5.24}$$

An example of the growth dynamics, phenotype distributions and intracellular distributions for this model are depicted in Figure 5.6. Figure 5.7 shows how the stationary distribution of intracellular states in the model varies with σ.

The Strong–Genetic Model

The strong genetic model can be constructed by explicitly setting the inherited molecules, k, to zero inside the GP–mapping probability calculation $p(.)$ in the construction of the matrix **P**. The epigenetic–memory model was motivated as an extension of the strong–genetic model in which the assumption of complete depletion of the molecules x and y is relaxed and we can

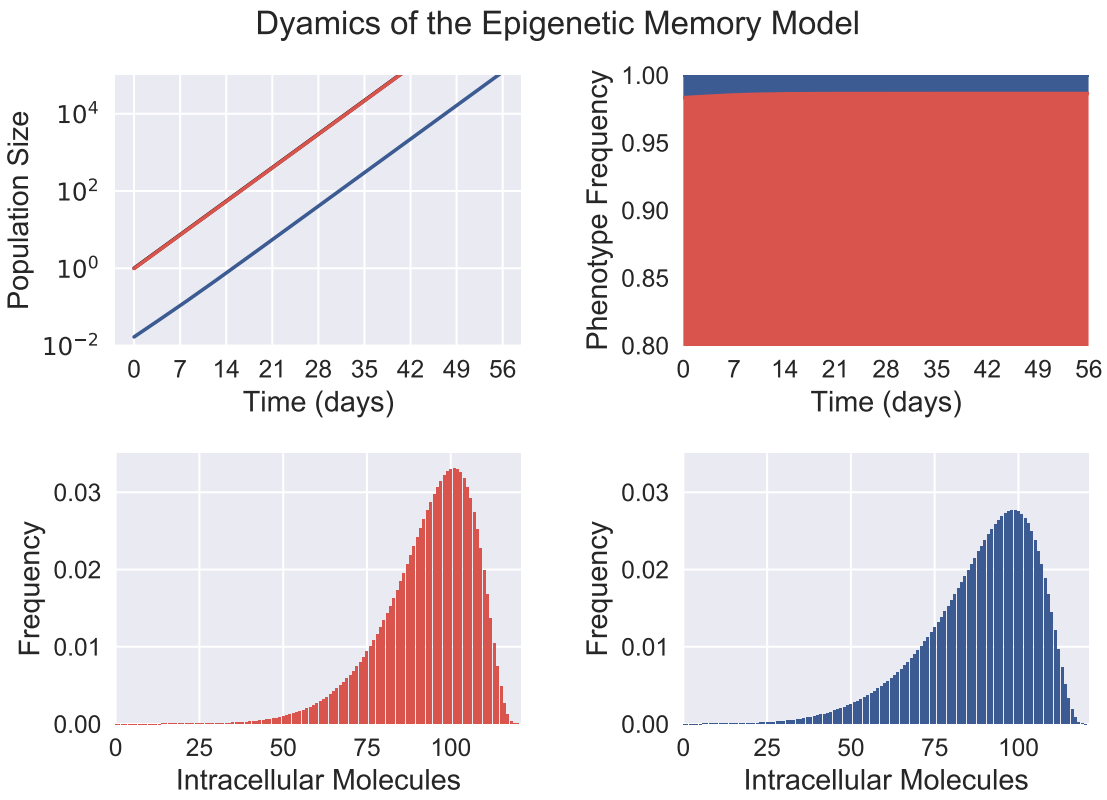

Fig. 5.6 **Dynamics of the epigenetic memory model.** The growth dynamics, phenotype distribution and intracellular state distributions of a population endowed with epigenetic–memory bet–hedging are shown. The memory decay parameter is $\sigma = 6.77 \times 10^{-7}$ s^{-1}. The growth (top right) and frequency (top left) of phenotype A is depicted in red. Phenotype B is depicted in blue. Bottom left shows the intracellular state distribution amongst phenotype A individuals. Bottom right shows this state distribution for phenotype B individuals.

view the strong–genetic model as equivalent to the epigenetic–model as σ grows[4]. Figure 5.7 shows an example of this relationship wherein taking $\sigma = 6.77 \times 10^{-5}$ s^{-1} results in very few individuals having more than zero intracellular molecules.

Constrained Genetic Memory Model

A constrained memory model can be constructed by setting $k = 0$ in all calculations of offspring probability (as for the strong genetic model) and switching the probabilities $p(.)$ to $1 - p(.)$ such that the genotype determines the likelihood of an individual taking the phenotype of its parent.

Exponential and Constrained Growth

The models above described exponential growth without bound. Such growth dynamics may be applicable in certain circumstances, for example the early growth stages of tumour, or the

[4]In the limit of large σ, the assumptions $\rho_j >> \sigma, j \in \{1,2,3\}$ that phenotypic determination is fast, may breakdown. For this reason, we do not consider $\sigma \to \infty$

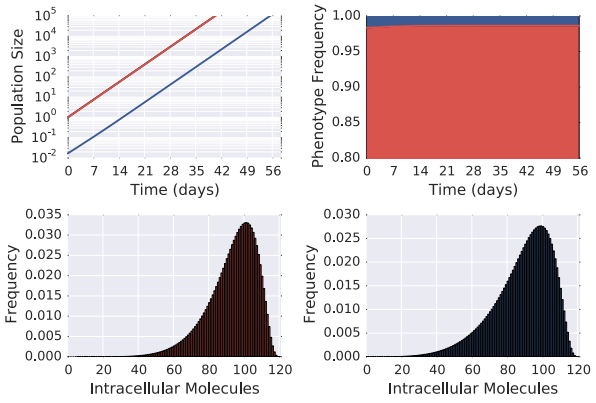

$$\sigma = 6.77 \times 10^{-7}$$

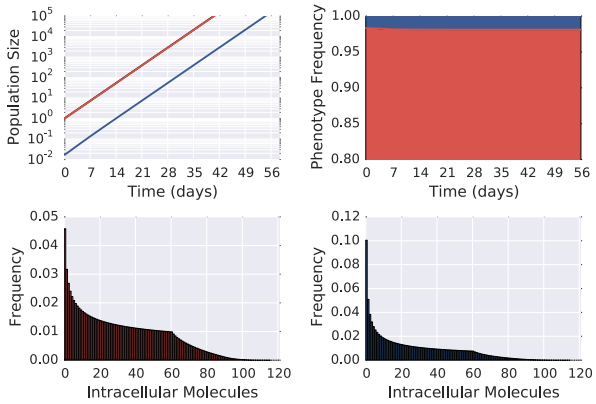

$$\sigma = 6.77 \times 10^{-6}$$

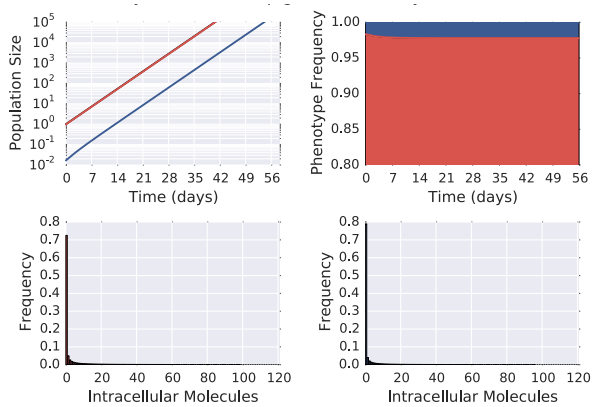

$$\sigma = 6.77 \times 10^{-5}$$

Fig. 5.7 **Varying σ alters the the stationary distribution of intracellular states within the population.** By increasing the decay rate σ, the stationary distribution of intracellular molecule numbers in the population is skewed towards zero, resulting in a diminished memory component. As σ grows, the epigenetic–memory mode of bet–hedging approximates to the strong–genetic form considered in the previous chapter. An example of this phenomenon with $\sigma = 6.77 \times 10^{-5}$ s^{-1} is shown.

outgrowth of small bacterial population in a new location. However, for many circumstances growth is likely to be inhibited by competition for substrate, physical space or some other ecological factor. To account for this in our modelling we can define a carrying capacity K and scale the reproductive rates by a factor $\left(1-\frac{A(t)+B(t)}{K}\right)$. Such a scaling results in logistic growth and makes the matrix $\mathbf{P}$ dependent on the population size. Here, we focus on exponentially growing populations and in Chapter 6 briefly explore the impact of carrying capacity on the population dynamics.

5.3 Parametrisation

We parametrise the model, as before, to qualitatively represent the dynamics of a population in which two phenotypes – a more proliferative, drug–sensitive phenotype A and a slowly proliferating, drug–resistant one B – arise stochastically in the population. We consider a parametrisation corresponding to tumour growth under experimental conditions. Specifically, we assume a weekly growth rate approximately corresponding to the population doubling in size twice. This parametrisation is chosen as a means to explore the effects of toxicity–induced dose–restrictions (common in cancers) and the potential for clinically viable drug fractionation strategies (for example, comparing weekly and daily doses). Biological justification for a weekly growth rate of approximately two tumour doublings comes from empirical observations of cancer cell-lines, both *in vitro* and *in vivo*, which give approximate tumour doubling times of between 1 and 4 days [273].

As the population update is given by the matrix $\mathbf{P}$ as defined for an arbitrary timestep (which we take as 1s), a scaling of this time step (for example to 4s or 8s) can be made to yield a parametrisation more closely resembling a slower growing disease. Under such a transformation, our results are unchanged and it is only the units of the average growth rate that differ. Thus, the results derived here are applicable to both bacterial infections and cancers (as well as other diseases) through an appropriate scaling of time.

We take the parameters for the proliferative phenotype as

$$b_A = 4.74 \times 10^{-6}\ \mathrm{s}^{-1} = 0.41\ \mathrm{day}^{-1} \tag{5.25}$$

$$d_A = 1.39 \times 10^{-6}\ \mathrm{s}^{-1} = 0.12\ \mathrm{day}^{-1} \tag{5.26}$$

which yields an average reproduction time of 2.44 days and a weekly log growth rate of

$$F_A = \log\left(\left((1+b_A)(1-d_A)\right)^{60\times60\times24\times7}\right) = 2.03 \tag{5.27}$$

For persister–like cells, no known approximation to the growth parameters are known. We assume that persister phenotype cells are characterised by a 10 fold reduction in both the birth

and death rates. Thus, we assume

$$b_B = 4.74 \times 10^{-7}\ \mathrm{s}^{-1} = 0.041\ \mathrm{day}^{-1} \tag{5.28}$$

$$d_B = 1.39 \times 10^{-7}\ \mathrm{s}^{-1} = 0.012\ \mathrm{day}^{-1}. \tag{5.29}$$

This yields an average reproduction time of 24.4days. The weekly log growth of the persister cells is then

$$F_B = \log\left(\left((1+b_B)(1-d_B)\right)^{60\times60\times24\times7}\right) = 0.203. \tag{5.30}$$

Finally, for a drugged environment we assume that persister–like cells are entirely impervious to the effects of a drug (b_B and d_B are unchanged) and that for proliferative cells, the drug is cytotoxic, effecting only d_A and leaving b_A unchanged. We take the drug induced death parameter to correspond to an eight fold increase in death rate,. Thus, under the treatment of the drug a population consisting only of phenotype A individuals is expected to shrink, the weekly log growth is

$$F_A^{\mathrm{drug}} = \log\left(\left((1+b_A)(1-d_A^{\mathrm{drug}})\right)^{60\times60\times24\times7}\right) = -3.86. \tag{5.31}$$

It should be noted that these parameters are not derived from empirical observations of bet–hedging, as few instances have been comprehensively studied outside of demonstrating the existence of persister–like cells. Instead, we determined these parameters from two representative criteria. The first is that the persisters are impervious to the drug and reproduce much more slowly than the sensitive cells. The second criterion is that the reproductive rate is sufficiently slow that we may consider fast– and slow– switching cycling environments that are clinically achievable.

As discussed in Section 5.2.1, we take $\sigma = 6.77 \times 10^{-7}\ \mathrm{s}^{-1}$ to correspond to an expected time of 10 weeks to achieve full depletion of the memory component. Later, we will explore the fitness effect of varying σ under different therapeutic regimes and molecular switches. To do this we will vary σ over the range $[10^{-11}\ \mathrm{s}^{-1}, 10^{-2}\ \mathrm{s}^{-1}]$.

Finally, we must determine initial conditions for the population size, molecular states and phenotype distribution. Owing to the form of our model, the precise population size is irrelevant and can be scaled. Thus, we initialise all of our simulations with an initial population size of $N = 1.0$. The initial phenotype distribution is more difficult to determine. It is appealing to initialise the population to correspond to a distribution determined solely by the genetic component of bet–hedging ($g = (x_b, y_b)$). However, doing so ignores the life–history of the population which has important implications for the phenotype distribution. Indeed, *in vitro* studies indicate that the age of a stationary–state inoculum of *E coli* can significantly impact on the proportion of persister cells present within it (see Luidalepp et al. [210]). This phenomenon is explored in more detail in Chapter 6. As the model we present here admits a stationary state for the phenotype distribution, both in continued exponential growth and growth towards a carrying capacity, and both with and without the memory component, we

initialise the phenotype distribution to match this steady state. In practice, this is achieved by first running a simulation to the steady–state phenotype distribution from an arbitrary starting population and then renormalising the population size to $N = 1.0$.

5.3.1 Robustness to Parameter Variation

Throughout this chapter, results will be derived for a single parameter set outlined above. Here, we provide a brief justification for why these results can be expected to hold for other parameter sets. First, we note that the population dynamics model outlined in Section 5.2.2 is a discrete time model which determines how the population changes over a timestep Δt. For this reason, changing the absolute magnitude of the parameters σ, b_i and d_i (for $i \in \{A,B\}$) whilst keeping their proportions the same will not alter our results. This is because any proportional scaling of the parameters is equivalent to redefining the time step, Δt, and our results are not explicitly dependent on the value of Δt (only the units of any derived values are).

Second, we note that changes to the relationship between the growth parameters, b_i and d_i (for $i \in \{A,B\}$), and the intracellular decay parameter, σ, will likely have a large impact on our results. It is for this reason that we explore alterations to σ in Section 5.7.

Finally, we must consider the relationship between the parameters b_i and d_i (for $i \in \{A,B\}$). Following our argument made above, it is only the relationship amongst parameters and not their absolute values that determine our results (by considering a rescaling of Δt). Thus, to explore the impact of parameter variation we consider fixed b_B and d_B and explore the impact of varying b_A and d_A. For brevity, we present this exploration of parameter robustness in Appendix 2.

A brief overview of the results of this analysis are as follows. For phenotype A, we considered each pair of parameters sampled from

$$b \in \{0.1b_A, b_A, 5b_A, 10b_A\}$$
$$d \in \{0.1d_A, d_A, 2d_A, 5d_A\}$$

for which the growth rate of phenotype A is greater than 1. These parameters were chosen to represent a broad range of scenarios for the relationship between the fitness of persister and non–persister cells, ranging from non–persister exhibiting a weekly growth rate only $2.9\times$ greater than non–persisters at the lower extreme and $141.1\times$ greater than non–persisters at the upper extreme. Using these parameters, the results of Section 5.4 presented below were re–derived. Through this exploration we found that the key results, specifically that organismal fitness is dependent on the specific combination of switch, genotype and memory mechanism, are preserved. However, we also found that bet–hedging is optimal for only certain parameter sets. Specifically, where the growth rate of phenotype A is sufficiently high, the fitness optimum for a fluctuating environment is attained for a monomorphic population of phenotype A. Following previously published theoretical work exploring bet–hedging [232]

we predict that an optimum bet–hedging strategy is in fact not monomorphic, but that the associated likelihood of phenotype A is sufficiently close to 1 that it cannot be captured by our model with $g_{\max} = 60$. From these observations we conclude that the results presented below will hold for a large space of parameter values encoding the relationship between persister and non–persister cells arising from bet–hedging, but that they are not universal. For example, we expect our predictions to fail to hold when persister cells are massively less fit than non–persisters. This observation further emphasises the need to empirically identify and characterise bet–hedging populations, especially in cases where they drive drug resistance. Further details are presented in Appendix 2.

5.4 The Selective Advantage of Memory

Having described a potential mechanism for phenotypic memory in bet–hedging and derived a model in which to study it, we now turn our attention to the evolutionary benefit of phenotypic memory. We compare the fitness of bet–hedging populations endowed with different mechanisms of memory, different molecular switches and different genotypes, under varying environmental regimes. We highlight that phenotypic memory can offer an adaptive advantage under a broad range of circumstances, but critically, the existence and magnitude of any fitness advantage contribution from memory is strongly dependent on the combination of GP–mapping, (epi–)genotype and environment.

Before we proceed it should be noted that both the strong adaptationist (or 'Panglossian') paradigm, that a trait providing adaptive benefit to an organism must have arisen through natural selection, or the even stronger analogue, that any adaptation that *could* offer a selective advantage *will* arise through natural selection, are both fallacious and predicated on a belief in the omnipotent 'guiding hand' of natural selection [125]. As such, the arguments provided in the following section, taken alone, provide no evidence supporting the existence of adaptive epigenetically inherited memory. However, coupled with the empirical experimental evidence for memory, and a theoretical but biologically reasonable *mechanism* through which memory can arise, the following evolutionary arguments lend credence to the potential existence of memory mechanisms.

Considerations of the adaptive benefit of phenotypic memory represent a first step towards exploring its potential implications in driving drug resistance, and thus, in determining therapies to mitigate resistance. In exploring clinical implications, we cannot derive more effective memory–dependent therapies without first providing strong empirical evidence that memory is a driver of drug resistance. However, we can explore the robustness of previously suggested therapeutic strategies. There is merit to this approach, as we can highlight a clinically important phenomenon: that the efficacy of a therapeutic regime is strongly dependent on the underlying driver of drug resistance (i.e. the GP–mapping), and further, incorrect assumptions may lead to therapeutic regimes that perform substantially worse than other available strategies. From this perspective, our modelling serves both to provide an argument to perform further,

model–driven experiments and also to provide a cautionary warning against transplanting efficacious therapeutic strategies between diseases, or potentially, even between patients.

5.4.1 Fixed Environments

We begin by considering the fitness of populations endowed with the different modes of bet–hedging in fixed environments. As in Chapter 4, we consider long–term, stable environments that strongly select for one of the two phenotypes. The environments are chosen to correspond to a hospitable and a drug–treated environment as described by the parameters introduced in Section 5.3. As discussed in Chapter 4, the evolutionary loss of bet–hedging can result in extinction when a later drug treatment occurs, but the structure of the GP–mapping can reduce evolvability, prevent evolutionary loss of bet–hedging and preserve drug resistance. We explore the impact of different mechanisms of phenotypic memory on this result and explore the evolutionary relationship between memory, molecular switch and genotype. As before, we consider genotypes satisfying $x_b + y_b = 60$ and each of the four molecular switches (AM, DC, DCx and DCy) defined in Figure 4.2. Further, we now consider each of the different mechanistic modes of bet–hedging described in Section 5.1.

Mirroring previous studies (for example, Jablonka et al. [164]) we define fitness as the expected log population growth rate over a fixed time–period. As the discrete time dynamics of our population are defined by a matrix model, the Perron–Frobenius theorem [205] gives that the growth rate is precisely the spectral radius of the population dynamics matrix. The model derived above simulates population dynamics on a molecular time scale to capture the dynamics of epigenetic decay of intracellular molecules. Measuring fitness as an expected per–second growth rate makes little intuitive sense for a population with a reproductive cycle on the order of days. Instead, we calculate fitness to be the log growth rate over a one week period

$$F = \log(\rho(\mathbf{P}^{60\times 60\times 24\times 7})), \tag{5.32}$$

where ρ is the spectral radius of the matrix.

Figure 5.8 shows plots of the population fitness, as determined by Equation 5.32, for population genotype, $g = (x_b, 60 - x_b)$, as x_b varies with $x_b \in \{0, \ldots, 60\}$. The parameters are as described in Section 5.3. Fitness curves for the epigenetic memory (green, solid), constrained–memory (orange, dotted) and strong–genetic (purple, dashed) models are shown for both the hospitable and drug–like environments. For bet–hedging with constrained memory and $x_b = 60$, an individual will always produce offspring with the same phenotype. This is equivalent to a clonal model of evolution without mutation. Thus, where $x_b = 60$, there exist multiple fitness values dependent on the initial distribution of phenotypes in the population. The value shown in the plot corresponds to the initial distribution resulting in maximum fitness (all A in the hospitable environment and all B in the inhospitable environment).

The key observation derived from these simulations is that the fitness is dependent on the combination of genotype, memory mechanism and molecular switch. Critically, this

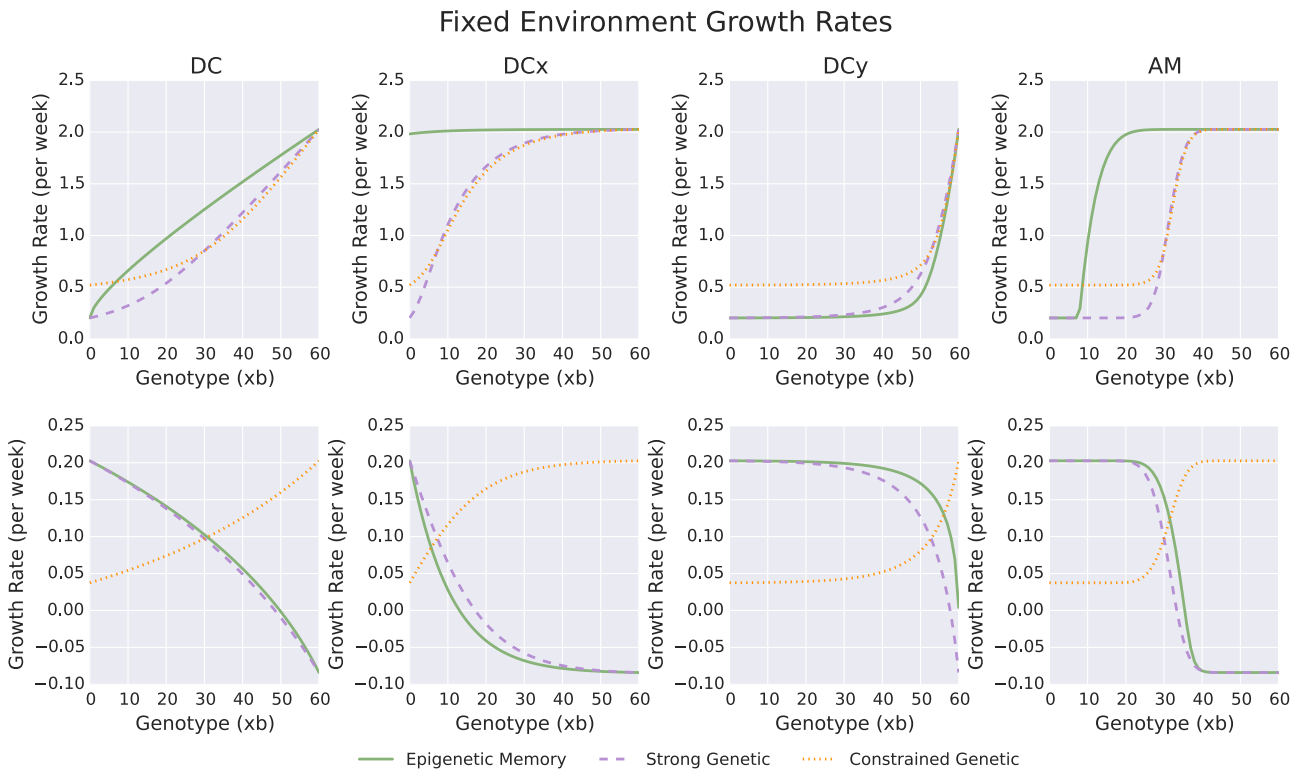

Fig. 5.8 **Memory, molecular switch and genotype combine to determine fitness in constant environments.** Subplots show the relationship between genotype (x_b) and growth rate for populations endowed with strong–genetic (purple, dashed), constrained memory– (orange, dotted) and epigenetic memory– (green, solid) bet–hedging. The curves are shown for the hospitable (top row) and drug–exposed (bottom row) environments. The molecular switch varies between the columns.

relationship is not additive. For example, in the hospitable environment, where $x_b \geq 30$, the population endowed with the epigenetic memory mechanism is optimal for the DC, DCx and AM switches. For the DCy switch, it is constrained memory that is optimal. However, when the environment is switched, epigenetic memory becomes deleterious (in comparison to the strong genetic mode of bet–hedging) for the DCx switch and constrained memory becomes optimal in large parts of the genotype space for all of the switches.

To capture the full extent of this non–additive relationship between the mechanistic components of bet–hedging, we generated the epi–fitness landscapes shown in Figure 5.9(A) and Figure 5.10(A). In these landscapes, each combination of genotype, memory mechanism and molecular switch is represented by a square and arrows indicate fitness conferring mutations. These landscapes reveal the epistatic fitness effects of mutations and indicate that, as in the genetic fitness landscapes explored in Chapter 3, all loci must be considered together in order to predict evolution.

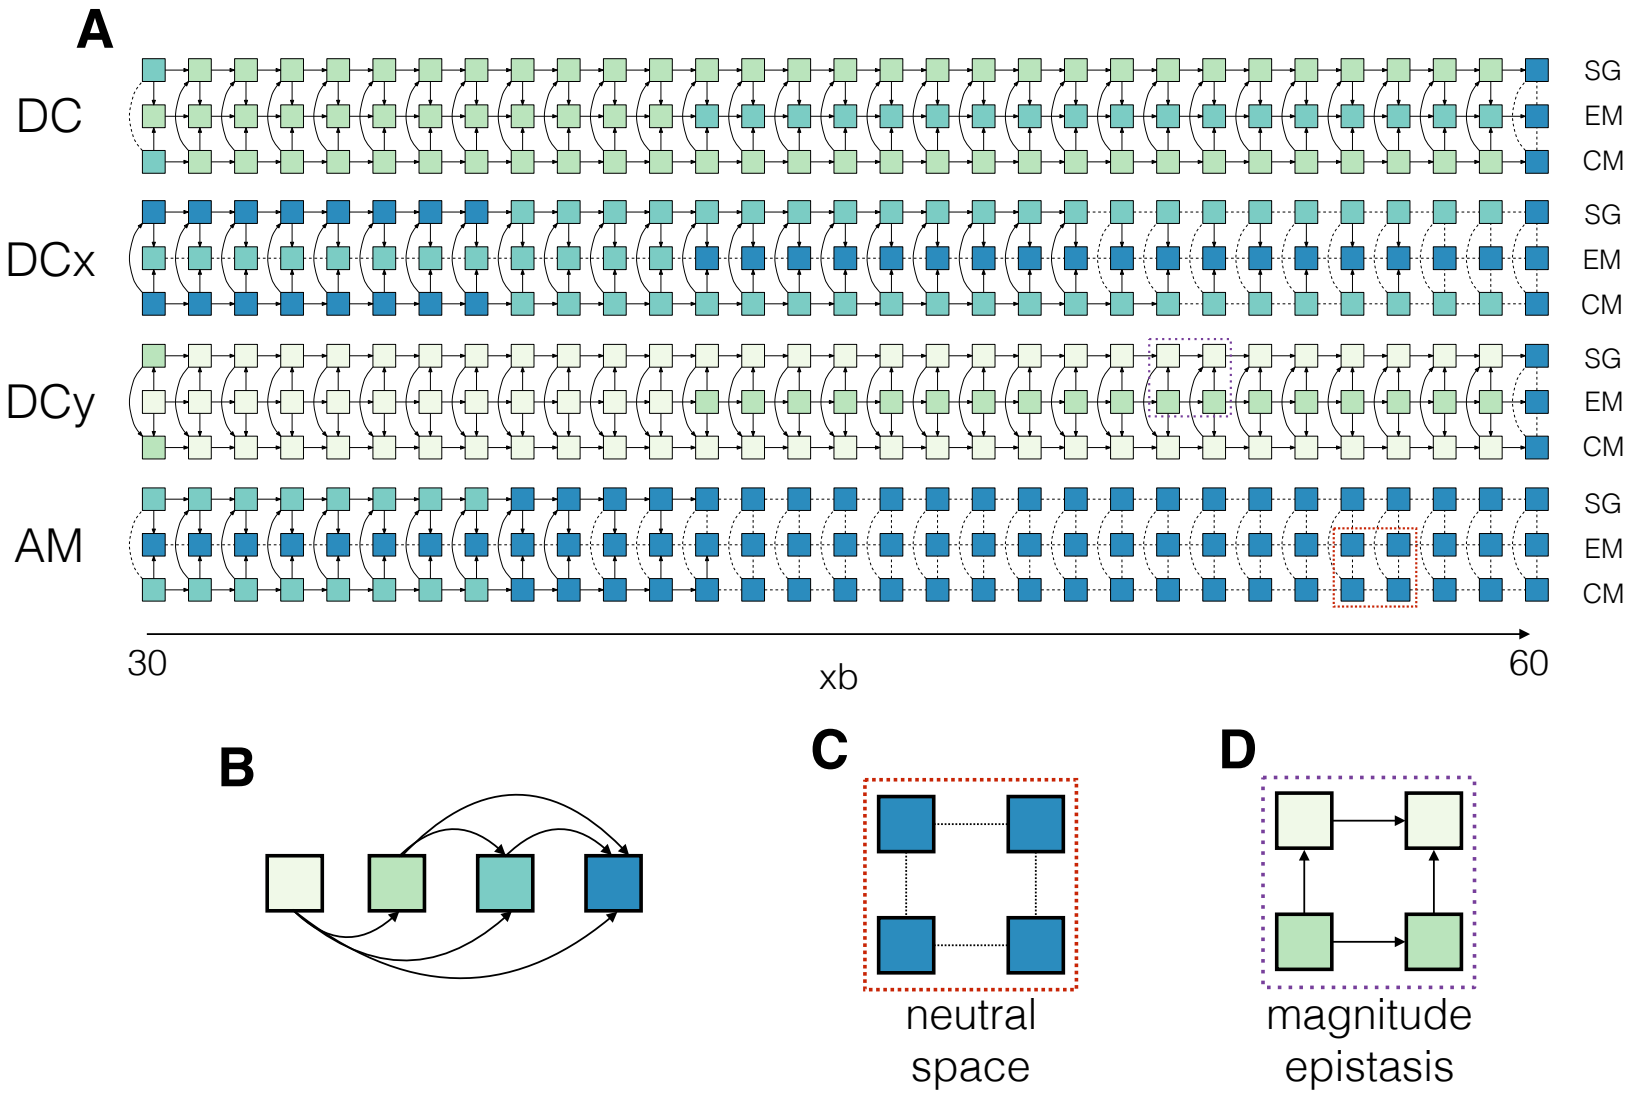

Fig. 5.9 **The fitness landscape for molecular switch, inheritance and genotype in the hospitable environment**. (A) A restricted portion of the fitness landscape is shown. Each square corresponds to a combination of genotype (x_b, varying from 30 to 60 on the x-axis), molecular switch (left–hand row label) and memory type (right–hand row label). Arrows indicate mutations that confer a fitness advantage, dashed lines indicate neutral mutations (defined as mutation resulting in a less than 0.1% change in fitness). The square colour indicates the fitness rank amongst the molecular switches according to the schema in (B). Where the fitness values differ by less than 0.1%, the squares are assigned the same color. (B) Schema depicting how colours represent the relative fitness ordering amongst the molecular switches. (C) An example neutral space in the landscape (marked by a red dashed square in (A)). (D) An example of magnitude epistasis from the landscape (marked by a purple dashed square in (A)).

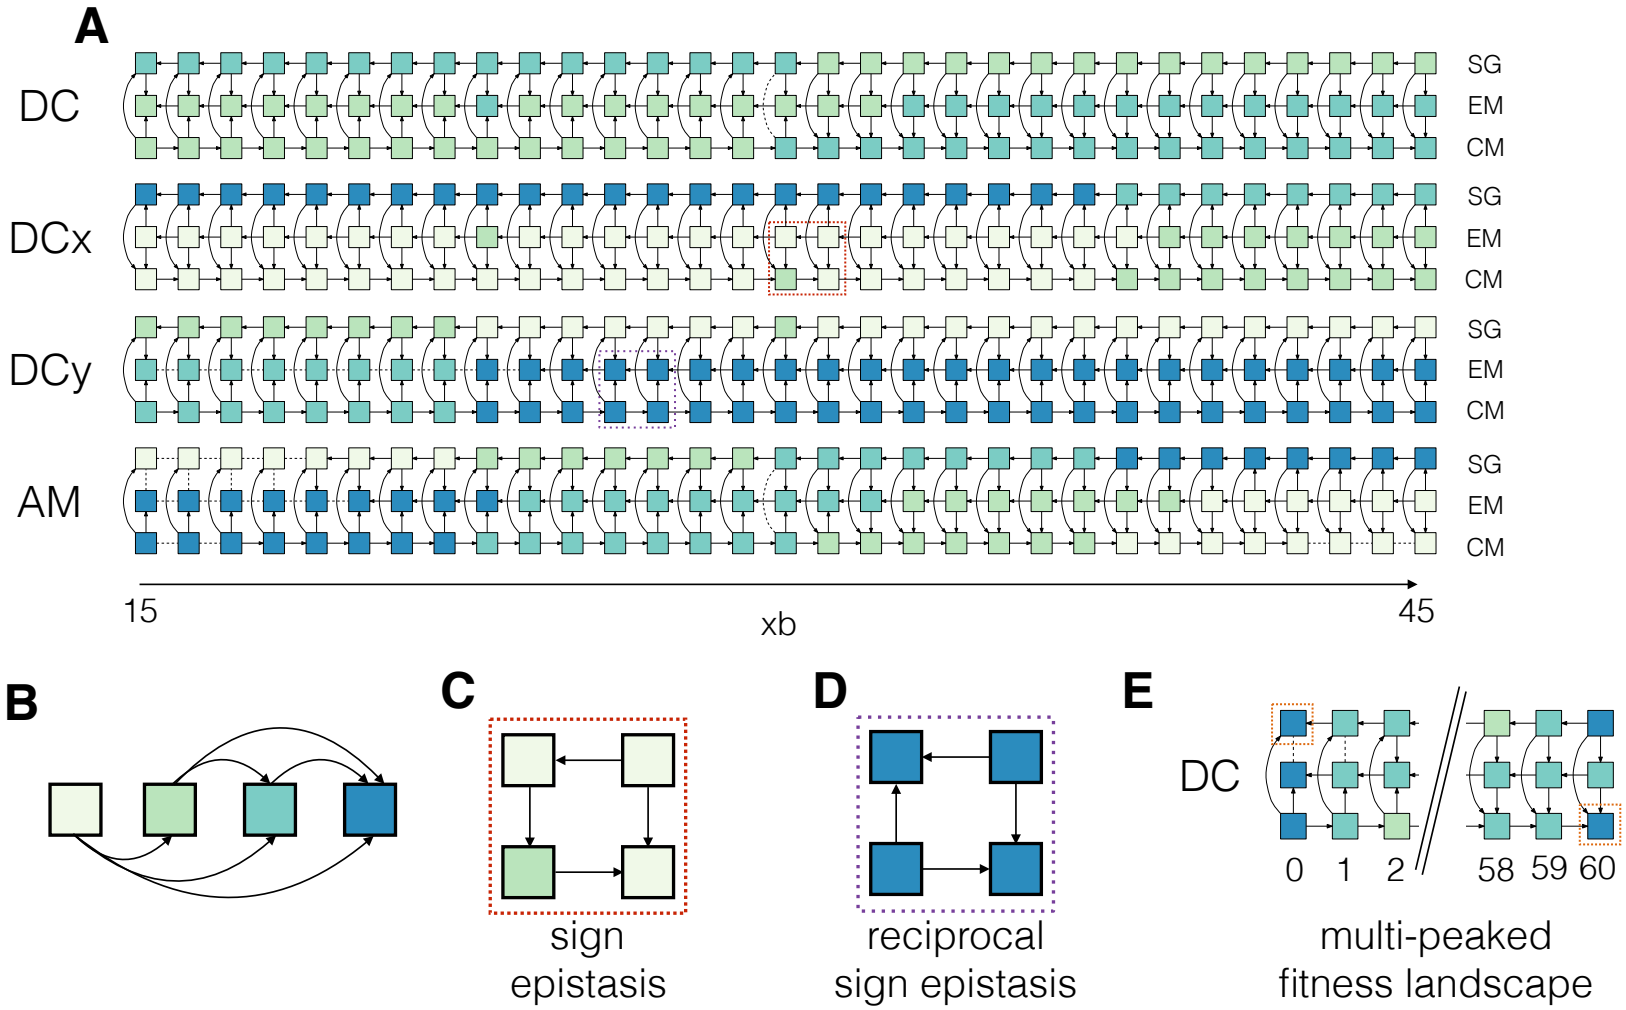

Fig. 5.10 **The fitness landscape for molecular switch, inheritance and genotype in the drug environment**. (A) A restricted portion of the fitness landscape is shown. Each square corresponds to a combination of genotype (x_b, varying from 15 to 45 on the x-axis), molecular switch (left–hand row label) and memory type (right–hand row label). Arrows indicate mutations that confer a fitness advantage, dashed lines indicate neutral mutations (defined as mutation resulting in a less than 0.1% change in fitness). The square colour indicates the fitness rank amongst the molecular switches according to the schema in (B). Where the fitness values differ by less than 0.1%, the squares are assigned the same color. (B) Schema depicting how colours represent the relative fitness ordering amongst the molecular switches. (C) An example of sign epistasis from the landscape (marked by a red dashed square in (A)). (D) An example of reciprocal sign epistasis from the landscape (marked by a purple dashed square in (A)). (E) A view of the landscape around $x_b = 0$ and $x_b = 60$ demonstrating the existence of two local optima of fitness (orange dashed squares).

Neutrality and Epistasis in the Fitness Landscapes

In Chapter 4, we demonstrated that robustness and redundancy in the molecular switches governing bet–hedging can arise through neutral or nearly mutations (Figure 4.4(A)), and argued that the molecular switch itself is subject to evolution. Where bet–hedging or memory are observed, we can ask whether they arose as fitness–conferring adaptations, as an evolutionary 'spandrel' [125] co–arising with some other adaptation, or as fitness–neutral phenomena that achieved prominence through drift. Where bet–hedging arises as a spandrel or neutral mutation, it may have been later co-opted for adaptive benefit through further mutations. The simulations we present here show patterns of such neutrality that hint at the potential for memory to be gained or lost through evolutionary drift. An example of neutrality between different memory mechanisms is shown in Figure 5.9(C). Figure 5.9(A) also contains examples of neutral genotypic mutations (horizontal dashed lines) and neutral switch alterations (preserved colour when the switch is altered).

Consider these fitness lanscapes from the perspective that GP–mappings are the subject of evolution [7]. Using the language of evolutionary theory: the mechanisms of (epigenetic) inheritance and molecular switching interact *epistatically*, just as two point mutations conferring resistance to an antibiotic might. Figures 5.9 and 5.10 show examples of such epistasis. Indeed, the landscape under drug exposure exhibits both sign epistasis (Figure 5.10(C)) and reciprocal sign epistasis (Figure 5.10(D)). As discussed in our review of fitness landscapes (Section 2.2.2), such epistasis has important implications for the accessibility and predictability of evolution, as it can induce rugged fitness landscapes [260, 261]. This is precisely what we observe; the landscape under drug exposure is multi–peaked (Figure 5.10(E)). From this observation we can infer that the different mechanisms of inheritance, specifically the strong–genetic and constrained–memory forms of bet–hedging, may arise stochastically in the evolutionary process. The evolution of these phenomena need not be repeatable or predictable.

These fitness landscapes demonstrate the non–additive fitness contributions of non–genetic mechanisms of phenotypic determination. Coupled with the genetic fitness landscapes studied in Chapter 3, the fitness landscapes can be considered to exist in duplicate – at the genetic and GP–mapping scales. The immediate consequence of this perspective is that a number of results that were derived for the genetic fitness landscapes are equally applicable at the mechanistic level. This demonstrates the power of abstraction in providing evolutionary insight that may be missed through a reductionist approach. The evolutionary effects of interactions between the gene and gene–interaction scales would likely be obscured by the overwhelming complexity of the protein–protein interaction network if we attempted to reconstruct it from empirical observations of intracellular dynamics.

Memory and Evolvability

Consider the paradox of bet–hedging that initially motivated our modelling: if bet–hedging decreases population fitness for extended periods of time, why is it not lost through natural

selection? In Chapter 4, we argued that the structure of the genotype–phenotype mapping can induce diminishing fitness benefits for successive mutations towards a one–phenotype strategy. Thus, the GP–mapping reduces evolvability and hence allows a population to maintain bet–hedging as a survival mechanism. As the GP–mapping is itself subject to evolution, we argued that this property may have itself arisen as an adaptation, the *evolution of reduced evolvability*. Phenotypic memory driven by epigenetic inheritance of intracellular molecules can augment this effect.

Consider now the population endowed with epigenetic memory and under the hospitable environment. For the AM and DCx switch, the introduction of epigenetic memory results in successive mutations $x_b \to x_b + 1$ having a further diminished fitness benefit. Thus, these mutations are less likely to fix in the population when epigenetic memory is present. It follows that memory can act to further reduce evolvability and maintain bet–hedging in the population over evolutionary time scales. This is clearly demonstrated by the increase in the size of the neutral spaces in Figure 5.9(A) (dashed edges) when moving from strong–genetic to epigenetic–memory mode of bet–hedging. Epigenetic inheritance and phenotypic memory can serve as alternative or complementary solutions to the paradox of loss of bet–hedging. Further, memory arising though epigenetic inheritance is doubly advantageous under the AM switch, preventing the loss of bet–hedging and increasing population fitness (in comparison to the strong–genetic case) in both the hospitable and drugged environment.

5.5 Memory in Fluctuating Environments

Previously, bet–hedging strategies have been demonstrated to offer a two–fold adaptive benefit: increasing fitness in fluctuating environments and buffering a population against rare catastrophic environmental change. Theoretical studies of bet–hedging have focussed primarily on the former role, deriving precise conditions on the parameters of phenotypic– and environmental–switching to determine when and how bet–hedging is an evolutionarily optimal strategy [164, 356, 232, 37]. These studies employed a model of bet–hedging with unconstrained epigenetic memory (Figure 5.1(C)), ignoring the GP–mapping and molecular mechanisms driving bet–hedging. Here, we briefly explore the impact of the different modes of bet–hedging memory and different molecular switches on fitness in a fluctuating environment. To our knowledge, ours is the first study to compare these mechanistically–restricted modes of bet–hedging with memory.

The dynamics of a population in a cyclic environment can be determined as follows. Consider a fixed population genotype x_b and denote by e_{hosp} and e_{drug} the environments defined by the parameters defined in Section 5.3. Denote the population update matrices corresponding to these environments by $\mathbf{P}_{\text{hosp}}$ and $\mathbf{P}_{\text{drug}}$. We consider cyclic environments consisting of W weeks of drug exposure followed $4-W$ weeks of drug–free growth (for $W = 0,1,2,3,4$). This cycle is assumed to be repeated *ad infinitum* and the population fitness is measured as the average log weekly growth rate. The population dynamics over a single

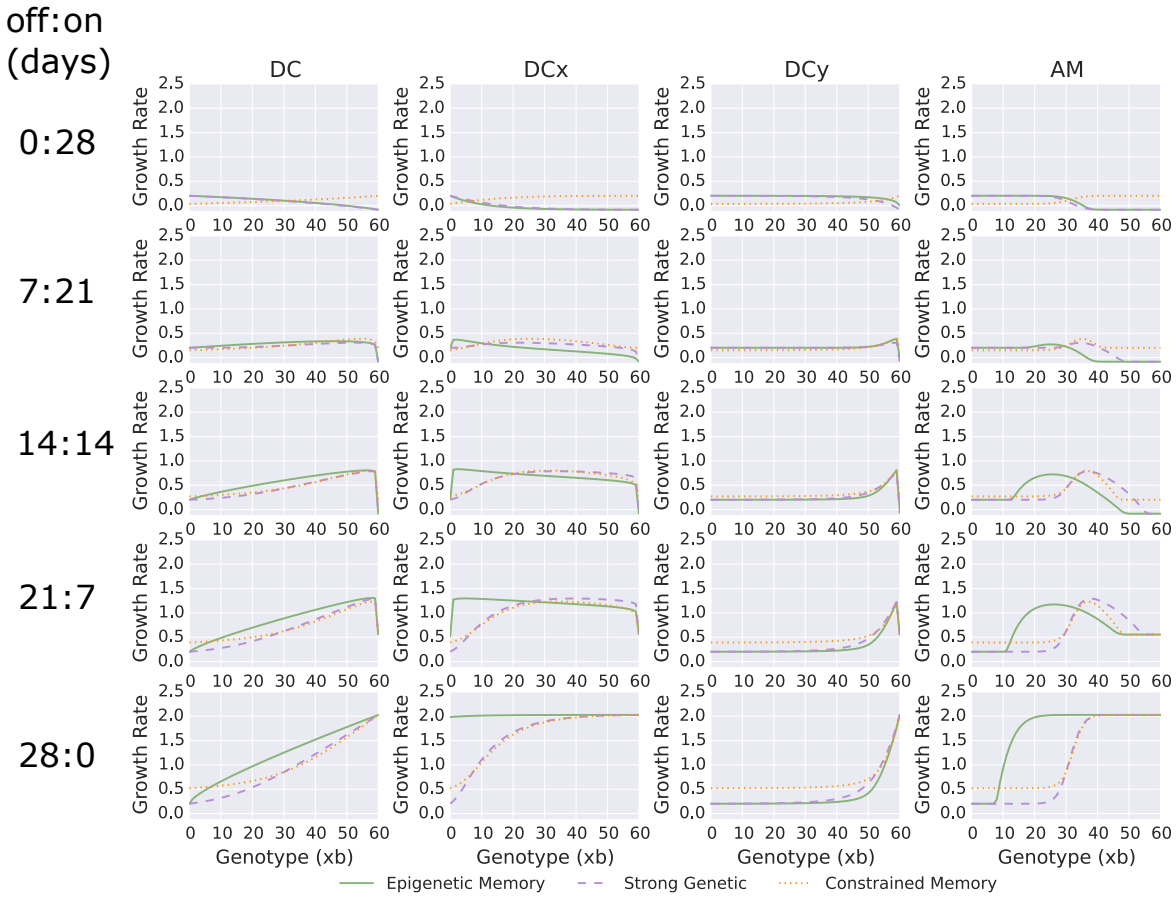

Fig. 5.11 **Memory, molecular switch and genotype combine to determine fitness in fluctuating environments.** Each subplot shows the relationship between population genotype x_b and growth rate in a fluctuating environment. Curves are shown for each memory mechanism (line colour/style, legend) and the molecular switch varies between columns. The environment fluctuates between drug on and drug off over the course of a 28day cycle (which is repeated). The days on/off are determined by the rows (labelled). Note that the first and final rows correspond to the curves in Figure 5.8

cycle are governed by the matrix

$$\mathbf{P}_{\text{cycle}} = \mathbf{P}_{\text{hosp}}^{60\times60\times24\times7\times W}\mathbf{P}_{\text{drug}}^{60\times60\times24\times7\times(4-W)}. \tag{5.33}$$

The expected log weekly growth rate can then be calculated from the eigenvalues of $\mathbf{P}_{\text{cycle}}$ as

$$F = \frac{1}{4}\log(\rho(\mathbf{P}_{\text{cycle}}), \tag{5.34}$$

where ρ denotes the spectral radius and the factor of $\frac{1}{4}$ normalises the log growth rate to a per–week average.

For each W, the genotype–fitness curves are plotted as a grid in Figure 5.11. Here, rows are determined by W, and columns by molecular switch. Each individual plot shows a curve for the strong–genetic (purple, dashed), constrained–memory (orange, dotted) and

epigenetic–memory (green–solid) modes of bet–hedging. Note that the first and last row of this figure correspond to the fitness in constant environments, and thus match the curves shown in Figure 5.8. The intermediate rows indicate how these curves change as a cyclic environment shifts between mostly hospitable and mostly drug–treated.

These fitness curves corroborate the results of previous bet–hedging studies, showing that bet–hedging can offer an adaptive benefit in fluctuating environments. Critically, we see that this result holds for each of the mechanistically restricted modes of bet–hedging memory. The value of x_b at which the optima of fitness occurs is dependent on both the switch and memory mechanisms (the curves under the AM switch exhibit the most striking examples of this). Hence, as in the case of fixed environments, we conclude that a purely genetic perspective, or even a coupled genetic and gene–interaction perspective, is not predictive of fitness. It is only when all of the functional mechanistic components of bet–hedging are considered together that fitness can be determined.

These curves also show that the maximum attainable fitness (for varying x_b) is approximately equal between the different mechanisms of bet–hedging, although the optimal genotype is both switch and memory dependent. Under the assumption that genetic mutations arise more frequently than those to the molecular switch or changes in the mechanism of inheritance, we can assume that the genetic optimum of fitness is likely to be found before structural mutations arise. In this case, mutations to the switch or inheritance mechanisms are likely to be either neutral or deleterious. This suggests, that alterations to these aspects of the GP–mapping are likely to be stable or slowly changing in large populations (but could evolve more rapidly after population bottlenecks).

Creating epi–fitness landscapes equivalent to those in Figure 5.9 and Figure 5.10 is cumbersome and offers little insight beyond that provided by these existing landscapes and the curves shown in Figure 5.11. This is because the structure of the landscape is dependent on the precise statistics of the environmental fluctuations, which have been chosen here for illustrative purposes. We forgo a comprehensive study of the large space of environmental regimes and instead consider next a more clinically–focussed question: Given a specific bet–hedging population, how can we design effective drug therapies. Instead of categorising the full space of possible environmental regimes, we instead highlight those regimes which minimise population growth under clinical restrictions concerning dose frequency and potential toxicity.

5.6 Implications for Therapy

Eco–evolutionary thinking has driven the development of an alternative paradigm for treating diseases with drug resistance driven by Darwinian evolution. This paradigm, known as adaptive therapy, is predicated on understanding the evolutionary and ecological mechanisms driving drug resistance and pre–empting them with dosing schedules that avoid resistance. One such ecological principle is that of *competitive release*. Phenotypic alterations that result in drug resistance often carry a fitness cost in untreated environments. Thus, where resistant

phenotypes co–exist with sensitive phenotypes in a population, they do so in small numbers, or are driven to extinction. When drug treatment occurs, sensitive individuals are killed, reducing competition and allowing the resistant population to grow. Under these circumstances, the clinical maxim of prescribing the maximal tolerable dose for as long as possible may in fact drive resistance or recurrence by reducing competition [304, 85, 84].

Motivated by previous mathematical modelling [103], Enriquez-Navas et al. [85, 84] used cancer cell–line xenograft models (for both ER+ and triple negative breast cancer) to verify an adaptive strategy wherein drug dose was reduced in line with tumour shrinkage (identified by magnetic resonance imaging (MRI)). The dose de–escalation strategy was observed to lead to control of tumour size, where dose–skipping and maximum–dose strategies led to progression. MRI revealed reduced necrosis and increased vascularisation in the controlled tumours, suggesting that dose de–escalation may prevent the emergence of a micro–environmental niche that favours resistant phenotypes. The adaptive paradigm is now being investigated in a pilot clinical trial as a potential therapy for castration resistant metastatic prostate cancer (ClinicalTrial.gov: NCT02415621) at the Moffitt Cancer Center.

Here, we explore the potential of adaptive therapies for treating diseases with bet–hedging–driven resistance. This approach is interesting for two reasons, the first being that effective treatment strategies for bet–hedging diseases are at present lacking. The second is that, as cancer treatments progress towards adaptive therapies, we may be inadvertently treating bet–hedging populations with them. A recent study used a barcoded population of EGFR+ non–small–cell lung cancer (NSCLC) cells to identify the degree of pre–existing resistant clones in a cancer population [160]. This study found a high number of pre–existing resistant clones, indicating the potential for adaptive therapies to be effective because it may prevent competitive release. However, this study also found that one of the pre–existing mechanisms of resistance was characterised by a gene expression signature for a mesenchymal phenotype with no apparent genetic driver. This pattern indicates the possibility of phenotype switching driving drug–resistance, for example, through a stochastic epithelial-mesenchnymal transition (EMT) [321, 362, 170]. A number of studies have proposed adaptive–like therapies for NSCLC [49, 95], and so understanding the efficacy of these strategies in situations where resistance is driven by bet–hedging may be clinically important.

5.6.1 Maximum Dose Density Fails for Bet–Hedging Populations

We begin with a brief exploration of the effect of the maximum–dose–density (MDD) strategy in treating diseases for which recurrence is driven by bet–hedging. The MDD strategy is characterised by the prescription of the maximum tolerable dose for the longest possible time, and aims to induce maximum cell death and potentially cure disease. At present, this approach is standard clinical practice for many diseases, including cancers. The failure of therapy and ultimate recurrence of disease under MDD therapy has been explored previously, both experimentally and through simulations of population growth [85, 84]. We provide a brief overview of the failure of MDD therapies in bet–hedging populations. Figure 5.12 shows

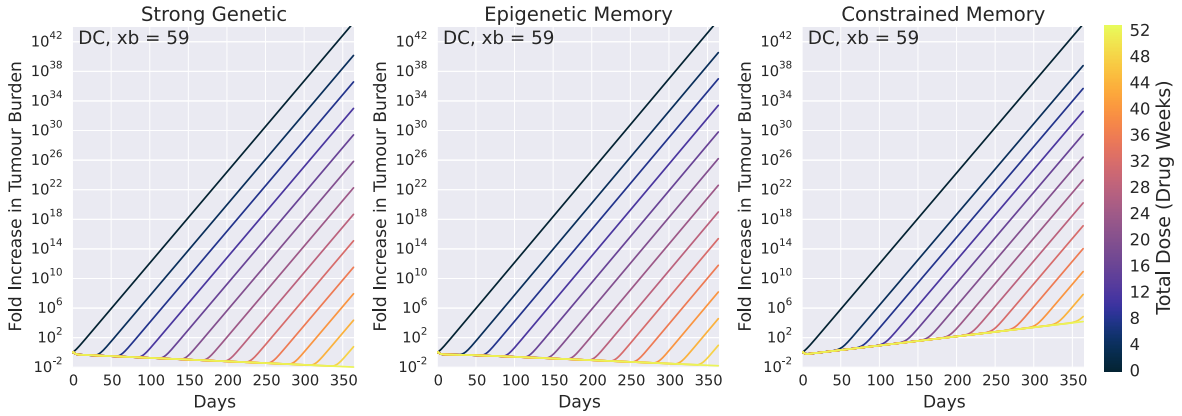

Fig. 5.12 **Maximum–dose–density (MDD) treatment fails for disease with bet–hedging–driven resistance.** Simulated growth curves are shown for a bet–hedging population exposed to an MDD treatment strategy with varying total dose (indicated by curve colour, see rightmost colourbar). The plots correspond to populations endowed with (left) strong–genetic, (centre) contrained–memory and (right) epigenetic–memory bet–hedging. The molecular switch is the DC switch and the genotype is $x_b = 59$. In all cases, a residual population persists and drives growth upon cessation of therapy. In the cases of strong–genetic and epigenetic–memory, the residual population is diminishing under therapy and could, in a clinical setting, experience stochastic extinction after a sufficiently long drug exposure. However, such an exposure may prove too toxic. Similar curves (not shown) are derived from simulations with alternative molecular switches.

growth curves for populations endowed with one of strong–genetic, constrained–memory or epigenetic–memory bet–hedging. The molecular switch (DC) and genotype (x_b) are chosen for illustrative purposes. For all modes of bet–hedging, the population either shrinks or grows slowly under drug exposure, but rebounds quickly after cessation of therapy. Where the population size is continually diminishing, a sufficiently long treatment regimen will ultimately lead to cure through stochastic extinction of the small disease population. However, owing to the toxicity of chemotherapeutic drugs, especially in the treatment of cancers, continual therapy is unethical or impossible, indicating that cessation of therapy and population outgrowth is likely inevitable. This phenomenon is robust to the precise mechanisms of bet–hedging and similar curves to those shown in Figure 5.12 (not shown) can be derived for each combination of molecular switch and inheritance mechanism.

5.6.2 Metronomic/Cyclic Therapy

The simplest alternative to a maximum–dose–density therapy, and one with a tunable degree of toxicity, is a cyclic (or metronomic) therapy. Under the cyclic paradigm, a dosing schedule is decomposed into intervals where a drug is taken and intervals where it is not. However, as noted by zur Wiesch et al. [368], erratic or irregular dosing schedules are unlikely to be correctly adhered to. It follows that to be clinically relevant, the cycling strategy must have a degree of regularity. The length of the intervals can be varied, for example lasting days

or weeks, and can impact the efficacy of the treatment. The dose of the drug taken for each interval remains fixed throughout the treatment, and the schedule is repeated until cessation for reasons such as disease progression or toxicity concerns.

Cyclic therapies have been explored as a potential means to preserve the efficacy of antibiotics where highly resistant clones emerge through *de novo* arising genetic mutations [30, 34, 255, 122, 368]. As discussed in Chapter 3, the principles underpinning these studies are equally applicable to designing cancer therapies. Here, we explore a restricted space of cycling strategies, considering only strategies in which a single cycle lasts 14 intervals. We explore the effects of both one day and one week intervals. To simulate toxicity restrictions, we assume the existence of a maximum tolerable 14 week dose, θ, measured in drug weeks. This approach places comparable restrictions on total dosing over 14 weeks, regardless of whether the interval length is one–day or one–week. We assume that cycles are repeated *ad infinitum* and define fitness as the the expected one–week log growth rate of the population, which can be calculated by the consideration of eigenvalues of the population update matrix defined in Section 5.2.2.

Define a cycling regimen, r, as a binary vector of length 14

$$r = (b_1, b_2, \ldots, b_{14}). \tag{5.35}$$

denoting the presence ($b_i = 1$) or absence ($b_i = 0$) of the drug for each interval. Denote the interval length, in days, by ι. Note that, regardless of the value of ι, the total 14-week dose is given by $\theta = \sum r$ drug weeks.

The one–interval population update matrices for the drug–free and drug–treated environments are then given by

$$\mathbf{P}_0^{\iota} = \mathbf{P}_{\text{hosp}}^{60\times60\times24\times\iota} \tag{5.36}$$

$$\mathbf{P}_1^{\iota} = \mathbf{P}_{\text{drug}}^{60\times60\times24\times\iota} \tag{5.37}$$

where the matrices $\mathbf{P}_{\text{hosp}}$ and $\mathbf{P}_{\text{drug}}$ are as described in Section 5.3. The log growth rate over a complete cycle is determined by the log of the spectral radius, ρ, of the matrix

$$\mathbf{P} = \mathbf{P}_{b_1}^{\iota}\mathbf{P}_{b_2}^{\iota}\ldots\mathbf{P}_{b_{14}}^{\iota}. \tag{5.38}$$

To ensure the cycling strategies with one–day and one–week intervals give comparable fitness measures, we normalise the log growth to a per–week average and thus define fitness by

$$F = \frac{1}{2\iota}\log\left(\rho(\mathbf{P})\right). \tag{5.39}$$

For each mode of bet–hedging, each molecular switch and each toxicity value ($0 \leq \theta \leq 14$), we evaluated the average population growth rate for each possible candidate cycling strategy. As such, a total of 2^{14} cycling strategies are evaluated for each interval length. Two

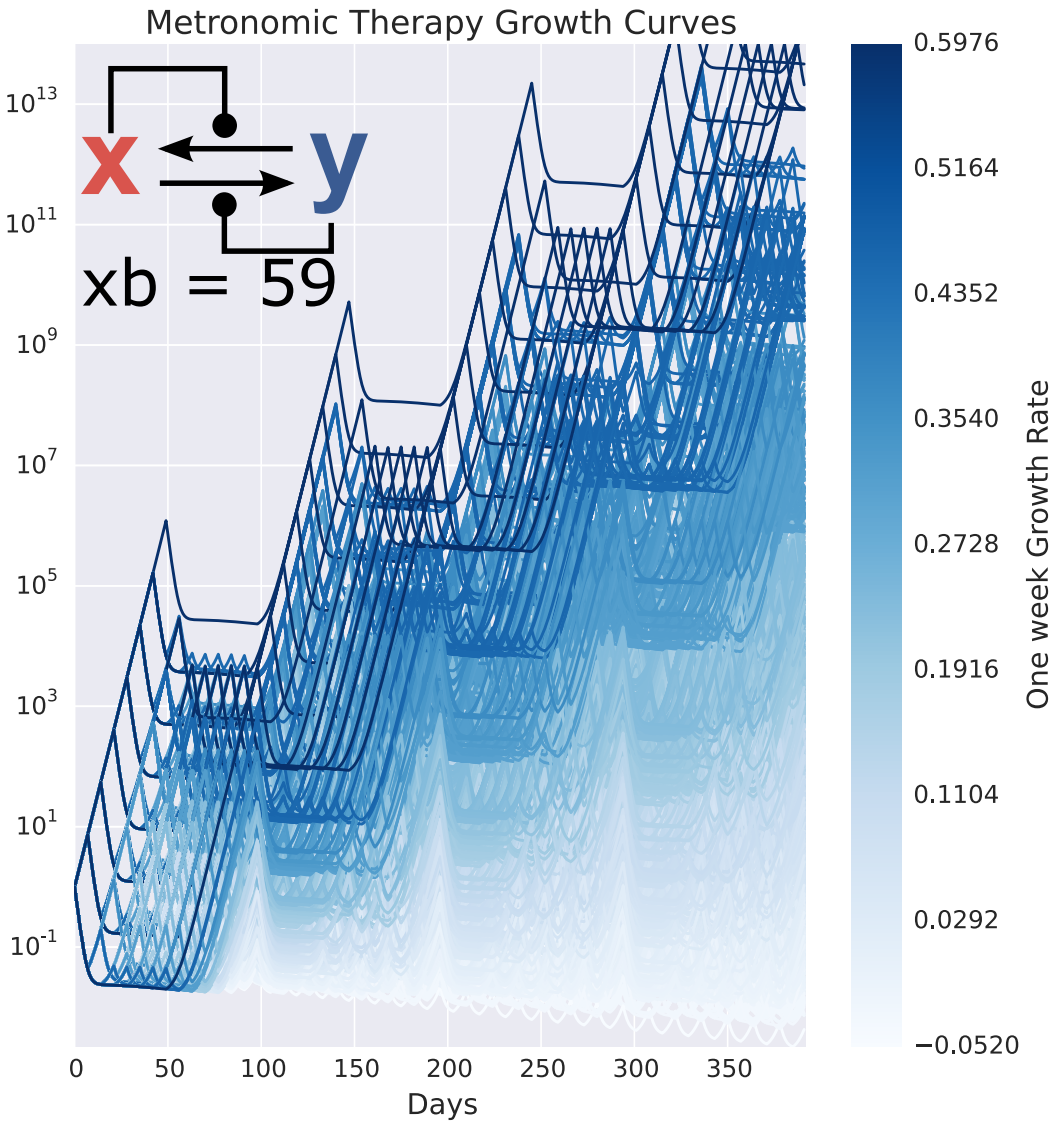

Fig. 5.13 **Growth curves for a bet–hedging population under a range of metronomic therapies.** Growth curves are shown for all 3432 14-week metronomic therapies with a total dose $\theta = 7$. The mechanism of bet–hedging is the epigenetic–memory model, the molecular switch is DC and $x_b = 59$, correspoding to memory–free proportion of phenotype A close to 99%. The memory decay parameter is $\sigma = 6.77 \times 10^{-7}$ s^{-1}.

representative genotypes, x_b (with $y_b = 60 - x_b$), were considered in our analysis. The first was chosen to correspond closely to a $p(x_b) = 0.99$ likelihood of stochastically producing individuals of phenotype A under the strong–genetic form of bet–hedging. These genotypes were chosen to mirror the phenomenon of persister–like cells and correspond to $x_b = 59$, $y_b = 1$ for the DC switch (giving $p = 0.983$), $x_b = 46$, $y_b = 14$ for the DCx switch (giving $p = 0.991$), $x_b = 59$, $y_b = 1$ for the DCy switch (giving $p = 0.909$) and , $x_b = 39$, $y_b = 21$ for the AM switch (giving $p = 0.991$) . Note that the values of p differ between the switches, resulting in differing results for the strong–genetic and constrained–memory models which are

dependent only on the value of p. Example growth curves for the 3432 metronomic therapies with $\theta = 7$ and one–week intervals are shown for the epigenetic memory model in Figure 5.13.

The second representative genotype x_b was chosen to correspond closely to a 50% split between the phenotypes in the absence of memory. As in Chapter 4, these genotypes were taken as $x_b = y_b = 30$ for the DC and AM switches (giving $p = 0.5$), $x_b = 7$, $y_b = 53$ for the DCx switch (giving $p = 0.488$) and $x_b = 53$, $y_b = 7$ for the DCy switch (giving $p = 0.511$).

We performed an exhaustive analysis to determine the average growth rate for all 2^{14} cycling strategies with $\iota = 7$. The results of this analysis are displayed in Figures 5.14 and 5.15 for the two representative genotypes. The scatter plots in these figures are arranged by rows corresponding to molecular switch and columns corresponding to inheritance mechanism. For each subplot, treatment strategies are plotted as a scatter plot with x–coordinate corresponding to θ and y–coordinate F, the weekly growth rate. The MDD strategies for each θ are marked with a dark hexagon. Plotted lines show the growth rates of the optimal strategy for each θ under different inheritance conditions (allowing a comparison along the rows of the figure).

The first conclusion we draw from this exhaustive survey is that the efficacy of metronomic therapies in subject to substantial variation, even where the total dose θ remains fixed. This phenomenon is clearly shown in the example growth curves for $\theta = 7$ in Figure 5.13. Critically, our exhaustive search finds that the MDD strategy is universally the least effective, regardless of the intracellular switch, the mode of bet–hedging or which of the (two example) genotypes is chosen. This indicates that the present approach of maximal–dose–density may be not just ineffective, but in fact the worst possible approach for a fixed drug dose. For example, the growth rates for MDD of high total dose (say $\theta \geq 12$ say) can often be matched by strategies with a substantially lower total dose.

The MDD strategy in our model is assumed to be repeated *ad infinitum*, As such, it is still a metronomic therapy, albeit with 'patchier' environmental switching. The other strategies in our analysis can be considered to be less patchy, or more *fractionated*, and achieve greater population control. From this observation we hypothesise that, for a fixed total dose, a more optimal metronomic strategy might be found by further fractionating the dose. We reason that a fractionated dose may be effective because giving a drug dose when there are very few proliferative cells causes only a slight decrease in population size. For example, see Figure 5.12, where a constant drug dose induces a slow decline in population size. That same dose may be more effective when the proportion of sensitive cells is higher.

To test this hypothesis we performed a second exhaustive survey of treatment strategies, but with one–day treatment intervals ($\iota = 1$). Figures 5.16 and 5.17 show the results of this analysis. For the fractionated dose, we see that the fitness distribution for the treatment strategies for fixed θ collapses so that the the MDD treatment is close to the optimal for each molecular switch and memory type. This gives a clear demonstration of the efficacy of dose fractionation in the treatment of disease with bet–hedging driven resistance.

To further demonstrate the efficacy of fractionation, and to clearly highlight the importance of strategy optimisation where fractionation is not possible, we produced Figure 5.18 and 5.19

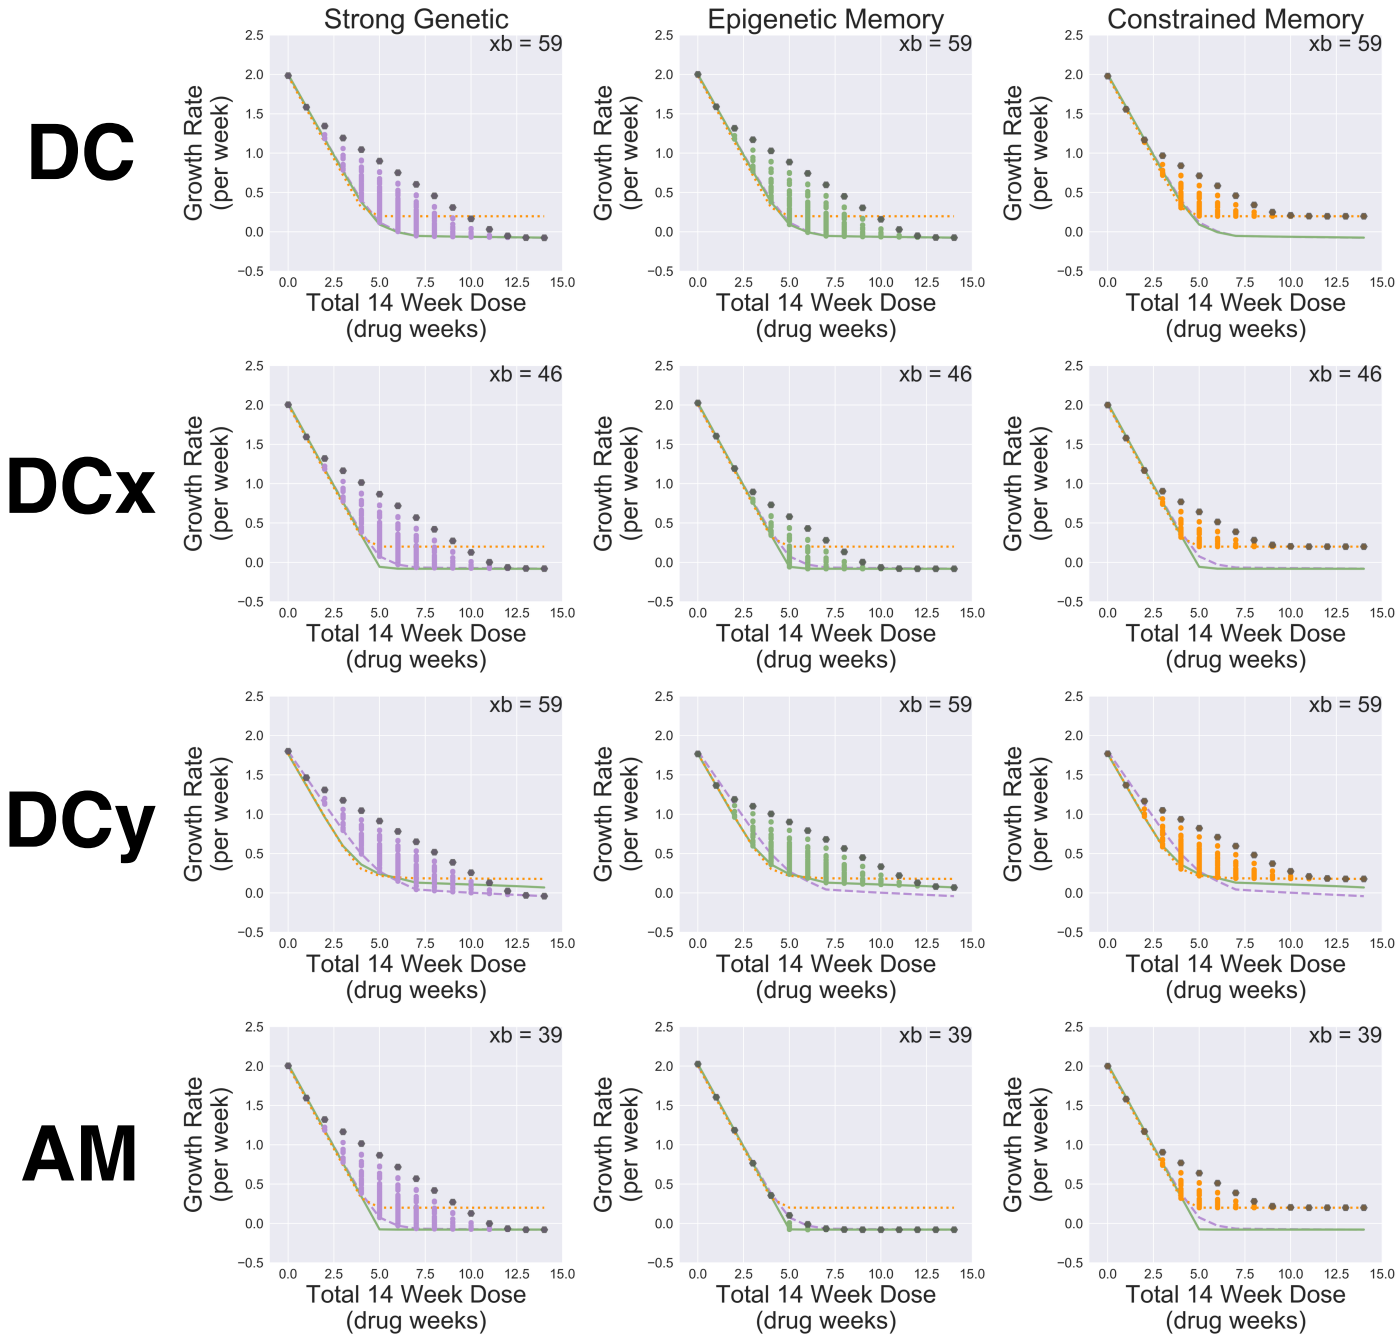

Fig. 5.14 **Efficacy of toxicity–restricted cycling strategies on a population with a small proportion of persisters.** An exhaustive analysis of all 14–week metronomic therapies, in which each week is designated as drug–on or drug–free, was performed. The analysis was performed for each molecular switch (rows), each memory mechanism (columns) and for a genetic component (x_b) corresponding closely to a 99% proportion of each phenotype. In each subplot, a point represents a strategy and is placed on the x-axis according to the total 14–week dose and y-axis according to tumour growth. Maximum–dose–density strategies (wherein drug is prescribed as a block until the toxicity limit is reached) are denoted by dark hexagons. In each subplot, the lines denote the dose–restricted minimum growth rate attainable for strong–genetic (purple, solid), epigenetic–memory (dashed, green) and constrained–memory (orange, dotted) modes of bet–hedging memory for the specified switch (row).

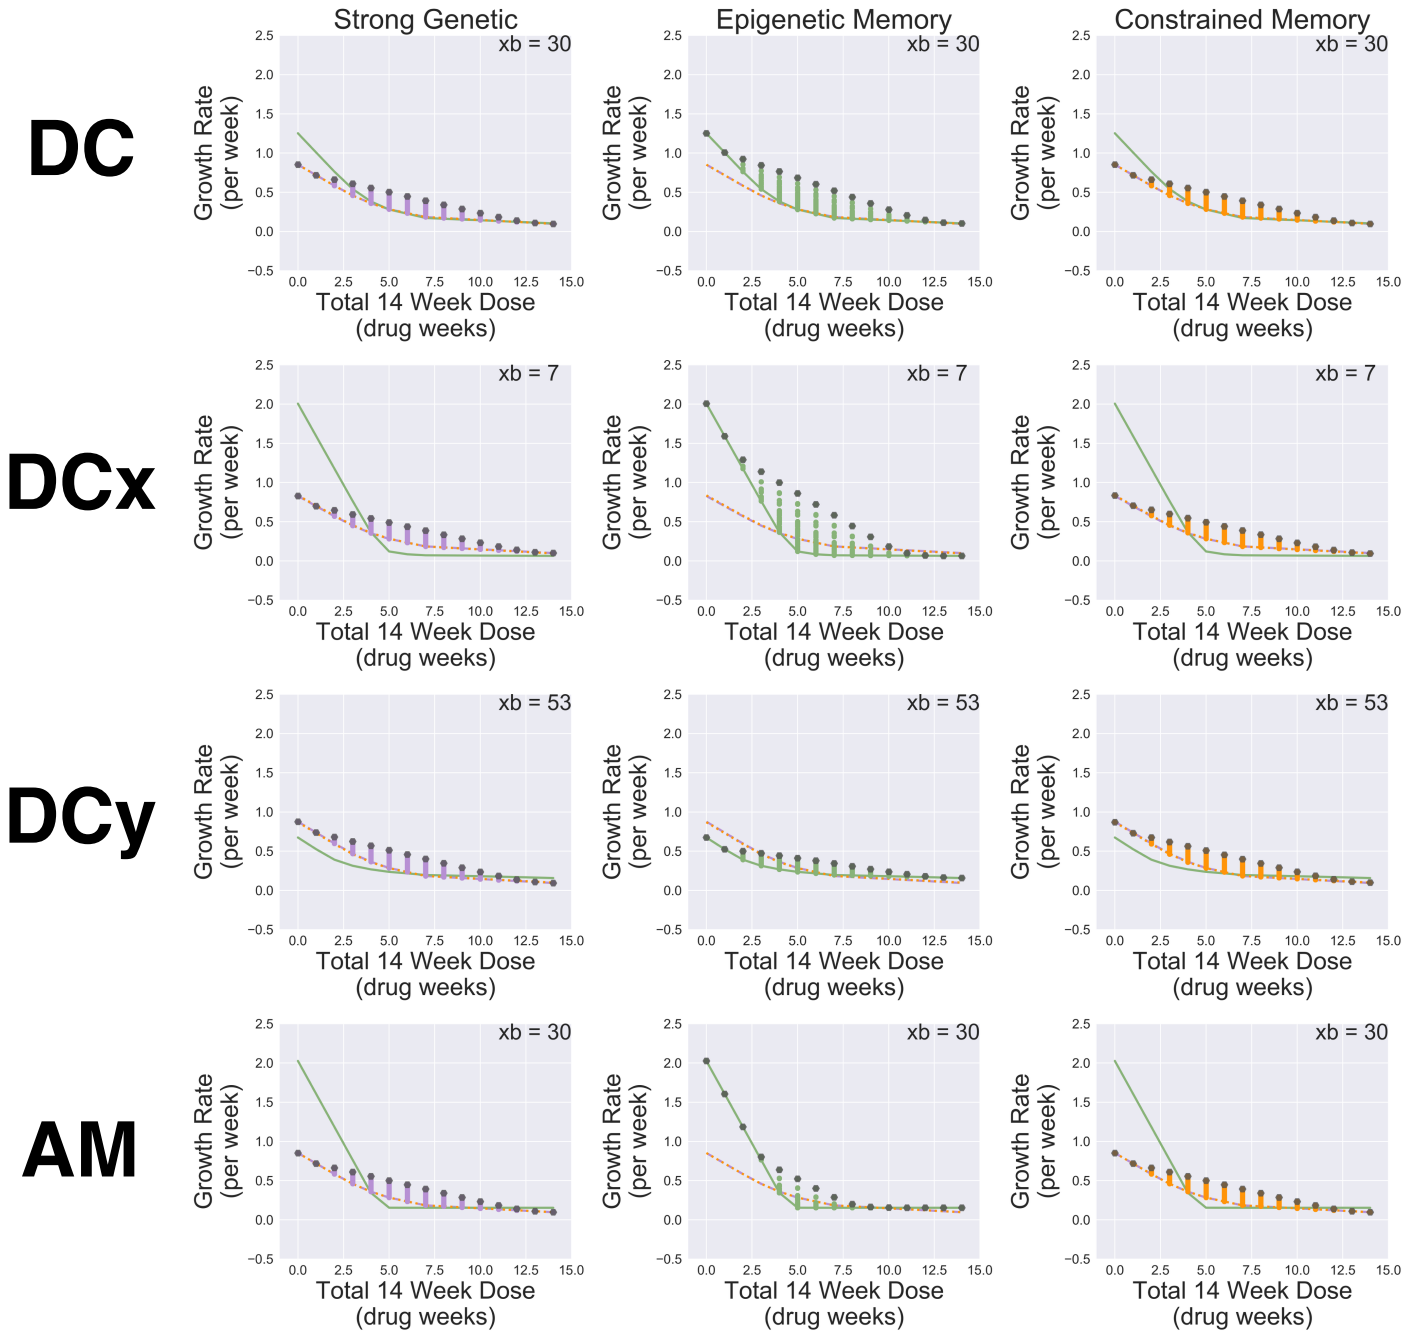

Fig. 5.15 **Efficacy of toxicity–restricted cycling strategies on a population with a large proportion of persisters.** An exhuastive analysis of all 14–week metronomic therapies, in which each week is designated as drug–on or drug–free, was performed. The analysis was performed for each molecular switch (rows), each memory mechanism (columns) and for a genetic component (x_b) corresponding closely to a 50% proportion of each phenotype. In each subplot, a point represents a strategy and is placed on the x-axis according to the total 14–week dose and y-axis according to tumour growth. Maximum–dose–density strategies (wherein drug is prescribed as a block until the toxicity limit is reached) are denoted by dark hexagons. In each subplot, the lines denote the dose–restricted minimum growth rate attainable for strong–genetic (purple, solid), epigenetic–memory (dashed, green) and constrained–memory (orange, dotted) modes of bet–hedging memory for the specified switch (row).

for the two example genotypes. These figures show, for each molecular switch (rows), ι value (columns) and memory mechanism (line colour/style), how treatment optimisation can reduce toxicity. For each sub–plot, the restricted dose θ for the MDD strategy is plotted against the minimum dose that is capable of achieving population growth control at least as effectively. Thus, deviations below the line $y = x$ illustrate how treatment optimisation can reduce toxicity whilst maintaining efficacy. The shift of the curves towards $y = x$ for the $\iota = 1$ fractionated dose indicates that optimal therapeutic strategies offer smaller improvements over the MDD strategy when the dose is fractionated. Thus, strategy selection is less important in these cases.

We note that as memory mechanism is changed, the optimal dose–restricted strategy may change too. To conclude our exploration of the space of treatment strategies, we compared the efficacy of the memory–specific optimal strategies when the mode of bet–hedging memory is changed. Remarkably, we find that for a fixed genotype, molecular switch and dose restriction θ, a treatment strategy that is optimal for one memory mechanism results in an expected weekly log growth rate within 0.1% of the true optimum strategy when the memory mechanism is changed. This result indicates that, whilst the memory mechanism is a strong determinant of the rate of population growth under the optimal therapy, the structure of the optimal therapy is memory–agnostic. This represents a positive result for the design of adaptive or metronomic therapies, as the nature of epigenetic memory and inheritance is likely to be difficult to discern experimentally.

5.7 Memory Decay and Optimal Therapies

Where epigenetic memory is present in the population, the degree of memory (as determined by the decay rate σ), may be subject to evolution. For example, the over– or under–expression of enzymes that decay the molecules x and y could serve as mechanism to alter σ. In this section, we explore effects of variations to the memory parameter σ on the efficacy of metronomic therapies and ask how the dynamics of memory are expected to evolve.

For each of the molecular switches, and with genotype parameters x_b chosen as before to correspond to a memory–free proportion of phenotype A of 0.99 or 0.5, we calculated the growth rates of an epigenetic memory population with $\sigma \in [10^{-11}\ \mathrm{s}^{-1}, 10^{-3}\ \mathrm{s}^{-1}]$ under the optimal metronomic therapy derived for $\sigma = 6.77 \times 10^{-7}\ \mathrm{s}^{-1}$ as well as the MDD therapy. These calculations were performed for three examples values for the restricted 14–week dose, $\theta \in \{3, 7, 11\}$. Figure 5.20 shows the relationship between σ and growth rate for epigenetic memory populations under the dose–restricted optimal therapy for $\sigma = 6.77 \times 10^{-7}\ \mathrm{s}^{-1}$(red line) and the dose–restricted MDD (blue line) and with x_b corresponding to a memory–free population of approximately 99% phenotype A individuals. Figure 5.21 shows the same results but with x_b corresponding to a memory–free population of approximately 50% phenotype A individuals. In all cases, variations in σ result in large changes in the growth rate.

Critically, we see that the fitness effect of altering σ is strongly switch dependent, to the extent that under different switches we may expect σ to evolve in entirely different directions.

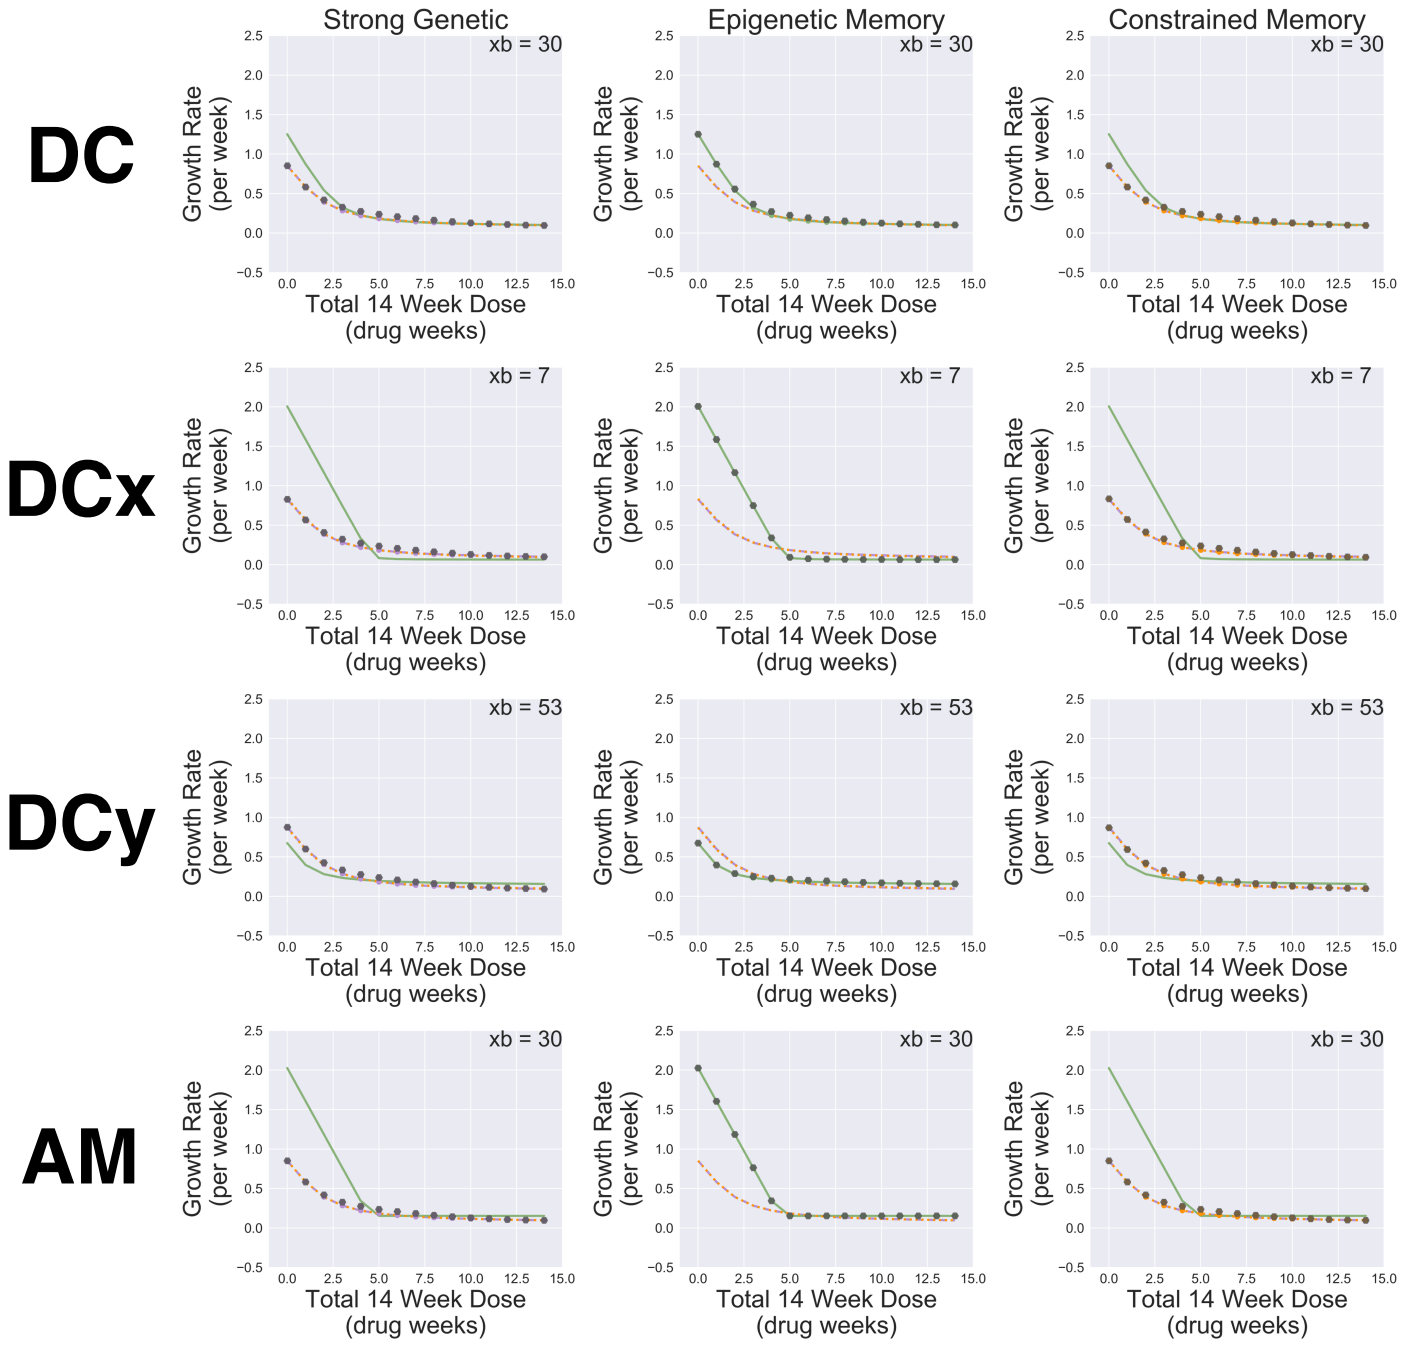

Fig. 5.16 **Efficacy of toxicity–restricted cycling strategies on a population with a small proportion of persisters.** An exhuastive analysis of all 2–week metronomic therapies, in which each day is designated as drug–on or drug–free, was performed. The analysis was performed for each molecular switch (rows), each memory mechanism (columns) and for a genetic component (x_b) corresponding closely to a 99% proportion of each phenotype. In each subplot, a point represents a strategy and is placed on the x-axis according to the total 2–week dose and y-axis according to tumour growth rate. Maximum–dose–density strategies (wherein drug is prescribed as a block until the toxicity limit is reached) are denoted by dark hexagons. In each subplot, the lines denote the dose–restricted minimum growth rate attainable for strong–genetic (purple, solid), epigenetic–memory (dashed, green) and constrained–memory (orange, dotted) modes of bet–hedging memory for the specified switch (row).

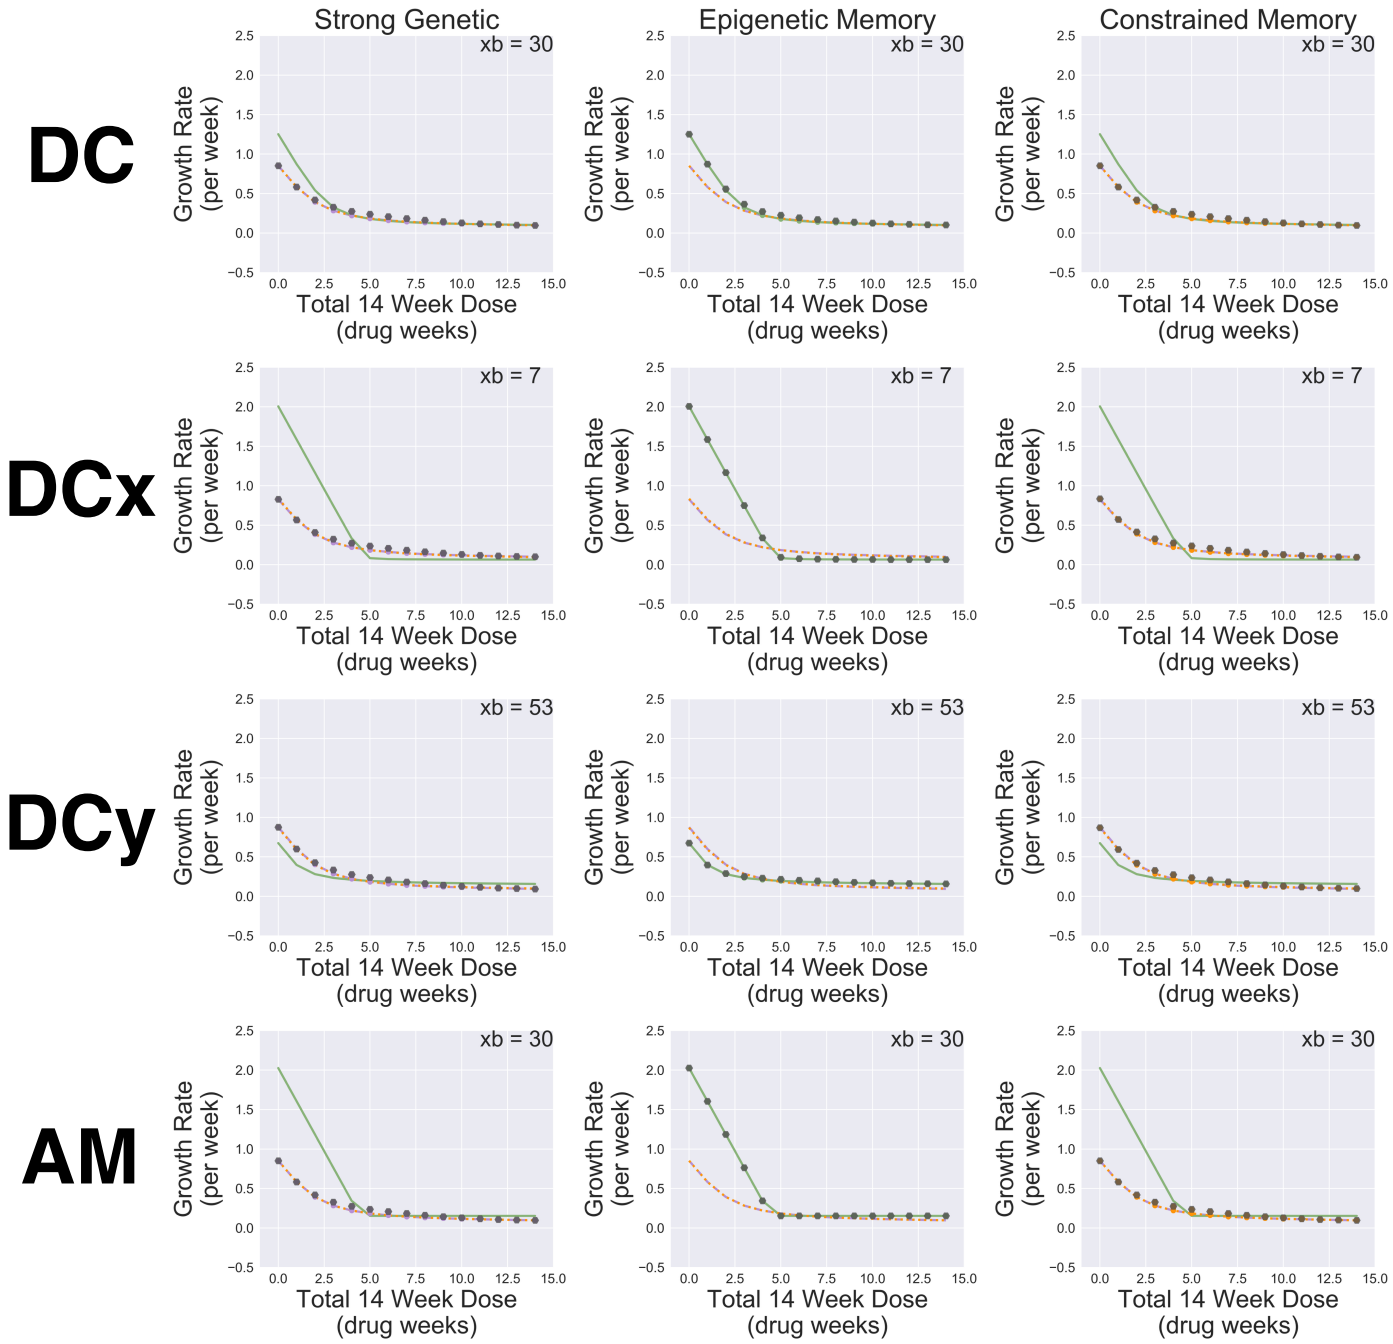

Fig. 5.17 **Efficacy of toxicity–restricted cycling strategies on a population with a large proportion of persisters.** An exhuastive analysis of all 2–week metronomic therapies, in which each day is designated as drug–on or drug–free, was performed. The analysis was performed for each molecular switch (rows), each memory mechanism (columns) and for a genetic component (x_b) corresponding closely to a 50% proportion of each phenotype. In each subplot, a point represents a strategy and is placed on the x-axis according to the total 2–week dose and y-axis according to tumour growth. Maximum–dose–density strategies (wherein drug is prescribed as a block until the toxicity limit is reached) are denoted by dark hexagons. In each subplot, the lines denote the dose–restricted minimum growth rate attainable for strong–genetic (purple, solid), epigenetic–memory (dashed, green) and constrained–memory (orange, dotted) modes of bet–hedging memory for the specified switch (row).

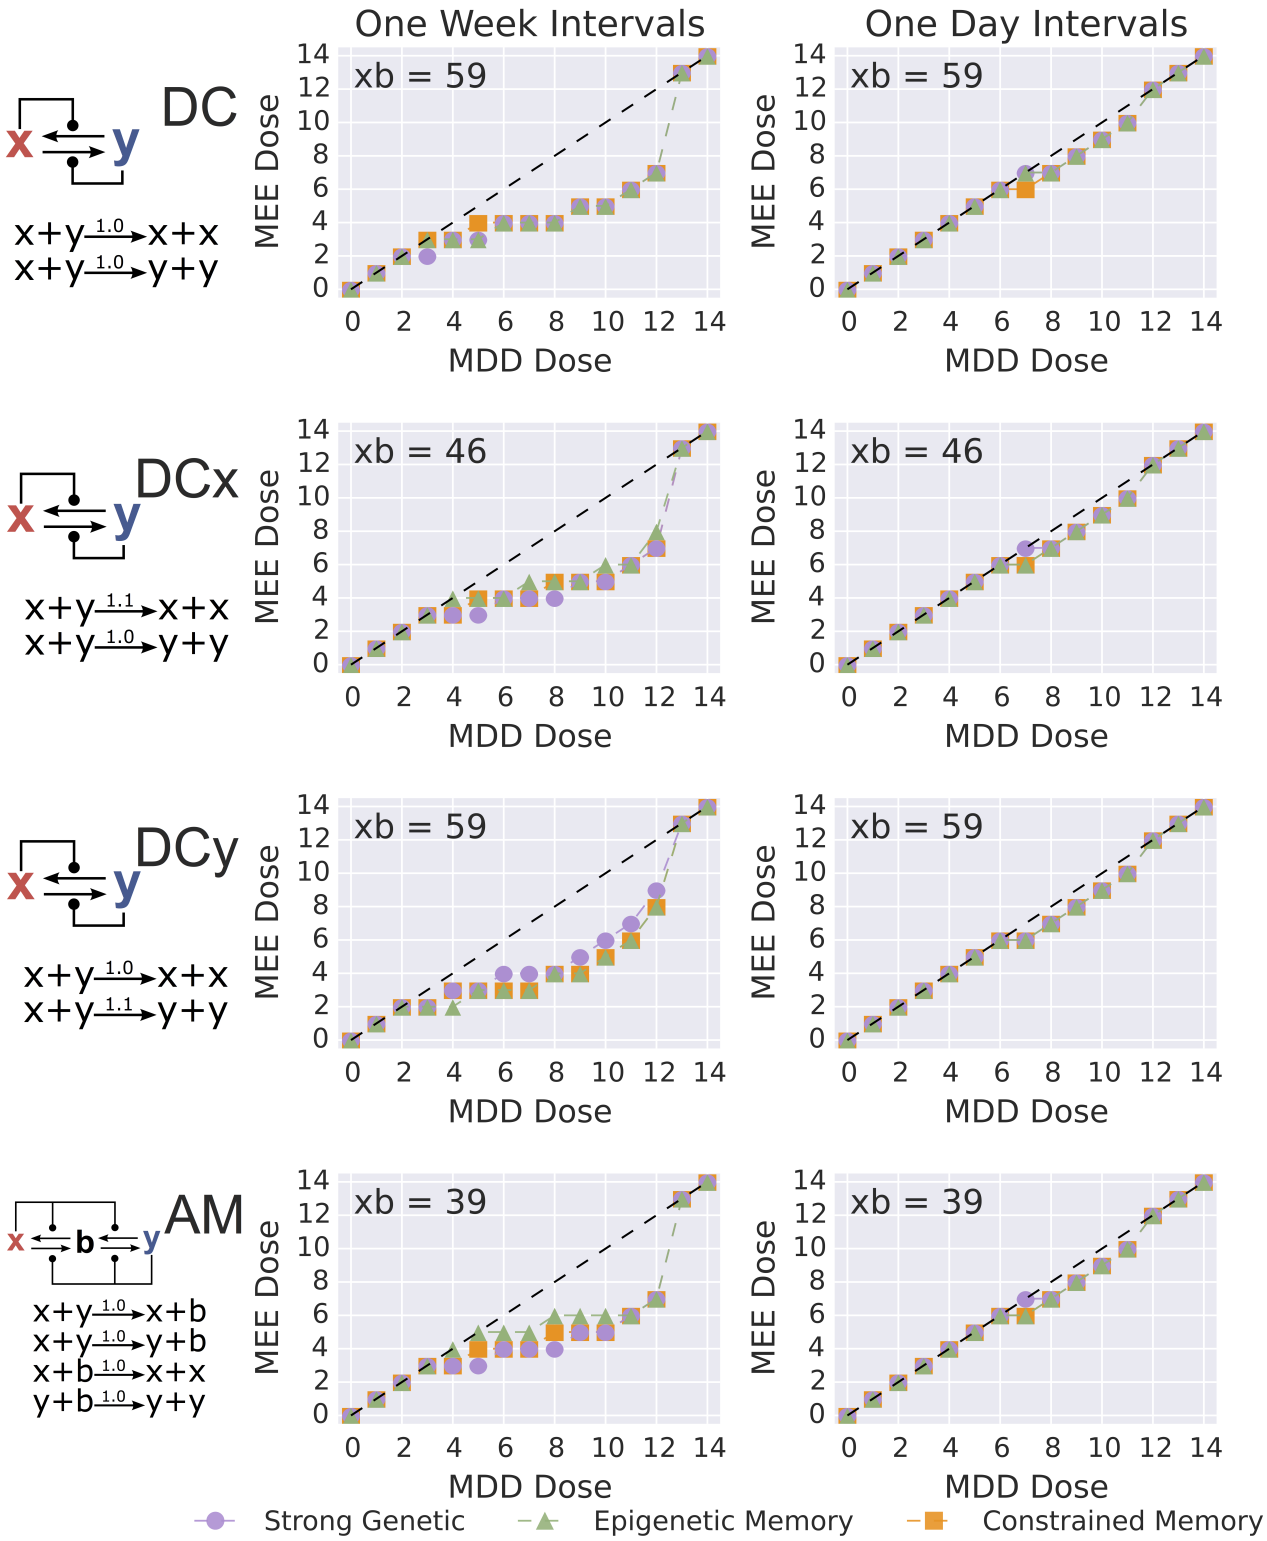

Fig. 5.18 **The maximum equivalent dose (MEE) for each maximum–dose–density (MDD) therapy.** For each molecular switch, the restricted dose values (θ) are plotted against the minimum dose for which the dose–restricted–optimal therapy results in a smaller growth rate than the MDD therapy. Deviations below the line $y = x$ indicate the degree to which toxicity can be reduced whilst maintaining the equivalent therapeutic benefit. The columns show the curves for drug intervals of one week and one day. The genotypes (x_b) are chosen to correspond closely to a memory–free likelihood of 0.99 for phenotype A.

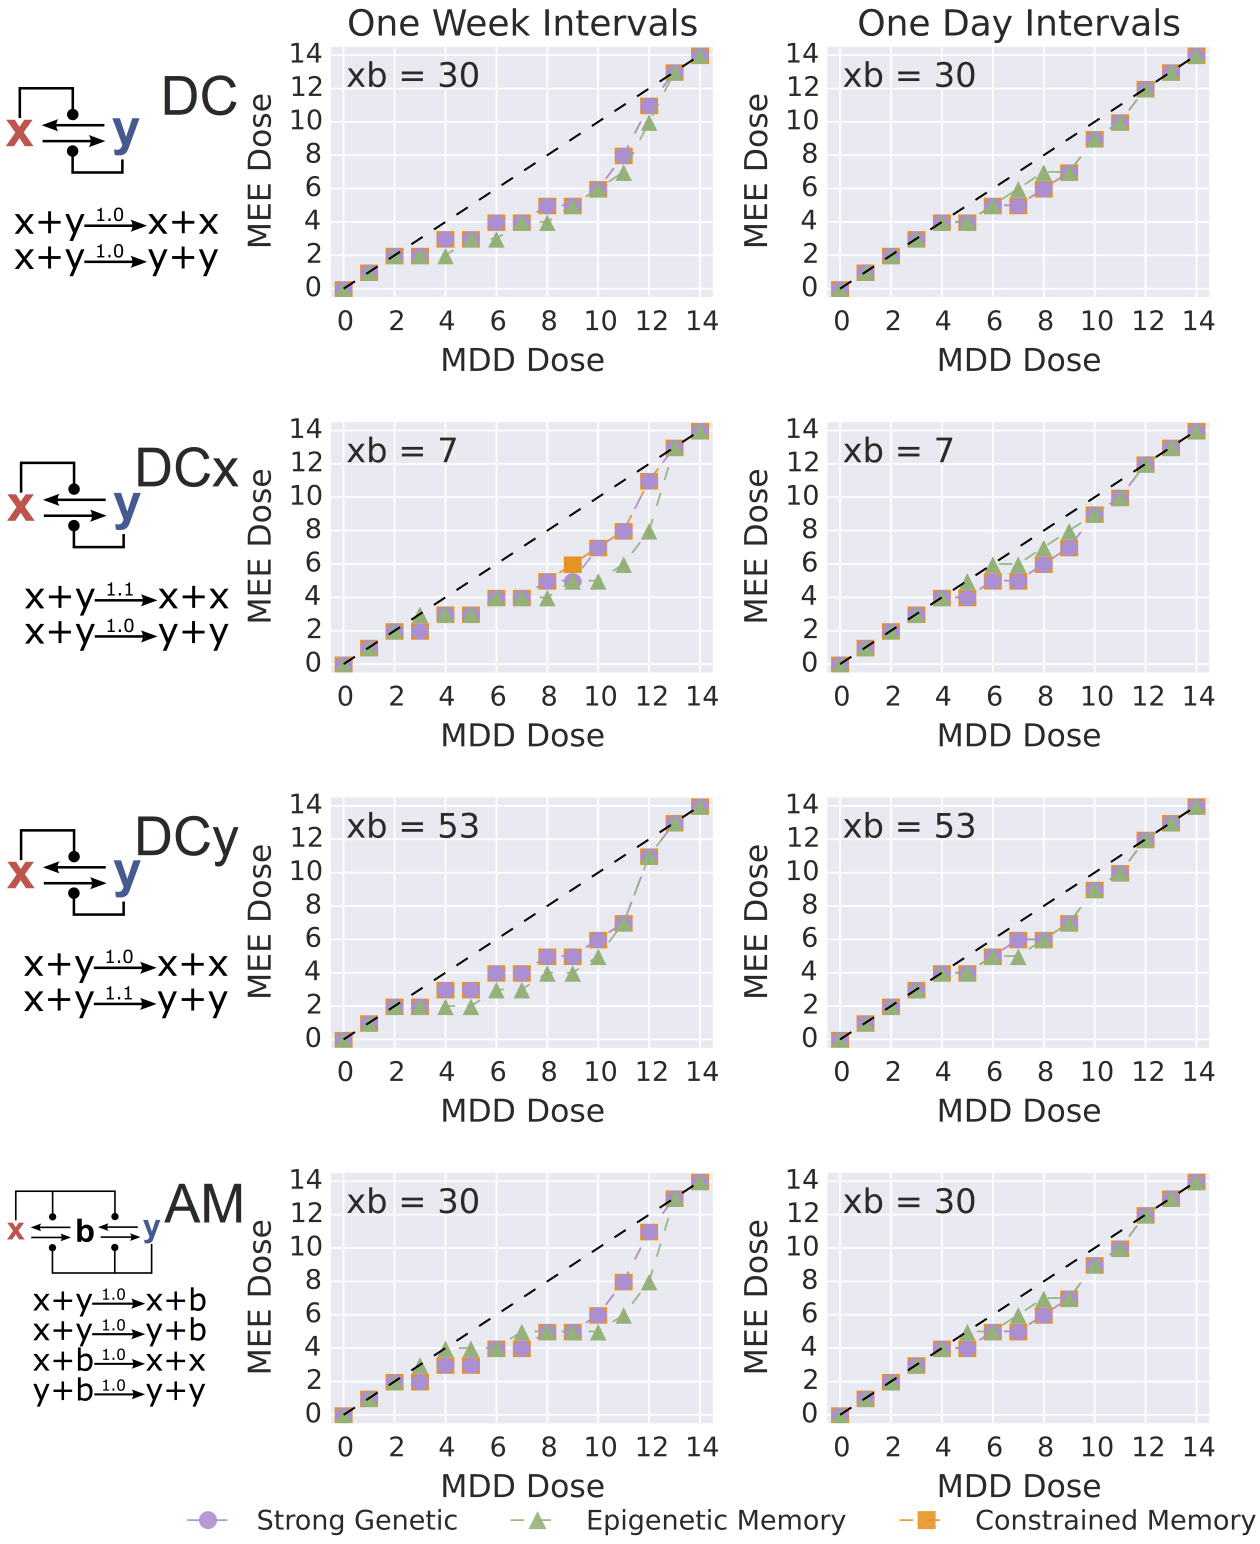

Fig. 5.19 **The maximum equivalent dose (MEE) for each maximum–dose–density (MDD) therapy.** For each molecular switch, the restricted dose values (θ) are plotted against the minimum dose for which the dose–restricted–optimal therapy results in a smaller growth rate than the MDD therapy. Deviations below the line $y = x$ indicate the degree to which toxicity can be reduced whilst maintaining the equivalent therapeutic benefit. The columns show the curves for drug intervals of one week and one day. The genotypes (x_b) are chosen to correspond closely to a memory–free likelihood of 0.5 for phenotype A.

Figure 5.21 shows an example of this phenomenon wherein, under the $\theta = 3$ MDD therapy, selection acts to increase σ under the DCy switch and to minimise σ under the DCx switch. Further, amongst the two figures we see examples of global optima of fitness (Figures 5.21, AM switch, $\theta = 7$ optimal therapy) and also minima of fitness from which evolutionary divergence of the σ value could occur (Figures 5.20, AM switch, $\theta = 7$ optimal therapy). Assuming mutations alter σ by only a small increase or decrease, these minimum points could render evolutionary trajectories inaccessible as they represent fitness valleys.

These results again highlight a non–additive relationship in the higher–order mechanisms governing the determination of phenotypes. In this case, it is the molecular switch and time–scale of intracellular decay that are observed to interact in unintuitive ways. We conclude that, where bet–hedging is investigated experimentally, and a memory component is verified to exist, that predictions about how this memory component will evolve cannot be made without further understanding of the underlying biology. We note that the non–linearity and switch–dependence of the relationship between σ and fitness is unintuitive, and it is only by considering evolution from the abstracted perspective of the base functional components that this relationship is clearly revealed. Again, we see the power of the abstracted, rather than reductionist, approach to biology.

The model GP–mappings we explore here assume that the resolution of the molecular switch to a stable configuration occurs sufficiently quickly that we can assume that it is instantaneous. With respect to memory, this assumption means that no intracellular molecules decay during the resolution of the switch. Where σ is large, for example $\sigma > 10^{-2}$ (the largest value used in our analysis), this assumption may break down. To explore the implications of weakening our assumptions would require the derivation of a more appropriate model. Such a derivation is beyond the scope of the work presented here. However, this observation does highlight a potential problem of studying biology through the algorithmic lens; that we must ensure that only sound assumptions are made in the generation of our models, and that these assumptions are not inadvertently violated in the analysis.

Finally, we note that the model presented here makes the simplifying assumption that the molecules x and y decay at an equal rate (σ). Where the decay rates of these two factors are permitted to vary independently, it seems likely that a more rugged evolutionary landscape of mutations to σ_x, σ_y and the molecular switch will emerge. A characterisation of this model represents a possible extension of the work presented here.

5.7.1 Dose Skipping Adaptive Therapy

The optimal therapeutic strategies identified by our exhaustive search are computationally expensive to derive. Further, they are dependent on the precise parameters underpinning the growth dynamics. By contrast, adaptive therapies of the type studied by Enriquez-Navas et al. [85, 84] are dependent only on the observed growth dynamics for the tumour. As such, these heuristic strategies are easy to determine as they do not require forward simulation of the growth dynamics, instead relying on estimates of the population size at fixed time intervals. We

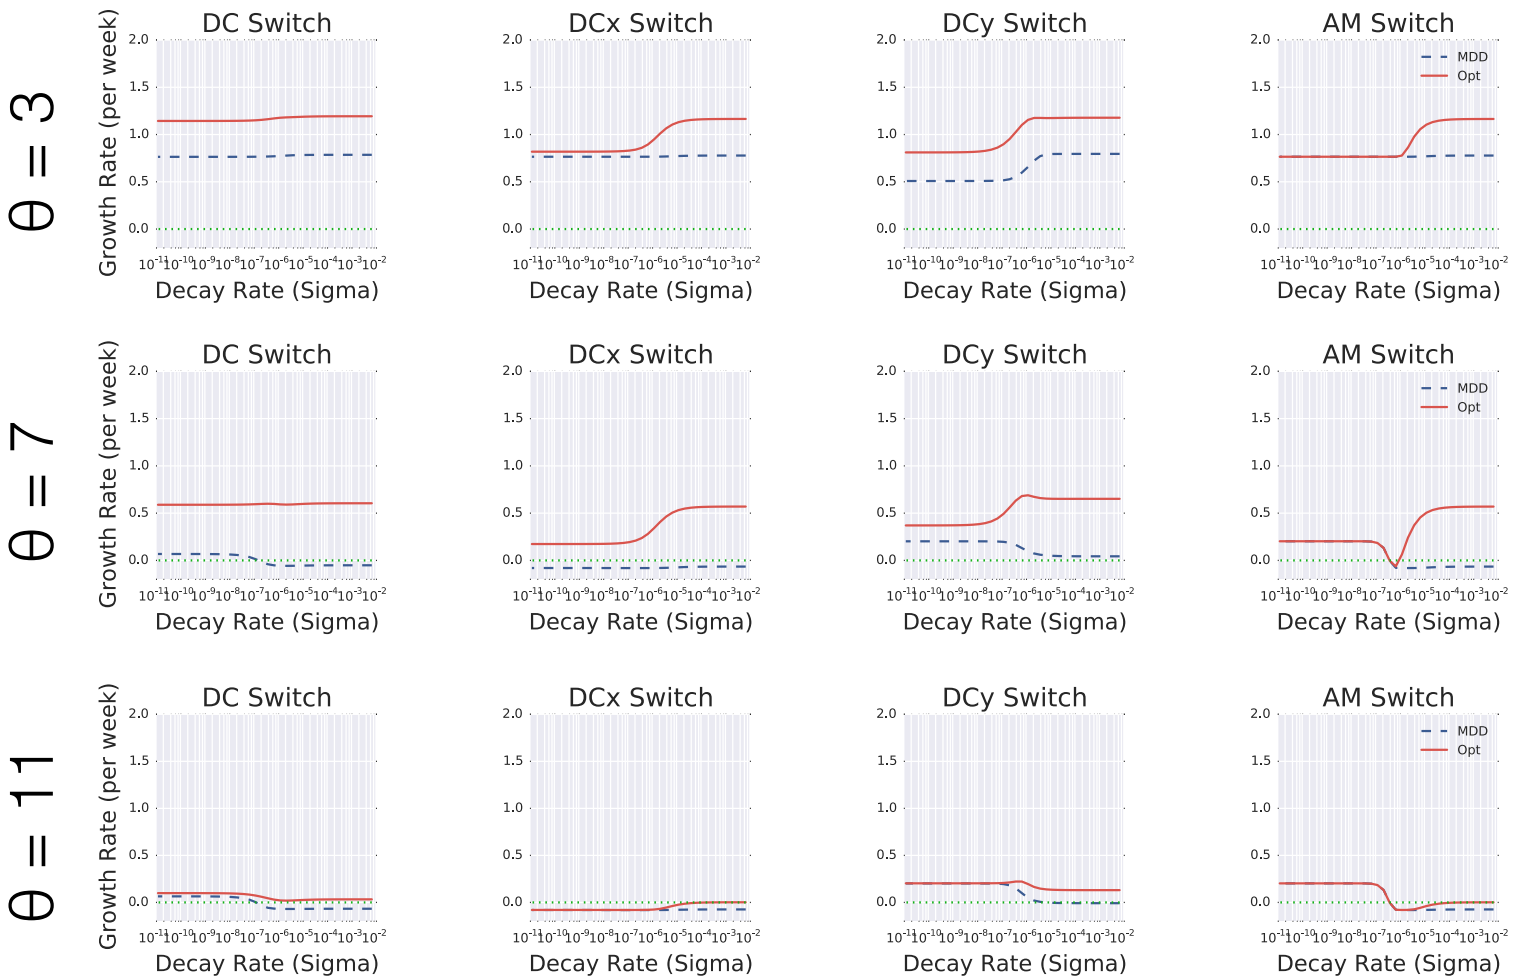

Fig. 5.20 **The fitness effects of varying σ are dependent on the molecular switch.** Curves depict how the population growth rate changes for the MDD (blue, dashed) and dose–restricted–optimal strategy (red, solid) as the memory–decay parameter σ is varied. Curves are depicted for each molecular switch (columns) and three example doses (rows). The dose–restricted–optimal strategy is computed for $\sigma = 6.77 \times 10^{-7}$ s^{-1}. The genotypes, x_b, are chosen to correspond to an approximate likelihood 0.99 of producing offspring of phenotype A under strong–genetic bet–hedging. The green, dotted line indicates the boundary at which growth rate is zero.

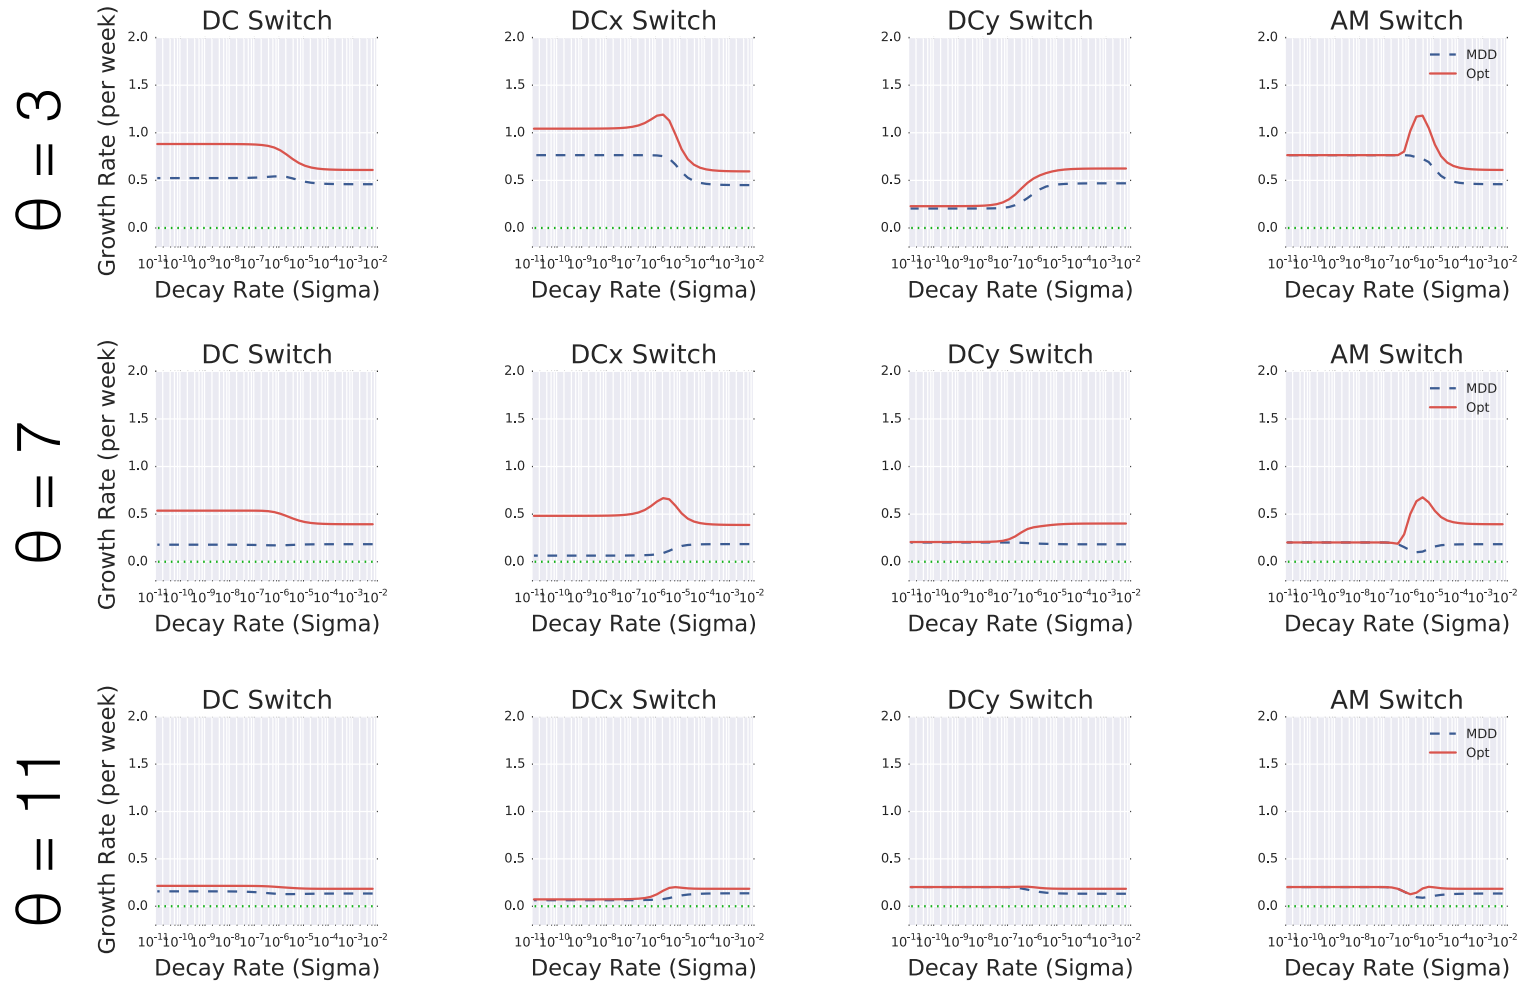

Fig. 5.21 **The fitness effects of varying σ are dependent on the molecular switch.** Curves depict how the population growth rate changes for the MDD (blue, dashed) and dose–restricted–optimal strategy (red, solid) as the memory–decay parameter σ is varied. Curves are depicted for each molecular switch (columns) and three example doses (rows). The dose–restricted–optimal strategy is computed for $\sigma = 6.77 \times 10^{-7}$ s^{-1}. The genotypes, x_b, are chosen to correspond to an approximate likelihood 0.5 of producing offspring of phenotype A under strong–genetic bet–hedging. The green, dotted line indicates the boundary at which growth rate is zero.

conclude our survey of the relationship between molecular switches, memory and therapeutic outcome with a comparison of optimal strategies with a dose–skipping adaptive therapy.

We implemented an adaptive strategy, mirroring Enriquez-Navas et al. [84], wherein a drug dose is given for a specific interval (of length ι, as before) only if the total population size N_T at interval T satisfies $N_T \geq 1.2N_{T-2}$. The initial two time points N_0 and N_1 are designated as drug dose intervals to mirror a clinically reasonable therapy wherein drug is administered as early as possible. To mirror our previous simulations of metronomic therapy, we imposed a dose–restriction on the adaptive therapy. For a dose–restriction θ, a drug–dose interval was prohibited if the sum of drug–dose weeks over the previous 13 intervals exceed θ. In this way, adaptive therapies can be short–sighted in comparison to an optimal strategy that saves a drug–dose for when it is most effective.

The results of these simulations of adaptive therapy are shown in Figures 5.22 and 5.23, for three example dose restrictions $\theta = 3, 7, 11$ (rows) and each of the molecular switches (columns). Here, the MDD (green), adaptive therapy (blue) and optimal therapy (red) growth curves are shown. We see that adaptive therapies are effective where the dose–restriction θ is high. For example, for $\theta \leq 7$, the adaptive therapy matches the optimal therapy in tumour control. For larger θ values, the adaptive therapy performs worse than the optimal in all cases. Most strikingly, adaptive therapy fails to control tumour for the DCy switch with $\theta = 11$ and $\sim$50% memory–free likelihood of individuals taking phenotype A, even though the optimal therapy can induce control. Adaptive therapy is less effective because it does not use all of the drug doses available under a high toxicity limit. However, despite the reduced dose, the adaptive therapy is often able to induce negative growth where the optimal therapy does so, albeit with a slower rate of shrinkage. It is only when the dose is heavily restricted, for example $\theta = 3$, that the adaptive therapy performs considerably worse than the optimal. The magnitude of the difference between the optimal and adaptive therapies in this case is switch dependent, again demonstrating the importance of considering the GP–mapping.

Similar results (not shown) can be derived for the alternative forms of bet–hedging. From these results we conclude that adaptive therapy may prove an effective heuristic strategy for controlling disease with bet–hedging–driven resistance whilst reducing toxicity. However, care must be taken as when toxicity is high, and as the available dose becomes severely restricted, the adaptive therapy may become sub–optimal. In this case, the adaptive strategy is too ‘greedy’ in assigning the available doses, resulting in long periods of drug–free growth at the end of the cycle period. This ultimately drives faster average growth in the population. Conversely, care must also be taken when toxicity is low, as the adaptive therapy may prove too conservative. Under certain conditions (of genotype, molecular switch memory mechanism), this may lead to tumour outgrowth where control can be achieved with a higher dose.

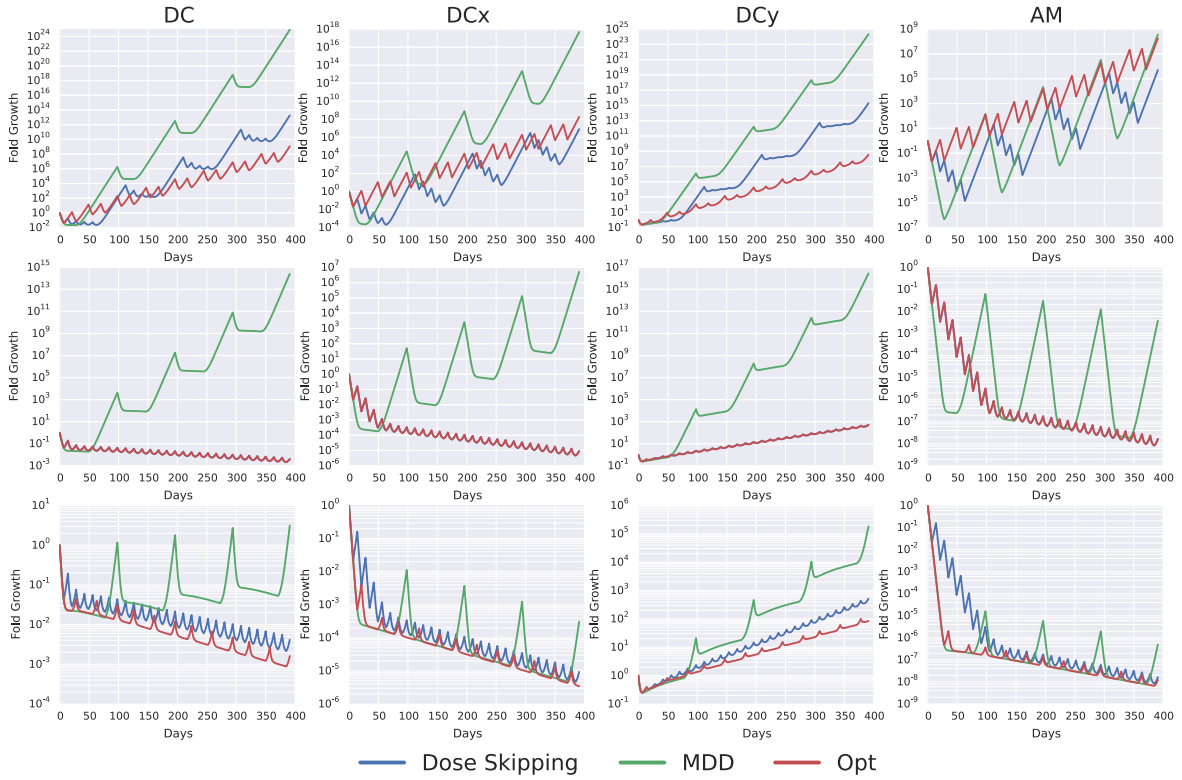

Fig. 5.22 **Comparions of MDD, adaptive therapy and optimal strategies for restricted doses.** One year long growth curves are shown for dose–restricted therapies. Maximum–dose–density (green), adaptive dose–skipping therapy (blue) and the optimal strategy (red) growth curves are shown on the same axes for each switch (columns) and doses $\theta = 3$ (first row), $\theta = 7$ (second row) and $\theta = 11$ (third row). In this example the population is endowed with the epigenetic memory mechanism ($\sigma = 6.77 \times 10^{-7}$ s^{-1}) and genotypes (x_b) corresponding closely to a memory–free switching probability of 0.99.

5.8 Implementation Details and Computational Considerations

The implementation of the models, derivation of the results and the production of all of the figures was performed using the Python programming language. A full annotated implementation of all models, the derivation of all results, and the code used to produce the figures is available from Github (see Appendix 3). Appendix 3 contains a Jupyter notebook that provides an overview of this code.

5.8.1 Performance

As outlined in Section 5.2.2, an explicit stochastic simulation was implemented but proved too computationally expensive to derive the required results (Figure 5.5). Instead, a discrete time matrix model was derived. The construction of this matrix, derivation of an exponential growth rate, and the determination of growth curves in this model execute quickly, on the order

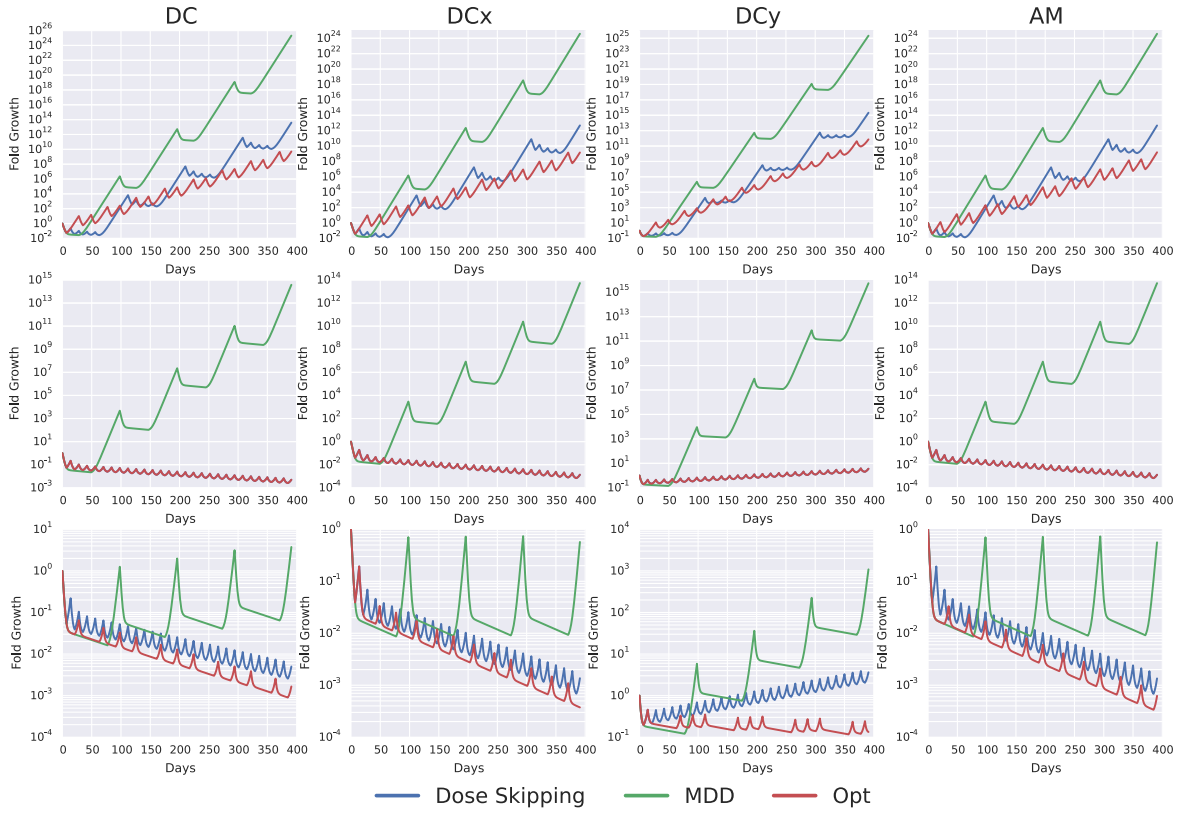

Fig. 5.23 **Comparions of MDD, adaptive therapy and optimal strategies for restricted doses.** One year long growth curves are shown for dose–restricted therapies. Maximum–dose–density (green), adaptive dose–skipping therapy (blue) and the optimal strategy (red) growth curves are shown on the same axes for each switch (columns) and doses $\theta = 3$ (first row), $\theta = 7$ (second row) and $\theta = 11$ (third row). In this example the population is endowed with the epigenetic memory mechanism ($\sigma = 6.77 \times 10^{-7}$ s^{-1}) and genotypes (x_b) corresponding closely to a memory–free switching probability of 0.5.

of seconds on a personal computer. However, the results presented here rely on exhaustive searches of genotype/switch/memory mechanisms (in the case of Section 5.4) or the exhaustive exploration of all 2^{14} metronomic therapies in the case of Section 5.6. For a fixed genotype, memory mechanism and molecular switch, these exhaustive searches take between one and six hours on a personal computer. The results presented here were evaluated by executing these exhaustive searches on a scientific computing cluster.

5.8.2 Testing

The deterministic model construction was tested by comparison to a hand–derived example for small x_b, y_b. The growth dynamics of the model was tested by comparison with analytically determined growth rates for specific parameter values. The implementation of bet–hedging within the model was tested by explicitly zeroing the molecules inherited by each individual and comparing to the memory–free bet–hedging model. The search for optimal metronomic

therapies was tested using a symmetric parameter set for which the optimal strategy is known. Through this chapter growth rates are determined from matrix eigenvalues numerically calculated by the SciPy packaged for scientific computing in Python [167]. We did not attempt to verify the correctness of this numerical calculation.

5.9 Discussion

We have explored the relative fitness contributions of three mechanistic processes that drive the determination of phenotypes in an evolutionary system; genetic inheritance, the GP–mapping and epigenetic inheritance. We began this chapter by highlighting assumptions of phenotypic 'memory' made in previous mathematical modelling of bet–hedging. We define this memory as a non–genetic dependence of an offspring's phenotype on that of its parent. Through our development of an evolutionary framework in Chapter 2, we identified that organismal phenotype must be determined solely from the contributions of environmental and inherited factors, as interpreted by a genotype–phenotype mapping. It was this perspective that underpinned the *strong–genetic* form of bet–hedging explored in Chapter 4, and highlighted the need for additional inherited factors in order for a bet–hedging population to exhibit memory. We have explored the role of epigenetic inheritance on the evolution of bet–hedging with specific focus on the different mechanistic modes of bet–hedging, which are rarely distinguished in the mathematical literature.

Where memory is included in mathematical models of bet–hedging (for example [232, 164, 356]), the (non–genetic) mechanisms through which phenotypic information is passed from parent to offspring are rarely explored. We have introduced a model of phenotypic memory wherein the epigenetic inheritance of intracellular factors serves as a mechanism to pass (imperfect) phenotypic information. This model follows naturally from the bistable switch model of bet–hedging introduced in Chapter 4 through the weakening of our model assumptions. Specifically, by weakening the assumption that intracellular factors (in our case, x and y) are entirely depleted at the time of reproduction. Instead, we consider these factors as subject to decay throughout the life of an individual and inherited, along with additional molecules produced by a genetic burst, by the offspring. This mechanism of epigenetic inheritance biases the phenotype of the offspring towards that of the parent when reproduction occurs sufficiently quickly. However, memory cannot be arbitrarily tuned (as in the model introduced by Müller et al. [232]) as it is dependent on organismal fitness, and thus environment.

In addition to the strong–genetic and epigenetic memory models of bet–hedging, we considered a third model, the constrained–memory model, in which epigenetic memory is fixed and genetically determined. Under this model, offspring take the same phenotype as their parents with likelihood, $p(g)$, determined by their genotype. We make no assumptions about the biological mechanisms that underpin this mode of bet–hedging, beyond noting that some other factors must be inherited (or indeed, environmental signals produced) that co–ordinate

the phenotype of parent and offspring. This model is a restricted version of the most general form of bet–hedging (Figure 5.1) studied by Müller et al. [232].

For these three models we explored the evolution of bet–hedging and its impact on drug resistance. This exploration represents an extension of previous bet–hedging studies, including our own presented in Chapter 4, through the consideration of different mechanisms of bet–hedging and memory.

5.9.1 The Evolution of Bet–Hedging in Constant Environment

We began by studying the behaviour of populations endowed with the three modes of bet–hedging and memory in fixed hospitable environments, aiming to generalise the results derived in Chapter 4. By considering the relationship between genotype (x_b), molecular switch and mode of epigenetic memory, we showed the contributions to fitness from these mechanisms to be non–additive. These results were summarised by the construction of mechanism (as opposed to genotype) fitness landscapes (Figures 5.9 and 5.10) that clearly show the existence of neutrality and epistasis in the landscapes. Previously, epistasis has been demonstrated to restrict evolutionary accessibility and introduce the potential for evolutionary steering (Chapter 3). Specifically, we identified combinations of molecular switch, genotype and environment for which epigenetic inheritance can be beneficial (DCx, $x_b = 20$, hospitable environment), fitness neutral (DCy, $x_b = 20$, hospitable environment) or deleterious (DCx, $x_b = 20$, drug–treated environment) when compared to the strong–genetic mode of bet–hedging. Similarly, we found that the fitness effect of changing from constrained to epigenetic–memory, or constrained–memory to strong–genetic bet–hedging, can also be advantageous, deleterious or neutral depending on the molecular switch, genotype and environment (Figure 5.10).

Our observations highlight non–additive inter–dependence between the factors underpinning bet–hedging. It follows that the fitness of a population, or predictions of its evolution, cannot be determined from observations of only one of these factors. Thus, a reductionist approach may fail to identify the key determinants in the evolution of bet–hedging (and its role in drug resistance). Certainly, if focus is given only to one of the factors explored here, for example the genes, then the role of the GP–mapping or epigenetic inheritance will be obscured. This also holds true for studies solely of the GP–mapping or solely of epigenetic inheritance. Conversely, as the details of each of the mechanisms of bet–hedging and memory become understood in ever finer granularity and complexity, the role of their functional interactions may be obscured. Just as focusing too narrowly on one aspect of organismal phenotypic determination does not give a clear, overall understanding of bet–hedging, neither does considering all biological/chemical components and their interactions. In short reductionism, coupled with increasingly more complex experimental tools, allows us (more than ever) to miss the wood for the trees. In Chapter 2, we introduced abstraction as an alternative paradigm that can remedy this problem. Here, we have applied this approach to clearly identify epistasis in a 'higher–order' landscape of biological mechanisms, highlighting the potential for divergence in their evolution (Figure 5.9), and demonstrating that each of the modes of bet–hedging with

memory can arise through Darwinian evolution. We believe that the identification of these non–linear relationships serves to prove the potential of the approach outlined in Chapter 2.

Finally, in our study of constant environments we highlighted a potential alternative (or complementary) solution to the evolutionary loss of bet–hedging paradox introduced in Chapter 4. Under the epigenetic–memory model, we find that the fitness benefit of successive mutations $x_b \rightarrow x_b + 1$) towards a one–phenotype strategy is further diminished when epigenetic memory is present in the system.

5.9.2 Fluctuating Environments

In previous studies, bet–hedging has been shown to confer a fitness advantage (over a one phenotype strategy) when environments fluctuate cyclically or unpredictably, suggesting that bet–hedging may evolve in such environments. Indeed, *de novo* evolution of bet–hedging was demonstrated by Beaumont et al. [26] through cycling selective pressures on the phenotypic property of colony morphology in the bacterium *Pseudomonas fluorescens*. However, the mathematical models commonly used to demonstrate the evolutionary advantage of bet–hedging are often predicated on an assumption of unconstrained phenotypic memory (Figure 5.1). We sought to partially verify these results in the case of restricted forms of bet–hedging memory.

We determined the expected growth rates for bet–hedging populations endowed with each of the mechanisms of memory, each molecular switch and each genotype in a set of cyclic environments. These environments consisted a cycle between d_1 days of drug–free environment followed by d_2 days of drug exposure, for $d_1 + d_2 = 28$ days and $d_1 \in \{0, 7, 14, 21, 28\}$. We found that, for each of the fluctuating environments, all of the modes of bet–hedging exhibited a higher fitness (for some x_b) than a one–phenotype strategy. Again, we considered combinations of memory, mapping and genotype to demonstrate that the optimal genotype x_b is strongly dependent on the underlying mechanisms driving bet–hedging and memory. This result again highlights the ineffectiveness of purely genomic studies of bet–hedging.

We found that each of the modes of bet–hedging and memory resulted in very similar fitness at the global optimum. This suggests that the present methodology of studying bet–hedging through models of unconstrained memory is not necessarily an incorrect one. We found that the epigenetic memory model results in higher fitness than the strong–genetic model at the majority of genotypes. If evolutionary convergence to a global optimum of fitness is expected to be fast (for example, as in the Strong Selection Weak Mutation model introduced in Chapter 3) then this will have little evolutionary impact. However, where genetic heterogeneity is high in the population, as is the case of many cancers [111], epigenetic memory reduces the fragility of individuals to the accumulation of mutational burden. These results provide at least some evidence that memory may be selected for in cancer populations. Empirical studies of bet–hedging or memory in the context of cancer are scarce. However Gupta et al. [134] estimated switching probabilities between lung–cancer cells exhibiting basal, luminal and stem-like phenotype (determined by gene expression) by training a Markov chain model

on the observed phenotype frequencies. They found that each of the phenotypes exhibited a unique set of transition probabilities, suggesting a degree of memory within the system. Our modelling highlights a potential flaw in the assumptions of Gupta et al. [134]: both genetic heterogeneity and epigenetic memory can result in heterogeneity in the transition probabilities of the model.

We have demonstrated the evolutionary advantage of bet–hedging only in a small number of cases, and for a single parameter set corresponding to a cytotoxic drug acting on a population of sensitive and persister–like cells. Thus, our results do not fully extend the mathematical results of [232] to the restricted modes of bet–hedging, as these results are parameter independent and hold for both cyclic and stochastically fluctuating environments. We opted to forgo a full mathematical treatment of these results in order to focus on the clinical question of the evolution of drug resistance. The mathematical derivation of a precise relationship between strong–genetic, epigenetic memory, constrained memory and unconstrained memory under different environmental regimes is a potential extension of the work presented here. Rather than pursue this approach, we opted to focus on the reverse question: given a bet–hedging population, how can we best choose an environmental regime to inhibit growth?

5.9.3 Implications for Therapy

To explore the relationship between bet–hedging and therapy we exhaustively explored a subset of all potential metronomic therapies. In particular, we considered all therapies consisting of 14-intervals in which each interval is designated as drug–treated or drug–holiday. Thus, this space consisted of 2^{14} potential metronomic strategies. This space was explored twice, for interval lengths $\iota = 7$days and $\iota = 1$day in order to assess the impact of dose fractionation on controlling bet–hedging driven disease. In our analysis we further subdivided the space of metronomic therapies by the total number of weeks of drug exposure (θ) over a 14–week interval. This subdivision allowed us to determine how much toxicity can be reduced whilst maintaining a specified degree of population control.

Our key finding is that, for a fixed intermediate θ value and $\iota = 7$, the variation in growth rate amongst the $\binom{14}{\theta}$ different metronomic therapies is large. In some cases (DC switch, $x_b = 59$, $\iota = 7$, $\theta = 7$) the variation in growth rates spans an order of magnitude with a majority of strategies resulting in population outgrowth but the optimal strategy driving population shrinkage. Further, we found that the maximum–dose–density strategy, in which the drug–dosing is prescribed continuously until the toxicity limit is reached, is the least effective treatment. This result is independent of the molecular switching mechanism, epigenetic memory or genotype and suggests that the current maximum–dose–density strategy may in fact be the worst possible approach for a restricted dose. Taken together, these results clearly indicate a need to reassess the present clinical practice for treating disease with bet–hedging driven resistance, and to instead consider the potential for metronomic treatment.

The MDD strategy for the $\iota = 7$ case is still a cyclic therapy, albeit one that switches most slowly of all those considered in our search. This observation led us to ask if dose fractionation

might drive better population control. We repeated the exhaustive analysis of length 14 drug cycles but with a smaller interval length of $\iota = 1$day and found that this alteration did not substantially improve the efficacy of optimal therapy in comparison to the $\iota = 7$ case. However, the distribution of growth rates associated with all therapies of dose restriction θ collapsed to a narrow distribution near the optimal. This finding suggests that where the dose is fractionated, the need to determine model derived optimal therapies is reduced. One interpretation of these results is that fractionated strategies are more robust to perturbations to the strategy. Such perturbations are likely common in clinical practice, for example because a patient must rearrange an appointment or forgets to take a pill, and in such instances fractionation could result in an improvement to clinical outcomes. Of course, there is a limit to the degree of fractionation. The pharmacodynamic properties of a drug, or simply the method through which it is administered, can limit the frequency with which the switch from drug–on to drug–off can be made. Thus, the fragility of metronomic therapies, and the need for model–derived optimal cyclic regimens, will likely be disease–, drug– and clinic– dependent.

By comparing the performance of optimal therapies for a given bet–hedging population when the molecular switch or mode of bet–hedging memory is changed, we found a less than 1% deviation in expected log weekly growth rate from the true optimal strategy. Indeed, we find that the growth rate of the optimal strategy between populations with different bet–hedging mechanisms is highly variable, but that the strategies themselves are relatively unchanged. Clinically, the precise mechanisms of bet–hedging are likely to be difficult to discern (at least in the short term), and this result can be considered good news. The precise mechanisms need not be known in order to derive (near–)optimal strategies, and incorrect assumptions regarding the dynamics of bet–hedging are unlikely to result in the design of massively suboptimal cycling regimen.

Finally, we altered the interpretation of the results of our exhaustive search to focus instead on dose reduction. Specifically, we asked for a specific dose restriction θ, what is the minimum dose with which we can achieve population control equivalent (or better) than the θ–restricted MDD therapy. For $\iota = 7$, we found that the dose can be nearly halved (for intermediate θ) whilst maintaining an equivalent degree of population control (Figures 5.18 and 5.19). Such a dose reduction can have important clinical impact. For example, where the total dose (as well as the dose over a specific time frame) is restricted, a dose–reduction can extend the treatment period, potentially increasing likelihood of stochastic extinction of the small resistant population. A second potential impact is a reduction in the cost of therapy. In recent years The National Institute for Health and Care Excellence (NICE), the UK body responsible for evaluating drugs for use in the NHS, has controversially de-listed a number of potentially life–extending drugs citing concerns over cost effectiveness [76]. Thus, therapeutic strategies that can minimise cost whilst maintaining efficacy offer a means through which to reintroduce these drugs to the clinic.

The Evolution of Memory

The epigenetic memory model of bet–hedging is dependent on the inheritance of intracellular molecules that are subject to decay throughout the life of an individual. The rate of this decay, determined by the parameter σ, determines the degree of memory and for large σ the model is equivalent to the strong–genetic form of bet–hedging studied in Chapter 4. As a mechanistic parameter, the rate σ could be subject to evolution. For example, if intracellular decay is driven by the presence of some enzyme, then over–production of this enzyme will increase σ, and where this increase is advantageous it may be selected for. We investigated the effect of alterations to σ on the population growth rate for both the MDD and optimal treatment strategies, both for $\iota = 7$ and $\iota = 1$ and for each of the molecular switches.

Critically, we found that there is not universal selection for increased or decreased σ. Instead, we see that the fitness effect of altering σ is strongly switch dependent, to the extent that the different switches can drive the evolution of σ in different directions. We also observe the existence of inaccessible evolutionary trajectories in the evolution of σ (Figure 5.20). These results again highlight the non–additive interactions between different mechanistic components. Indeed, the observation of slow or fast intracellular decay in an experimental system cannot be interpreted without understanding of the associated switching mechanism. Predictions of how the rate of decay, and the associated memory, will evolve are also likely to be impossible without understanding other aspects of phenotypic determination.

Consider the case wherein a hypothetical targeted therapy can alter the dynamics of the molecular switch. In this case, evolutionary steering could provide potential therapeutic benefit where, for example, a first drug selects for high σ and a second alters the molecular switch such that the high σ value is deleterious. Although hypothetical, this example demonstrates the potential for therapies that exploit trade-offs not only at the genetic scale, but at the scale of the GP–mapping. As many of the presently most effective targeted therapies work to inactivate or disrupt intracellular signalling pathways, this concept of 'mechanism' steering could be applicable to a number of disease, most notably, cancers.

Adaptive Therapy

To conclude our exploration of alternative therapeutic strategies, we compared the optimal metronomic strategy to a dose–skipping adaptive therapy inspired by [84]. These strategies prescribe a drug dose for an interval only when the population growth over the previous two intervals is sufficiently high. We found that, where drug dose restrictions are not severe, an adaptive strategy compares well with the overall optimal strategy. Thus, a heuristic approach is likely sufficient. This result is particularly good news for cancer therapies, where it is suggested that adaptive therapy will be effective at controlling disease with genetically–driven resistance by preventing competitive release. Our results suggest that mischaracterising the driver of resistance (genetic or bet–hedging) does not result in a major decrease in treatment efficacy.

In the case that restrictions on total drug dose are very severe, we find that these results begin to break down. In particular, the adaptive approach is too 'short–sighted' and prescribes the tolerable drug–dose too quickly. This results in rapid outgrowth where drug cannot be prescribed. In this case, a more conservative optimal therapy is shown to be significantly more effective. Whether this holds for adaptive therapies remains to be demonstrated experimentally. However, in a previous study, modelling combination therapies for EGFR+ NSCLC, we observed similar results for genetically driven resistance [Manuscript in preparation].

5.9.4 Conclusions

By applying the algorithmic perspective outlined in Chapter 2, we identified a number of previously undescribed modes of bet–hedging with memory. Previous studies have focused on a model of unconstrained phenotypic memory, but restrictions to this assumption yield models of bet–hedging with varying degrees of memory, for example the strong–genetic, constrained–memory and epigenetic–memory models described here. Identifying the differences between these models shows the potential of the algorithmic approach. It was only by considering evolution from the perspective of the basic functional components that the role of epigenetic inheritance, which was overlooked in previous studies, became clear.

By exploring the behaviour of bet–hedging populations under each of these models we identified the existence of non–additive relationships between the different mechanistic determinants of phenotype. We encapsulated these relationships in higher–order fitness–landscape graphs that clearly demonstrate the existence of neutrality, epistasis and the potential for ruggedness in the landscape of biological mechanisms (i.e GP–mapping, phenotypic memory and genotype). Through these 'mechanism' landscapes, we have related the nature of evolution at the GP–mapping scale to previously–derived models of evolution at the genetic scale and in doing so, presented a means to re–interpret the the fitness landscape results presented in Chapters 2 and 3 in the context of evolution at this mechanistic level. Further, the existence of two inter–connected fitness landscapes (in fact one higher–dimensional landscape) suggests a duality between evolution at the gene and GP–mapping scales.

In traditional evolutionary modelling, genetic alterations are interpreted by the GP–mapping (coupled with environment) to determine phenotype. However, through the correspondence presented here this relationship can be considered from an alternative perspective. A fixed *genotype* can serve the role of a *mapping* through which mutations to the GP–mapping are interpreted. This represents an appealing mathematical symmetry in evolution. Of course, both of these scales will evolve in parallel but the symmetry suggests a potentially useful experimental approach: To hold one aspect constant (i.e the genotype) whilst letting the other evolve. Experimental transplantation of genotypes could permit such an approach.

We have also explored the relationship between bet–hedging and metronomic therapy. Here, the results are clear: The present methodology of maximum–dose–density is not only poor, but potentially the worst possible. This holds true regardless of the mode of bet–hedging memory, the molecular switch or the population genotype. We find two potential approaches

to improve therapy. Firstly, dose fractionation improves population control. Secondly, where fractionation is not possible, the derivation of model–predicted optimal therapies can offer improvement. Critically, we find that these optimal strategies are robust to alterations of molecular switch or mode of bet–hedging memory, indicating that a precise understanding of the mechanistic drivers of bet–hedging is not necessary. Further, where the tolerable dose is not too low, we find that a heuristic adaptive therapy produces results comparable to an optimal strategy. Adaptive therapies have already been suggested (and are presently undergoing clinical trial, NCT02415621) as an efficacious therapy where diseases exhibit resistance driven by genetic mutations. Our results indicate that inadvertently applying these adaptive therapies where the resistance mechanism is instead bet–hedging will not itself result in a major failure of therapy.

In this chapter and in Chapter 4, we have seen that traditional cytotoxic treatments for diseases with bet–hedging resistance are unlikely to result in a cure. Indeed, it seems that control of the population growth is the best achievable objective for any cyclic approach. To fully cure these diseases we must develop drugs that interfere with the molecular mechanisms driving the stochastic determination of phenotypes. Likewise, to exploit ruggedness in the mechanism landscape we require drugs capable of targeting these molecular mechanisms. A necessary first step towards this goal is to clearly identify bet–hedging where it is indeed driving resistance. At present, few clinical examples have been identified outside of model bacterial organisms and *in vitro* studies of cancer cell lines (in which the drivers of bet–hedging are yet to be molecularly characterised). In the next chapter, we explore how properties of the GP–mapping are manifest at the population scale and how by observing the dynamics of a disease population, we may be able to identify bet–hedging in the clinic.

Chapter 6

The Impact of Genotype–Phenotype Mappings at the Population Scale

6.1 Introduction

In the previous chapters we have explored how aspects of the GP–mapping can influence evolution, driving drug resistance and ultimately causing the failure of therapy. Through models of genetic resistance and bet–hedging we demonstrated that the design of effective therapies is often predicated on understanding the GP-mapping. Unfortunately, many functional aspects of the mapping are at present unknown, owing to the overwhelming complexity of the biological processes that govern phenotypic determination. Indeed, most empirical studies consider only the effects of genetic aberrations, either in isolation or jointly, through knock–out studies or from a statistical perspective that identifies quantitative trait loci (QTLs). These studies of the GP–mapping are restricted, at least in part, owing to a dearth of model systems, difficulty identifying when and where different mechanisms of phenotypic determination are present, and a lack of empirical measurements at the phenotypic scale.

Where each individual in a population is endowed with an identical GP–mapping, average–case behaviour of the mapping can be determined by considering the population as a whole. For example, a simple statistic such as the proportion of different phenotypes can be used to infer properties of a bet–hedging GP–mapping. However, more complex population–scale patterns could also yield information. The spatial arrangement of genetic clones, of phenotypically distinct individuals, or the relationship between these arrangements, can be used to understand the GP–mapping. Here, we explore how such population–scale manifestations, which we call the *meta–phenotype*, can be studied to elucidate properties of the mapping. We provide two different examples in which meta–phenotypic properties of a population can yield insight into the mechanisms of the GP–mapping without the need to study complex intracellular processes.

We begin by exploring the distribution of phenotypes in a non–spatial model of a bet–hedging population and demonstrate how this distribution shifts when restrictions are placed on growth. From this observation we infer a relationship between bet–hedging and growth rate

that explains a phenomenon of bet–hedging populations previously described as 'puzzling' in the biological literature [204]. Specifically, we determine a simple model–derived hypothesis for the observation that the proportion of persister cells in a population of *E. coli* expands by more than an order of magnitude as the population approaches confluence.

Next, we explore how genetic sequencing can be used to infer mechanisms of resistance, even where those mechanisms are not genetically driven. Silent genetic mutations are known to accumulate in a population in a clock–like fashion. Thus, as a population evolves, its life history is recorded in the genome through the accumulation of such aberrations. Each time an individual reproduces, its complement of mutations is inherited by the offspring and further mutations can occur. Thus, mutations will exist in a population at a frequency determined by both the time at which they arose and the fitness of the clone in which they arose. By considering full distributions of mutational frequencies, which we refer to as *mutational spectra*, properties of population life–histories can be inferred. For example, the role of natural selection or neutrality in the growth of a population has been shown to leave distinct signatures in the mutational spectrum of a population [309, 354]. We will demonstrate that genetic and bet–hedging driven drug–resistance can be differentiated through consideration of mutational spectra following the outgrowth of a resistant population.

6.2 Impact of Exponential and Stationary Growth

In a review of persister cells and dormant disease, Lewis [204] highlight an unusual property of the *in vitro* growth dynamics of a population of persister and non–persister *E. coli* cells. Specifically, the ratio of persister to non–persister cells is not fixed, but rather changes throughout the growth phases of the population. In the early exponential phase, the non–persister cells dominate and increase as a proportion of the total population. Experiments, wherein *E. coli* are passaged to maintain exponential growth, result in a continuously expanding proportion of non–persisters, with the persister cells becoming undetectable in the population [177]. When the cells are not passaged, the growth dynamics differ. In the 'mid–exponential' (more closely resembling *linear*) phase of growth, wherein reproductive rate slows as the population size approaches a carrying capacity, the proportion of persisters within the population begins to expand. Finally, in the stationary (confluent) phase, the proportion of persisters reaches a maximum frequency in the population, settling at approximately 1%. Lewis [204] highlights that this phenomenon is presently unexplained. Quorum sensing – the regulation of gene expression through population–density–related extra–cellular signalling – was excluded as a mechanism driving this effect because the inclusion of spent growth medium (which should contain higher concentrations of such signalling factors) was found to have no effect on the proportion of persisters at different growth phases. Here we find, through mathematical modelling and simulation of population growth, that the meta–phenotypic pattern of growth identified by Lewis [204] arises naturally from the interaction of bet–hedging and the changing fitness differential of the two phenotypes under competition.

6.2.1 Modelling Growth Phases

To model the *in vitro* growth experiments for a bet–hedging population we used a simple two–phenotype bet–hedging model similar to that introduced in Chapter 4. We extend this model by introducing a carrying capacity K and scaling the birth rates of both phenotypes by $\left(1-\frac{N}{K}\right)$, where N is the total population size. This scaling results in logistic growth. For illustrative purposes we ignore the specific GP–mapping and the intracellular processes that drive bet–hedging and instead fix a likelihood value

$$p = \mathbb{P}\left[g \overset{\text{GP}}{\to} A\right]. \tag{6.1}$$

that we vary freely without consideration of the GP–mapping or genotype. As outlined in the previous two chapters, because we will not be considering how the value p evolves, the specific GP–mapping that generates this switching behaviour will not alter our results.

We consider two models of bet–hedging, the strong–genetic and constrained–memory models studied in the previous chapters. Where b_i, d_i are the birth and death rates (respectively) for the phenotypes $i \in \{A,B\}$, the discrete–time model population dynamics are given by

$$\begin{pmatrix} A(t+\Delta t) \\ B(t+\Delta t) \end{pmatrix} = \begin{bmatrix} (1-d_A)(1+c(t)b_A)p & (1-d_B)(1+c(t)b_B)p \\ (1-d_A)(1+c(t)b_A)(1-p) & (1-d_B)(1+c(t)b_B)(1-p) \end{bmatrix} \begin{pmatrix} A(t) \\ B(t) \end{pmatrix} \tag{6.2}$$

for the strong–genetic model and

$$\begin{pmatrix} A(t+\Delta t) \\ B(t+\Delta t) \end{pmatrix} = \begin{bmatrix} (1-d_A)(1+c(t)b_A)p & (1-d_B)(1+c(t)b_B)(1-p) \\ (1-d_A)(1+c(t)b_A)(1-p) & (1-d_B)(1+c(t)b_B)p \end{bmatrix} \begin{pmatrix} A(t) \\ B(t) \end{pmatrix} \tag{6.3}$$

for the constrained–memory model. Here, the population update matrices are time–dependent owing to a dependence on the competition–induced scaling factor

$$c(t) = \left(1 - \frac{A(t)+B(t)}{K}\right) \tag{6.4}$$

for carrying capacity K.

The birth and death parameters for this model are taken as in Chapter 4. In particular, time–steps are taken as $\Delta t = 1$hr and the birth and death rates as $b_A = 1.0\ \text{h}^{-1}$, $b_B = 0.015\ \text{h}^{-1}$ and $d_A = d_B = 0.005\ \text{h}^{-1}$ respectively. Empirical *in vitro* experiments begin by diluting a stationary–state population of *E. coli*, plating the population and observing growth. Thus, the initial growth dynamics of the population need not match those of a population that has experienced a sustained period of exponential growth. This may occur as, before the plating, the population is likely held at constant size (during the process of freezing/thawing a cell-line for example) and we attempt to simulate this experimental design. Growth simulations start by simulating 48 hrs of model growth with a small carrying capacity of $K = 10^3$. The initial

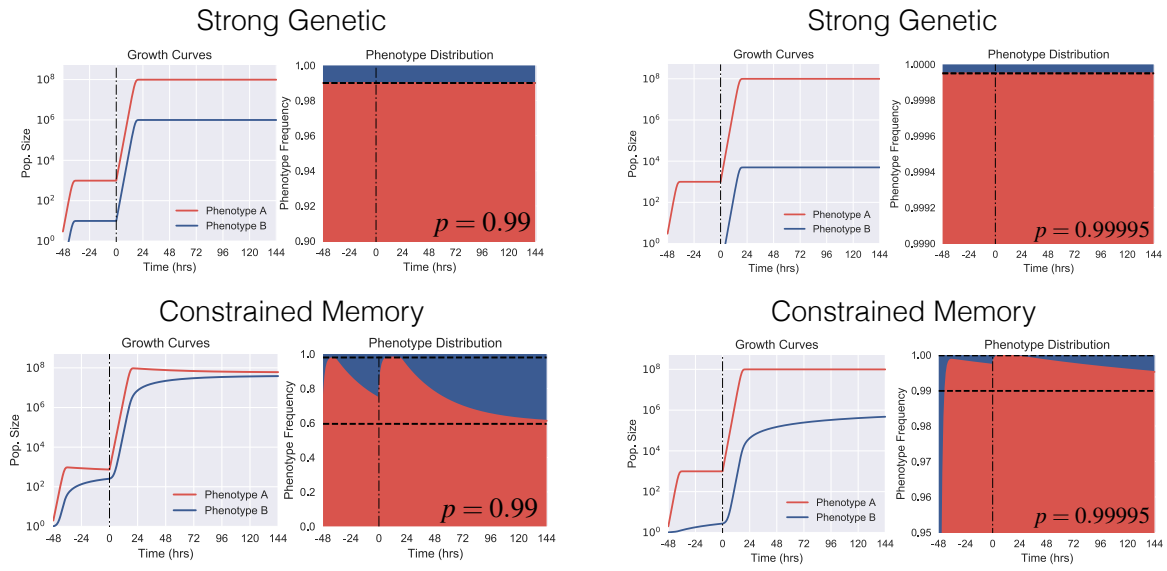

Fig. 6.1 **Growth dynamics of *in silico* bet–hedging populations with carrying capacity.** Growth curves and phenotype distributions over a period of 192 hrs of simulated *in vitro* growth are shown for four different bet–hedging populations. The top row depicts strong–genetic bet–hedging and the bottom row depicts constrained–memory bet–hedging. The inset p value denotes the switching probabilities, as defined in Equation 6.1. Each simulation begins with 48 hrs of simulation at small carrying capacity $K = 10^3$ before this value is increased to $K = 10^8$ (denoted by the vertical dashed line). Horizontal dashed lines denote the stationary distribution of phenotype A for exponential growth (higher dashed line) and the population at carrying capacity (lower dashed line).

conditions for this pre–simulation are taken by initialising a population of proportion p many A-phenotype individuals and $1 - p$ many B-phenotype individuals.

After 48 hrs of simulation, the carrying capacity K is increased to $K = 10^8$, allowing the population to enter a period of sustained exponential growth, followed by a slow down in growth rate and finally stationary phase. The population size and phenotype distributions are recorded over the initial 48 hrs, with $K = 10^3$, and an additional 144 hrs with $K = 10^8$. These dynamics were recorded for both the strong genetic and constrained memory models and for $p = 0.99$ and $p = 0.99995$. Population growth curves and the associated phenotype distributions are presented in Figure 6.1.

6.2.2 Results

The immediate observation is that the dynamics described by Lewis [204] are not captured by the strong genetic model, but that the constrained memory model closely matches these dynamics. Indeed, consider the start of the simulated *in vitro* growth ($t = 0$ hrs) where the carrying capacity is increased and the population begins to grow in an exponential phase. The stationary distribution of phenotypes in the exponential phase contains a larger proportion of the proliferative phenotype (red, A) than the population in the stationary growth phase for $K = 10^3$. Thus, we witness an initial decrease in the proportion of persister–like (blue, B)

individuals, consistent with the dynamics described by Lewis [204]. As the population approaches confluence, the reproductive rate (of both phenotypes) slows, resulting in a sub–exponential growth rate wherein the fitness differential between the two phenotypes is decreased. During these stages, the persister phenotype individuals grow to a larger proportion, again matching the dynamics described by [204].

The growth dynamics for both models can be easily explained by considering the contribution of the GP–mapping and fitness to the stationary distribution of phenotypes (i.e meta–phenotype). As the population approaches carrying capacity, the effective (scaled) birth rates b_A, b_B are diminishing and thus the difference between these values becomes smaller. For the constrained memory model this induces a shift towards a greater number of phenotype B individuals, whilst for the strong–genetic model no such change in frequency occurs. The precise dynamics of the phenotype frequencies can be derived from the population update matrices.

First, for the strong–genetic model, the time–dependent population update is given by

$$\mathbf{P}_{\mathrm{SG}} = \begin{bmatrix} (1-d_A)(1+c(t)b_A)p & (1-d_B)(1+c(t)b_B)p \\ (1-d_A)(1+c(t)b_A)(1-p) & (1-d_B)(1+c(t)b_B)(1-p) \end{bmatrix} \tag{6.5}$$

where $c(t) \in [0,1]$ varies as the population grows. For a fixed value of c, a stationary phenotype distribution can be calculated. Denote $w_A(c) = (1-d_A)(1+c(t)b_A)$ and $w_B(c) = (1-d_B)(1+c(t)b_B)$. Then, solving for the eigenvalues of $\mathbf{P}_{\mathrm{SG}}$ for fixed c yields

$$\begin{aligned} 0 = |\mathbf{P}_{\mathrm{SG}} - I\lambda| &= \begin{vmatrix} w_A(c)p - \lambda & w_B(c)p \\ w_A(c)(1-p) & w_B(c)(1-p) - \lambda \end{vmatrix} \\ &= \left(w_A(c)p - \lambda\right)\left(w_B(c)(1-p) - \lambda\right) - w_A(c)w_B(c)p(1-p) \\ &= \lambda\left(\lambda - w_A(c)p - w_B(c)(1-p)\right). \end{aligned}$$

Thus,

$$\lambda^*(c) = w_A(c)p + w_B(c)(1-p) \tag{6.6}$$

is the dominant eigenvalue for each $c \in [0,1]$. Solving for the associated unit eigenvector yields

$$\mathbf{P}_{\mathrm{SG}} \begin{pmatrix} p \\ 1-p \end{pmatrix} = \lambda(c) \begin{pmatrix} p \\ 1-p \end{pmatrix} \tag{6.7}$$

following the same argument as in Equation 4.8. This stationary distribution is independent of c and stays fixed through the exponential, mid and stationary growth phases. This is why the strong–genetic model does not exhibit the dynamics described by [204]. Conversely, for the

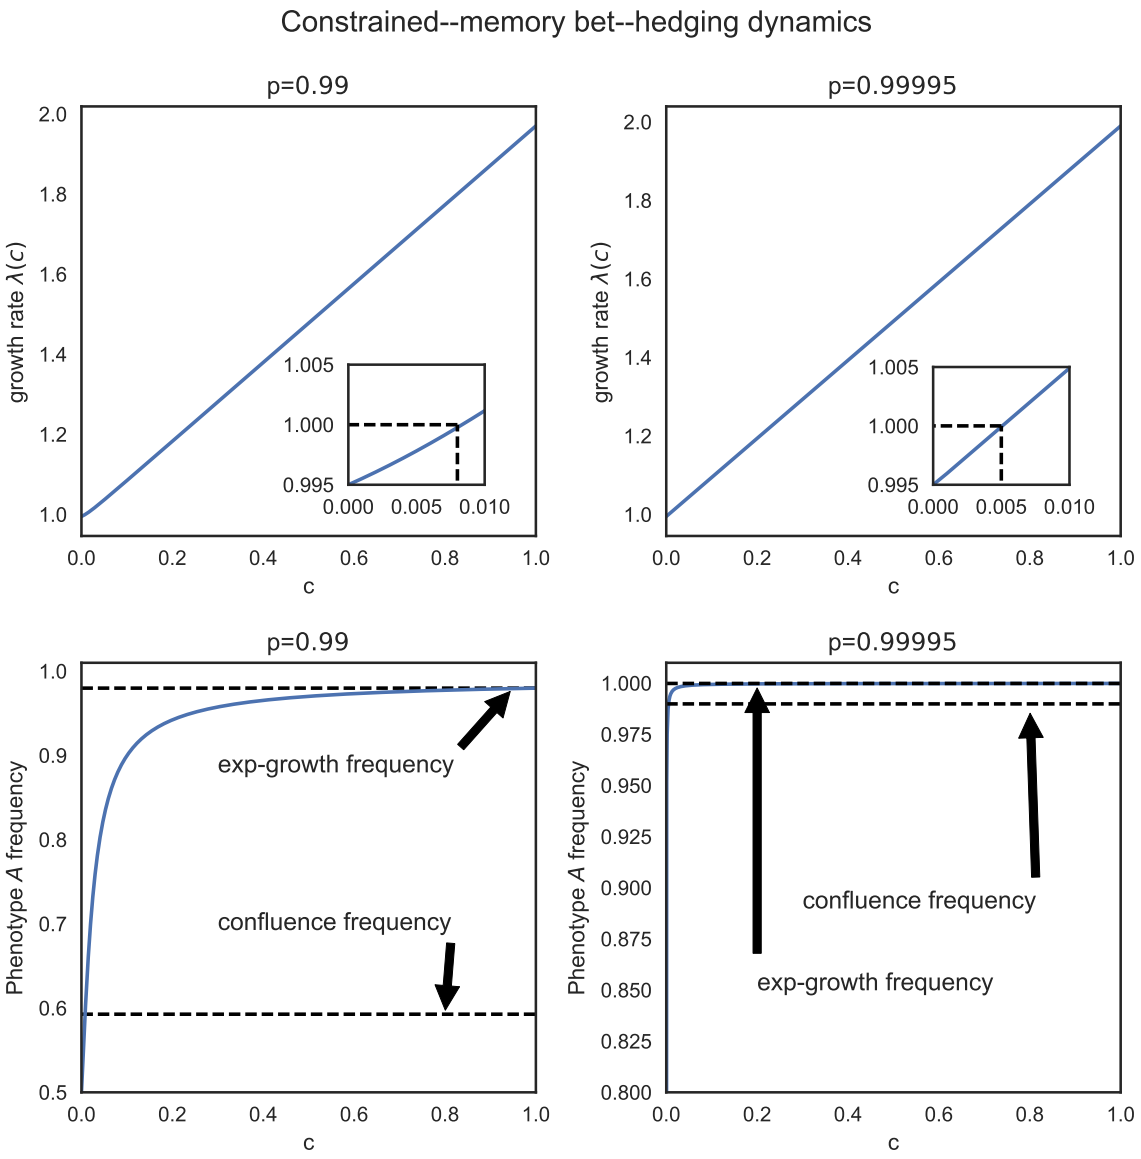

Fig. 6.2 **Growth rates and phenotype frequency for varying** c**.** The top row depicts the relationship between the carrying capacity induced competition c and the population growth rate λ^*. Inset shows the value of c for which $\lambda^* = 1$, this is the minimal value of c that is attained when the population reaches confluence. The bottom row depicts the stationary phenotype distributions for $c \in [0,1]$, the annotated horizontal lines indicate the phenotype distribution in exponential growth ($c = 1.0$) and at confluence (c as shown in the top row). The value of p is shown above each plot.

constrained memory model, $\mathbf{P}_{\mathrm{CM}}$, we have

$$0 = |\mathbf{P}_{\mathrm{CM}} - I\lambda| = \begin{vmatrix} w_A(c)p - \lambda & w_B(c)p \\ w_A(c)(1-p) & w_B(c)(1-p) - \lambda \end{vmatrix} \tag{6.8}$$

$$= \lambda^2 - w_A(c)p\lambda - w_B(c)p\lambda + w_A(c)w_B(c)p^2 = w_A(c)w_B(c)(1-p)^2 \tag{6.9}$$

$$= \lambda^2 - w_A(c)w_B(c)p\lambda + (2p-1)w_A(c)w_B(c). \tag{6.10}$$

Solving this equation gives a dominant eigenvalue of

$$\lambda^*(c) = \frac{\big(w_A(c)+w_B(c)\big)\,p + \sqrt{\big(w_A(c)+w_B(c)\big)^2 p^2 - 4(2p-1)w_A(c)w_B(c)}}{2}. \tag{6.11}$$

Unlike the strong–genetic case, the eigenvector corresponding to this dominant eigenvalue is not independent of c. This can be verified through a numerical solution. Figure 6.2 shows how the frequency of phenotype A individuals in the stationary distribution vary as c varies. Unlike the strong–genetic case, the phenotype distribution changes as the population approaches a carrying capacity and growth slows. The stationary distributions that are relevant to the *in vitro* growth experiments are those for the exponential growth phase ($N(t) \ll K$ such that $c(t) \approx 1$) and for the population at confluence. As cells experience death, $c = 0$ is never achieved. Instead, the value of c that is attained at confluence can be determined by solving

$$\lambda^*(c) = 1. \tag{6.12}$$

This value of c is dependent on b_A, b_B, d_A, d_B and p and can be determined numerically. For example, for our parameters $\lambda^*(c) = 1$ when $c = 0.008$ for $p = 0.99$ and when $c = 0.005$ for $p = 0.99995$ (Figure 6.2, top row insets). Critically, we see that the difference in phenotype distribution described by Lewis [204] arises naturally in a system where phenotypic memory is present in bet–hedging.

The population dynamics produced by our simulations do not completely capture the growth dynamics reported by Luidalepp et al. [210] who simulated continued exponential growth by repeated replating and observed that persister cells continued to diminish until they were undetectable. Our model predicts that under exponential growth, the proportion of persister cells should reach a stationary state. Whether this disparity is the result of a deficiency in our model or of technical limitations in empirical studies is unclear. For example, it may be that a stationary distribution is attained in empirical studies, but that persisters exist at such low frequency as to be undetectable. Our analysis partially supports this hypothesis. To attain a 1% proportion of persister cells at confluence (as described by [204]) we would require $p = 0.99995$ (See Figure 6.1). In this case, the exponential phase frequency of persister cells is approximately 0.01%, potentially falling below the experimental detection threshold. Further, if the proportion of persisters is very low then their numbers are likely subject to significant drift, further complicating their detection in the experimental population.

Through mathematical modelling with explicit consideration of the GP–mapping, we have demonstrated that bet–hedging in *E. coli* likely has some form of phenotypic memory. Coupled with previous empirical observations that a single cell can reconstitute the original population distribution (or at least, sensitivity to drug), we infer that this memory is likely not genetic, but rather epigenetic. By coupling experiments with mathematical modelling it may be possible to infer the dynamics of the bet–hedging mechanisms at the individual level by considering properties of the whole population (i.e. the meta–phenotype). We have

shown that the phenotype distributions in the exponential and confluent growth phases are determined by the parameters b_A, b_B, d_A, d_B and p. If the birth and death rates were known, then the value of p might be calculated via a parameter sweep. If these rates are not known, it may still be possible to perform a joint estimate of the parameters through Approximate Bayesian Computation (ABC) [315]. The success of such an approach is dependent on a close collaboration with an experimentalist, as factors such as the dynamics of the 'pre–experiment' procedures (thawing cells, growth to a small carrying capacity) would need to be controlled for.

6.2.3 Implications for Experimental Evolution

Previous experimental results, as well the results presented here, highlight the potential impact of re–plating or growth to confluence on the dynamics of a population. Where bet–hedging (with memory) is present in a population, the phenotype distribution will differ dramatically depending on the phase of growth. If experiments are designed explicitly to identify or explore bet–hedging, this effect can be readily accounted for. However, where experiments are performed to determine other properties of a population, e.g drug resistance, the underlying phenotypic distributions might not be accounted for.

Consider the experimental derivation of the fitness landscapes used to parametrise our model of sequential drug therapies in Chapter 3 [228]. In these experiments the initial growth period prior to the addition of antibiotic was overnight (16 hr) and followed by a 22 hr incubation in the presence of drug to determine an average growth rate via a microplate spectrophotometer. Such an experimental protocol measures growth exclusively in the early or mid–exponential phases and thus may underestimate or fail to account for drug–tolerance induced by persister cells. No mention of bacterial persisters is made in this study nor in those performed by Weinreich et al. [346] or Tan et al. [318] to derive empirical fitness landscapes. Failing to account for persisters has little impact on the *evolutionary* insights provided by these studies, as *E. coli* is used as a model organism for *genetically driven* adaptation. However, where attempts are made to extrapolate clinical protocols from these *in vitro* experiments, for example through collateral sensitivity analysis or evolutionary steering (Chapter 3), it is vital that the potential impact of the experimental protocol on measured drug resistance is accounted for. *In silico* modelling, specifically with explicit mechanistic simulation of the drivers of drug resistance (here, the GP–mapping), provides an ideal framework to quickly and efficiently explore the implications of experimental design in these cases.

6.3 Mutational Spectra and GP-Mappings

The properties of the GP–mapping can leave distinct genetic signatures in a population, even where phenotypic heterogeneity is not driven by genetic mutations. Each time a cells divides, genetic mutations stochastically occur. As only 1% of human DNA is exonic, and as there is

redundancy in the translation between base pair triplets and codons [334][Ch. 3], the majority of these mutations are neutral and have no phenotypic effect. As such, these mutations accrue throughout the life–history of a population in a clock–like fashion. The distribution of these mutations can provide insight into the nature of past selective pressures and the evolution of the population.

One example of how these clock–like mutations can inform us about evolution is the identification of proteins under stabilising selection via the consideration of dN/dS ratios [190]. Here, the ratio of the number of silent mutations (dN) to non–silent mutations (dS) is measured in the genes coding for a specific protein. Where the ratio is significantly less than one there is a signature of purifying selection for the protein, as non–silent mutations are removed from the population by natural selection. Identifying that a protein is under purifying selection implies that the protein is critical to the function of the cell (or organism) and this approach has been used to identify potential drug targets for novel cancer therapeutics [211]. Conversely, where the ratio is greater than one, the protein is identified as likely undergoing an evolutionary divergence. Where the ratio is close to one, evolution is identified as occurring neutrally, without strong selective pressure acting on the coded protein. Measuring the dN/dS requires a number of corrections, for example for the different rates of transition and transversion mutations, and may miss strong selection where, for example, purifying and divergent selection have both occurred in the life–history of the population. Despite these complications, the dN/dS ratio serves as a useful test for neutral evolution in the study of population genetics.

In cancer studies, mutational spectra have been used to characterise the extent of selection a tumour experiences during outgrowth (prior to therapy) [309]. Williams et al. [354] noted that, where a selectively advantageous mutation arises in a population, it does so in an individual carrying a complement of public (present in every individual in the population), sub–clonal (present in some, but not all individuals) and private (present only in that individual) mutations. This individual will give rise to a clonal lineage that outgrows non–mutant cells, resulting in a disproportionate representation of the mutations of that clone when the whole population is sequenced. The key contribution of Williams et al. [354] is the derivation of the distribution of mutational frequencies that is expected to arise during bulk sequencing of a tumour wherein all cells have equal fitness (i.e. a neutral expansion). Deviation from this frequency distribution can be considered evidence of strong selection acting during the growth the population and can be tested statistically. Bulk whole–genome sequencing of primary tumours is presently common clinical practice, and a wealth of data derived in this manner is available as part of The Cancer Genome Atlas project [350]. For this reason, the techniques derived by Williams et al. [354] are well suited for application to a large database of cancer genomic data.

Here we derive, by consideration of sequentially matched pre– and post–therapy mutational spectra for a simulated cancer population, a method to identify whether a resistance–conferring mutation or a non–genetic mechanism is responsible for drug resistance. To do this, we implement a model of neutral initial tumour outgrowth (similar to that used by Williams et al. [354]) to generate mutation distributions for a neutrally growing population. We then simulate

both selection for a specific clone, characterised by a single resistance conferring mutation occurring during the outgrowth, and selection for phenotypes that are randomly assigned to the population in a mutation–agnostic manner (i.e. bet–hedging). These two simulations result in predictions of mutational spectra following treatment for populations with genetic or bet–hedging driven resistance. We propose that comparison of these spectra could be used to highlight the drivers of resistance in populations in the same way that comparison with the neutral mutational spectrum was used by Williams et al. [354] to characterise the degree of selection acting on tumours during their initial outgrowth.

6.3.1 A Model of Mutational Accumulation

We implement a model of a growing population in which a number of neutral mutations arise at each cell division. Each cell in the population is represented by a list indicating the presence or absence of each mutation and at division this list is extended by a random number of new mutations and inherited by both daughter cells. We assume that no genetic site experiences multiple mutations, and hence, back mutations and mutational loss do not occur. This assumption is equivalent to there being infinitely many genetic sites and is known as the *infinite sites model* in population genetics [212, 181]. This model has been used extensively in population genetics studies of cancer, for example in the work of Williams et al. [354], Graham and Sottoriva [126] and others [309] where it has been demonstrated to capture the heterogeneity of neutrally arising mutations in growing cancer populations [355].

The number of new mutations is assumed to be Poisson distributed with parameter $\lambda = \mu = \mathbb{E}(\text{mutations per division})$. This assumption can be justified as follows. Suppose there are N genetic loci which each have independent probability p of mutating at a given cell division. In this case the number of new mutations arising at cell division is binomially distributed as $\text{Binomial}(N, p)$. For large N and small p, which we have explicitly assumed as part of the infinite sites assumption, the binomial distribution can be approximated by a Poisson distribution with parameter $\lambda = pn$ (this is the “Poisson limit theorem" and is easy to prove, for example as in [250]).

Finally, note that in reality mutations in cancer cells can be lost through structural or copy number alterations. The analysis performed here can be applied in cases exhibiting such copy number aberrations by only considering mutations arising in sections of the genome exhibiting normal copy number. This is because our analysis does not depend on specific mutations, but instead on the phenomenon, observed in all cancers, that neutral mutations accumulate throughout the genome during the growth of the tumour [15].

To simulate the population dynamics, we implement a Gillespie algorithm to simulate the exponential growth of a disease population. Initially, the population comprises a single mutation–free cell and an expansion is simulated. In reality, a founder cell for a cancer population will harbour a number mutations that are accumulated prior to neoplastic expansion. For the purposes of our model, these mutations can be ignored (or added in later) as they will be carried by all cells in the simulation.

We consider two drivers of phenotypic heterogeneity, genetic mutation and bet–hedging. For the genetic model, a specific mutation is designated as non–neutral and transforms individuals from phenotype A to phenotype B. In reality, resistance conferring mutations are thought to arise stochastically during the growth of a tumour, and thus can arise early, for example, driving primary resistance to a drug, or late, establishing a small drug resistant sub–population that re–establishes the tumour following a substantial reduction in tumour mass from therapy. To account for the role of early– or late– arising resistance–conferring mutations, we considered three cases for when a resistance conferring mutation arises. In the first, the 1000th mutation to arise is designated as resistance conferring, in the second, the 5000th mutation to arise is designated as resistance conferring and in the third, the 10000th mutation to arise is designated as resistance conferring. These values for the fitness conferring mutation number are chosen to correspond to an early, middle and late arising mutation within our model, which cannot simulate the accumulation of more than approximately 100,000 mutations tractably. This does not change the nature of our results as, as we will demonstrate below, where the fitness conferring mutation arises later in the population dynamics, the difference in the mutational spectra of genetic and bet–hedging–driven resistance are more pronounced. For the bet–hedging model, phenotypes are determined randomly and independently of the genetic mutations. We took a model of constrained memory bet–hedging with $p = 0.99$, following the methodology of previous studies. As before, phenotype A is assumed to be proliferative but drug sensitive whilst phenotype B is a slowly growing, drug–tolerant phenotype. For the drug–free environment, the birth rates of the phenotypes are denoted by b_A and b_B respectively whilst the death rates are denoted by d_A and d_B respectively. As before, we assume that a drug acts to increase the death rate of the phenotype A individuals.

We simulate drug treatment of the two population types. For the population with genetically driven heterogeneity, we simulate growth in a drug–free environment to a size of 25,000 cells after which the environment is switched to a drug–treated environment. The simulation is then continued until the population surpasses 25,000 cells (mirroring progression). For the bet–hedging population, an initial expansion to 25,000 is simulated, followed by 8 weeks of simulated drug therapy. Following those 8 weeks, the environment is switched back to a drug–free environment and the simulation continues until the population size reaches 25,000. Where the population reaches 25,000 cells within the 8 week drug period, the simulation is terminated. In both cases, the population mutational distribution is recorded immediately prior to the introduction of drug and at the end of the simulation, mirroring the collection of matched pre– and post– bulk tissue collection that occurs in a clinical setting. We note that a population size limit of 25,000 cells is small for a cancer population. This limit is chosen so that the model remains tractable to simulate, an approach that has been taken in previous modelling studies [354]. Pseudocode outlining the simulations of mutational accumulation for populations with mutation-driven resistance and bet-hedging-driven resistance are shown in Algorithm 4 and Algorithm 5 (respectively).

Algorithm 4 Population dynamics for the accumulation of mutations where resistance is driven by a single genetic driver mutation

procedure GROWTHSIMULATIONGENETIC(N_f, N_d, b_A, b_B, d_A, d_B, drugEffect, μ)
    ts, history ← [], []
    pop1, pop2 ← [2], []
    sizes ← ([],[])
    $t \leftarrow 0$
    druggedFlag ← False
    while len(pop1)+len(pop2) $\leq N_f$ and len(pop1)+len(pop2) > 0 **do**
        ts.append(t); sizes[0].append(len(pop1)); sizes[1].append(len(pop2))
        τ, pop1, pop2 ← Update(pop1, pop2, b_A, b_B, d_A, d_B, *mu*)
        $t \leftarrow t + \tau$
        if not druggedFlag and len(pop1) + len(pop2) $\geq N_d$ **then**
            druggedFlag ← True
            $d_A \leftarrow d_A \times$drugEffect (if cytotoxic)
            $b_A \leftarrow b_A \times$drugEffect (if cytostatic)
        end if
        history.append(pop1, pop2)
    end while
    return ts, history, sizes
end procedure

procedure UPDATEGENETIC(pop1, pop2, b_A, b_B, d_A, d_B, μ)
    Sample $r_1, r_2 \sim$ Uniform$[0,1]$
    $a \leftarrow [b_A \times$len(pop1), $d_A \times$len(pop1), $b_B \times$len(pop2), $d_B \times$len(pop1)$]$
    $a_0 \leftarrow$sum(a)
    $\tau \leftarrow \frac{1}{a_0} \log\left(\frac{1}{r_1}\right)$
    $r_2 \leftarrow a_0 r_2$
    if $r_2 \leq a[0]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop1)$\}$
        reproducer←pop1[i]
        pop1 ← pop1[:i) + pop1[$i+1$:]
        o1, o2, isDriver ←Reproduce(rep, μ)
        if isDriver **then**
            pop2.append(o1,o2)
        else
            pop1.append(o1,o2)
        end if
    end if
    if $r_2 \leq a[1]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop1)$\}$
        pop1 ← pop1[:i) + pop1[$i+1$:]
    end if
    if $r_2 \leq a[2]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop2)$\}$
        reproducer←pop2[i]
        pop1 ← pop2[:i) + pop2[$i+1$:]
        o1, o2, _ ←Reproduce(rep, μ)
        pop2.append(o1,o2)
    end if
    if $r_2 \leq a[3]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop2)$\}$
        pop2 ← pop2[:i) + pop2[$i+1$:]
    end if
    return τ, pop1, pop2
end procedure

procedure REPRODUCE(rep, μ)
    Sample $M \sim$ Poisson(μ)
    o1 ← rep
    muts ← Take(M, primes) (Consume M from a global list)
    for m in muts **do**
        o1 ← o1$\times m$
    end for
    o2←o1
    isDriver ← False
    if The specified 'driver' mutation count has been surpassed **then**
        isDriver ← True
    end if
    return o1, o2, isDriver
end procedure

Algorithm 5 Population dynamics for the accumulation of mutations where resistance is driven by bet-hedging.

procedure GROWTHSIMULATIONBH(N_f, N_d, b_A, b_B, d_A, d_B, drugEffect, μ, p)
    ts, history ← [], []
    pop1, pop2 ← [2], []
    sizes ← ([],[])
    $t \leftarrow 0$
    druggedFlag ← False
    while len(pop1)+len(pop2) $\leq N_f$ and len(pop1)+len(pop2) > 0 and not 8 weeks of drug treatment passed **do**
        ts.append(t); sizes[0].append(len(pop1)); sizes[1].append(len(pop2))
        τ, pop1, pop2 ← UpdateBH(pop1, pop2, b_A, b_B, d_A, d_B, μ)
        $t \leftarrow t + \tau$
        if not druggedFlag and len(pop1) + len(pop2) $\geq N_d$ **then**
            druggedFlag ← True
            $d_A \leftarrow d_A \times$drugEffect (if cytotoxic)
            $b_A \leftarrow b_A \times$drugEffect (if cytostatic)
        end if
        history.append(pop1, pop2)
    end while
    return ts, history, sizes
end procedure

procedure UPDATEBH(pop1, pop2, b_A, b_B, d_A, d_B, μ)
    Sample $r_1, r_2, r_3 \sim$ Uniform$[0,1]$
    $a \leftarrow [b_A \times$len(pop1), $d_A \times$len(pop1), $b_B \times$len(pop2), $d_B \times$len(pop1)$]$
    $a_0 \leftarrow$sum(a)
    $\tau \leftarrow \frac{1}{a_0} \log\left(\frac{1}{r_1}\right)$
    $r_2 \leftarrow a_0 r_2$
    if $r_2 \leq a[0]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop1)$\}$
        reproducer←pop1[i]
        pop1 ← pop1[:i) + pop1[$i+1$:]
        o1, o2 ←Reproduce(rep, μ)
        if $r_3 < p$ **then**
            pop2.append(o1,o2)
        else
            pop1.append(o1,o2)
        end if
    end if
    if $r_2 \leq a[1]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop1)$\}$
        pop1 ← pop1[:i) + pop1[$i+1$:]
    end if
    if $r_2 \leq a[2]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop2)$\}$
        reproducer←pop2[i]
        pop1 ← pop2[:i) + pop2[$i+1$:]
        o1, o2 ←Reproduce(rep, μ)
        if $r_3 < p$ **then**
            pop2.append(o1,o2)
        else
            pop1.append(o1,o2)
        end if
    end if
    if $r_2 \leq a[3]$ **then**
        Sample $i \sim$ Uniform$\{0,\ldots,$len(pop2)$\}$
        pop2 ← pop2[:i) + pop2[$i+1$:]
    end if
    return τ, pop1, pop2
end procedure

procedure REPRODUCE(rep, μ)
    Sample $M \sim$ Poisson(μ)
    o1 ← rep
    muts ← Take(M, primes) (Consume M from a global list)
    for m in muts **do**
        o1 ← o1$\times m$
    end for
    o2←o1
    return o1, o2
end procedure

Model Running–Time Optimisation

The naive modelling approach of maintaining a list of mutations for each cell in the population is prohibitively slow for the simulation of larger populations, as each new mutation must be recorded as absent or present in each cell when it arises. We circumvent this issue by introducing a mutation model based on the fundamental theorem of arithmetic. In this model, each subsequent mutation is represented by the next in a list of prime numbers and each individual is represented by an integer value. When a set of new mutations arise within an individual at birth, that individual's integer value is multiplied by each of the prime values for the mutations.

Upon cessation of the simulation, the complement of mutations within an individual can be determined by performing trial division with increasing primes. As the fundamental theorem of arithmetic ensures that each integer has a unique decomposition into prime factors, each individual in the population is guaranteed to be decomposed into its mutation set. The benefit of this approach is two-fold. The first is that the population dynamics, wherein the mutations require only multiplication, can be quickly performed and the (easily parallelisable) task of extracting the mutational signatures can be deferred until later. The second is that storing the mutational load in this way is more space–efficient than a list–based approach wherein (at least) a single bit is required for each mutation in each individual.

It may well be the case that a more efficient solution exists than the one presented here, for example using a space–reducing encoding of the mutation list. However, the prime number approach taken is fast enough for our purposes, easy to explain and validate, and somewhat mathematically elegant.

Parametrisation

The model is parametrised to correspond to an *in vitro* cell culture experiment, for example as performed by Hata et al. [145] wherein a potential mechanism of bet–hedging was identified. As such, we take $b_A = 0.042\ \mathrm{h}^{-1}$ and $d_A = 0.012\ \mathrm{h}^{-1}$ which gives a population doubling time of approximately 2.3 days. As before, we assume $b_B = 0.1 b_A$ and $d_B = 0.01 d_A$ to mirror a persister–like phenotype.

We considered two parameter sets for the effect of a drug. The first corresponds to a cytotoxic drug that increases the death rate of the phenotype A cells by a factor of 10. The birth rate of the phenotype A cells, as well as the parameters for the phenotype B cells, remain unchanged. Examples of drugs that behave in this manner include chemotherapies that induce death by targeted dividing cells. The second parameter set corresponds to a cytostatic drug that decreases the birth rate of phenotype A cells to zero, leaving the death rate of phenotype A cells, as well as the parameters for the phenotype B cells, unchanged. Examples of cytostatic drugs include some targeted therapies, such as erlotnib, that induce cell–cycle arrest [157].

For the mutation rate parameter μ, we note that estimates of mutation rate derived from genomic sequencing of cancers indicate that it can vary over two orders of magnitude, between

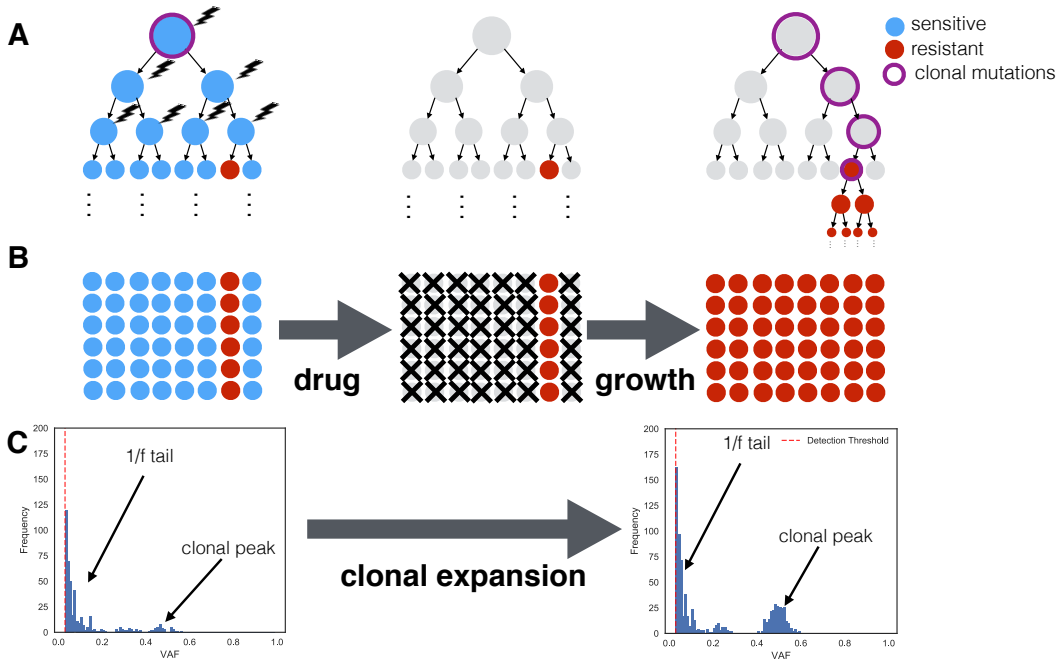

Fig. 6.3 **Schematic of the mutational spectra where drug resistance is genetically driven. A)** Schematic phylogenetic trees illustrating how selection for a subclonal mutation (subtree) results in a great number of clonal/truncal mutations (circled in purple) after selection. **B)** Natural selection for the resistant clone at the population scale. **C)** Model–derived VAF histograms of the pre–therapy population and the population after outgrowth in the presence of the drug. The expanded clonal peak after selection is highlighted. Pseudocode for the simulation of mutational accumulation with genetically driven resistance is shown in Algorithm 4.

10^{-5} and 10^{-7} mutations per base pair [354]. Here, we take an intermediate value within this biologically reasonable range and set a mutation rate of 10^{-6} per base pair per division. As the human exome comprises approximately 30 million base pairs this corresponds to approximately 30 exonic mutations per division. As mutations can be deleterious (and thus lost through selection), we take a conservative estimate of an average of $\mu = 10$ mutations per division, corresponding to the stochastic loss of every two out of three mutations. Whether this estimate is biologically reasonable remains an open question as the strength of negative selection in cancers remains a topic of debate [19, 215, 342].

Sampling Mutational Frequencies

The population is represented as a list of individuals wherein each is represented by its list of mutations. To mimic empirical data collection, we transform this population into a format that mirrors the data collection from empirical studies. In cancer genomics, mutations are identified and characterised by their *variant allele frequency* (VAF) – the proportion of alleles in the population that carry the mutation. The distribution of VAFs for the mutations identified in a sample can be used to identify the extent of natural selection in the life history of the population.

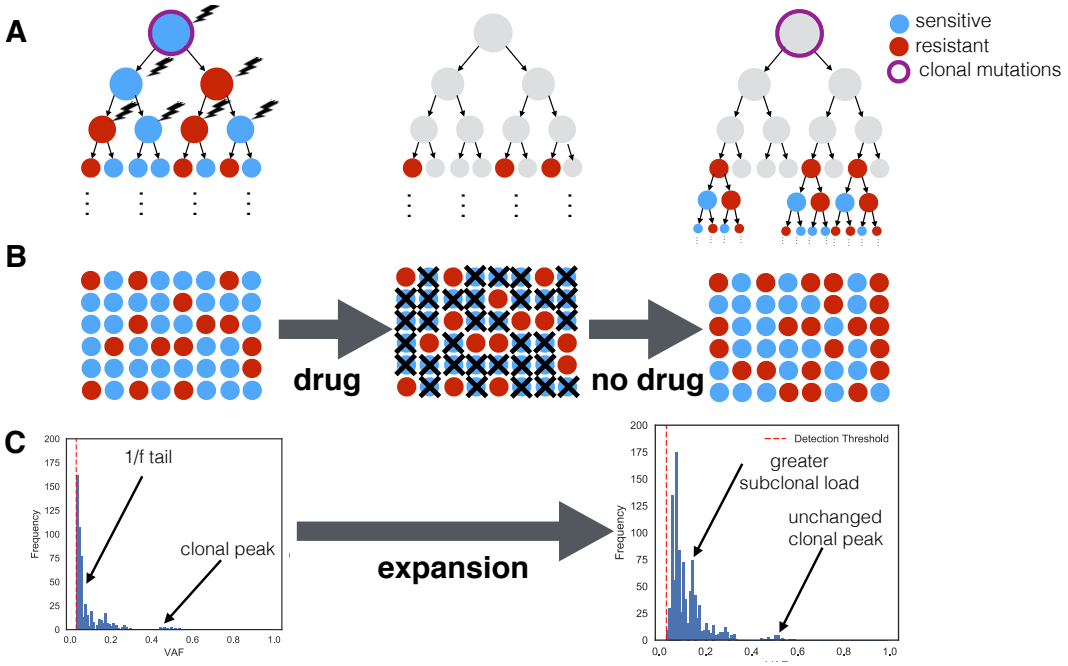

Fig. 6.4 **Schematic of the mutational spectra where drug resistance is driven by bet–hedging. A)** Schematic phylogenetic trees with nodes coloured according to the stochastically determined phenotype. Selection corresponds to an unbiased sub-sampling of the population leaving the number of clonal/truncal mutations unchanged. **B)** Natural selection for the resistant population initially homogenises the population, however, the phenotypic heterogeneity returns when therapy ceases. **C)** Model–derived VAF histograms for the pre–therapy population and the population after therapy ceases and outgrowth occurs. A greater number of mutations appear in the 1/f but the number of clonal mutations is unchanged. Pseudocode for the simulation of mutational accumulation with bet-hedging-driven resistance is shown in Algorithm 5.

In empirical studies, the VAFs are measured from next generation sequencing (NGS) of bulk tumour samples. The VAF is approximated as the ratio of reads (short DNA fragments, mapped to a location in the genome) that carry the mutation to reads that do not. This calculation is subject to noise introduced by a number of factors. Firstly, contamination by normal tissue introduces more alleles into the sample that do not carry specific mutations, resulting in a high number of reads that carry no mutation and systematic underestimation of the VAFs. Provided the normal contamination is not too severe, a correction can be applied to account for this. For this reason, we ignore contamination in our transformation of the model population to VAF distribution. Secondly, copy–number alterations can result in changes in VAF. For example, a clonal mutation in a tetraploid (three copies of each chromosome) population will exist in every individual but in only a third of alleles. This results in a VAF of 0.33, as opposed to 0.5 which is associated with a clonal mutation in a diploid population. This problem is confounded as different regions of the cancer genome often have different copy number owing to both gains and losses. However, as the study of mutational spectra is agnostic of specific mutations, areas of the genome that exhibit copy number changes can be ignored (and only mutations in diploid regions considered). For this reason, we assume that

the population in our model is diploid. Coupled with the infinite sites assumption, this means that where a cell harbours a mutation in our model, we assume that it contributed one mutated allele and one wild–type allele to the population.

An additional source of error in the approximation of VAFs is that the number of reads covering a specific base pair of the genome is often low (on the order 50-200 reads for whole genome or whole exome sequencing). As reads can correspond to any allele present in the population, this can induce shifts in the VAF owing to sampling bias. We simulate this sampling bias by means of binomial sampling. This assumption can be justified by noting that genomic sequencing in cancer (in a simplified sense) works as follows [298]. First, a small population of cells likely comprising both cancer cells and normal tissue, is extracted and the cells are lysed to pool their DNA together. The DNA is then fragmented by sonication. Suppose that a fraction, q, of these DNA fragments corresponding to a specific site in the human genome harbour a specific mutation. The fragments are then amplified by PCR. If we assume that the PCR amplification transforms each original fragment into the same number of fragments in the final pool, then a proportion q of this final DNA fragment pool will harbour the mutation. Note that in reality, PCR amplification can and likely does induce bias in the proportion q as some fragments are, owing to their chemical properties, more amenable to amplification through PCR. However, modelling the precise molecular mechanism of PCR and the associated biases is difficult, and many population genetics studies ignore this bias [309, 354, 355]. Finally, note that from the final DNA fragment pool not all fragments will be measured/read by the sequencer. Instead, only a fixed number, d, known as the depth of sequencing at the site will be read. These fragments can be modelled as extracted randomly and independently from the pool, and thus the number of reads harbouring the mutation can be modelled as binomially distribution with parameters d and q.

We compute the VAF for each mutation as follows. First, the fraction of cells in the population that harbour the mutation is computed. This value is the *cancer cell fraction* (CCF). Next, we sample a number of reads supporting the mutation according to

$$r_m \sim \text{Binomial}\left(d, \frac{\text{CCF}}{2}\right) \tag{6.13}$$

where d is the *depth* of sequencing (the average number of reads mapping to a given base pair of the genome). Note that the CCF is divided by 2 as each cell harbouring a mutation is assumed to contribute one wild–type and one mutant allele to the population. In our simulations we take $d = 120$, a reasonable[1] but good depth coverage for whole exome sequencing. From r_m, we compute the VAF as

$$\text{VAF} = \frac{r_m}{d}. \tag{6.14}$$

The mutational spectrum, in the case of our study, is the distribution of VAFs in the population. Figures 6.5(C) and 6.6(C) show example VAF distributions generated by our model. Here,

[1]Reasonable in the sense that is cost effective (less than $2000 for to sequence a human genome as 120x depth).

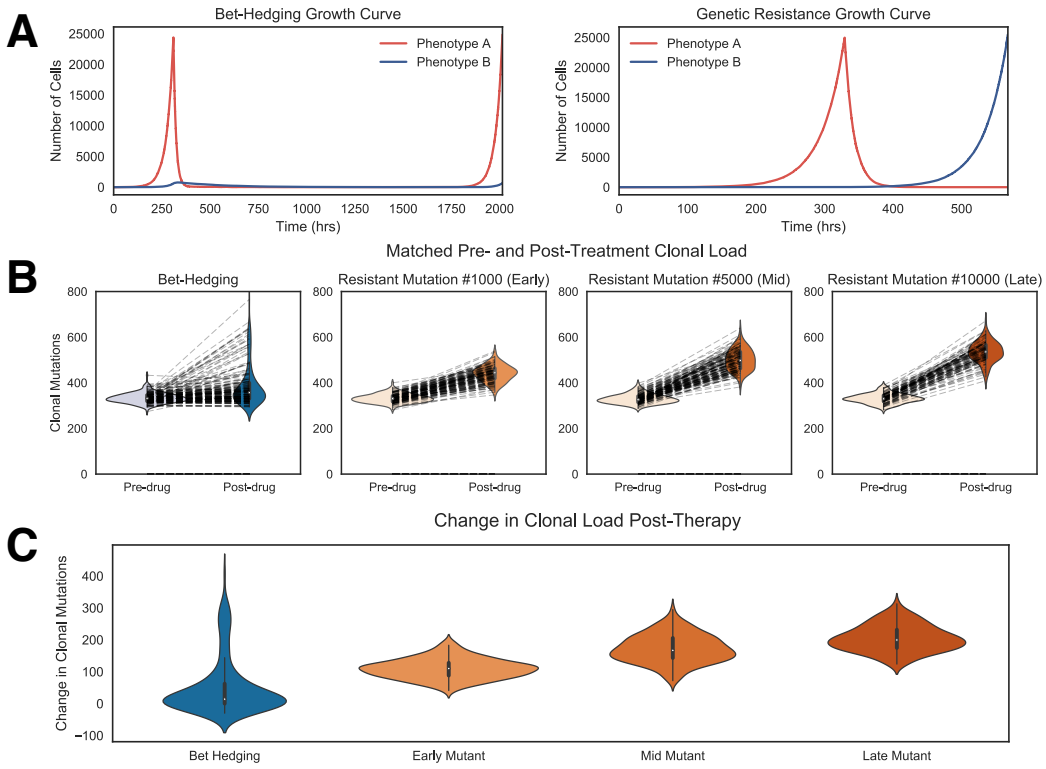

Fig. 6.5 **Mutational spectra of cytotoxic therapy. A)** Example growth curves for the bet–hedging population (left) and the population with resistance occuring at the 5000th mutation (right) under cytotoxic therapy. In both cases, the drug is introduced when the population reaches 25000 cells. For the bet–hedging population, the drug is simulated for 8weeks and then removed and the population grows in a drug–free environment to a size of 25000 cells. For the population with genetic resistance, the simulation is performed under drug treatment until the population reaches 25000 cells. **B)** Violin plots show the distribution of clonal mutations in the pre– and post–therapy populations. Plots are shown for 200 simulations of the bet–hedging population and for genetic resistance arising early (mutation 1000), mid (mutation 5000) and late (mutation 10000). Dashed lines connect the pre– and post–therapy clonal mutation number for individual simulations. **C)** Violin plots show the change in the clonal mutation number for the matched plots in **B**. The bar plots internal to the violins indicate the median (white dot) and quartiles.

we see the characteristic $1/f$ distribution described by Williams et al. [354] wherein the VAF histogram comprises two parts. The first is a binomial sampling distribution around 0.5 corresponding to the clonal mutations and the second is the power law distribution wherein the likelihood of a mutation having VAF f follows a $1/f$ power law distribution.

6.3.2 Mutational Signatures of Genetic Resistance and Bet–Hedging

The results of our simulation study are summarised in Figure 6.5 for the cytotoxic drug and Figure 6.6 for the cytostatic drug. In both cases, an example growth curve is shown (A) for both the bet–hedging population and the population with resistance driven by (the 5000th, 'mid') genetic mutation. Each simulation was performed 200 times. Part (B) shows violin plots

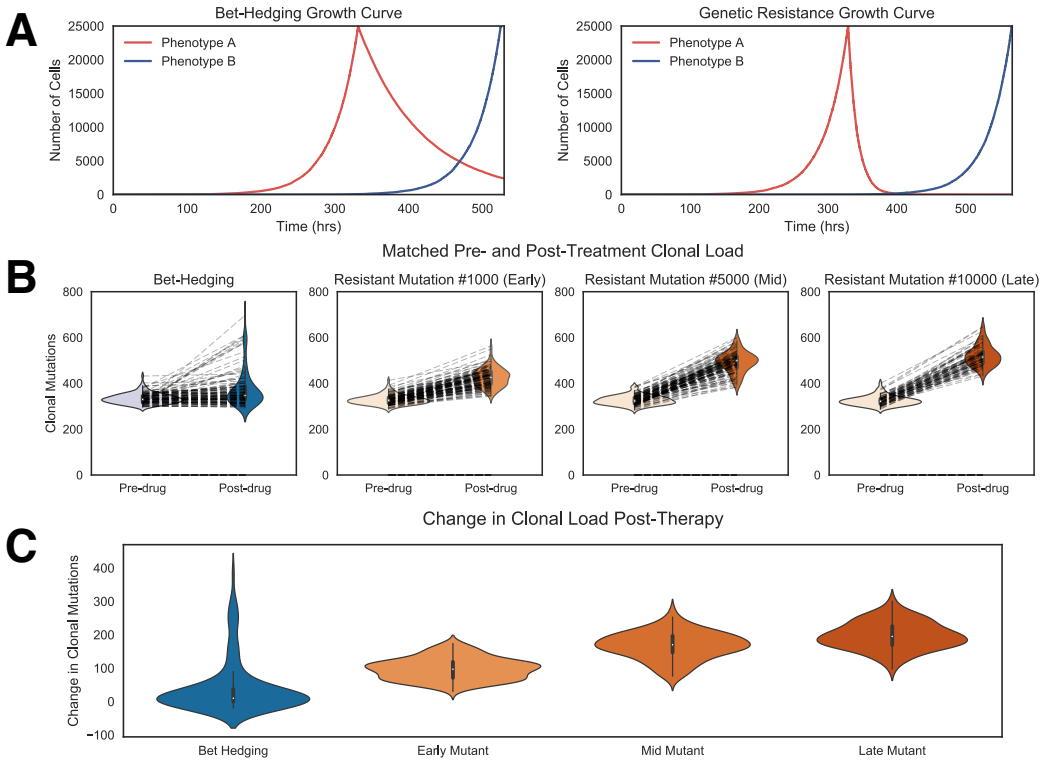

Fig. 6.6 **Mutational spectra of cytostatic therapy. A)** Example growth curves for the bet–hedging population (left) and the population with resistance occuring at the 5000th mutation (right) under cytotoxic therapy. In both cases, the drug is introduced when the population reaches 25000 cells and the population growth is simulated until the population exceeds a size of 25000 cells. **B)** Violin plots show the distribution of clonal mutations in the pre– and post–therapy populations. Plots are shown for 200 simulations of the bet–hedging population and for genetic resistance arising early (mutation 1000), mid (mutation 5000) and late (mutation 10000). Dashed lines connect the pre– and post–therapy clonal mutation number for individual simulations. **C)** Violin plots show the change in the clonal mutation number for the matched plots in **B**. The bar plots internal to the violins indicate the median (white dot) and quartiles.

of the pre– and post–therapy number of clonal mutations, defined as those with a variant allele frequency greater than 0.35. For each simulation, a random number of clonal mutations was added to the population to simulate the accumulation of mutations before neoplastic expansion. This number of clonal mutations was chosen to be distributed $\sim \text{Binomial}(600, 0.5)$. Dashed lines indicate matched pre– and post–therapy measurements. For the three cases of genetically driven mutations there is a clear increase in the number of clonal mutations for each simulation and each drug type. Further, the magnitude of the increase is greater where the resistance conferring mutation arises early, precisely as predicted (Figure 6.3).

For the bet–hedging population, there is a distinct split between those simulations in which there is little change in the clonal mutation number and those with a large increase. Observation of the growth curves (not shown) reveals that in some cases, including those in which a large increase in the clonal mutational load occurs, the population size becomes small during drug exposure – on the order of ten cells. In these cases, a large increase in the number of clonal

mutations can arise where all of the stochastically arising surviving cells group closely together (sharing many mutations) in the phylogenetic history of the population. This stochastic event can drive an increase in the clonal load without a genetic driver of resistance, indicating that the differentiation of bet–hedging and genetically driven resistance from a single pair of patched pre– and post–therapy samples may not be possible. This does not preclude the use of mutational spectra in a clinical setting. Indeed, recent experimental techniques known as *patient–derived xenografts* (PDX) allow the culture of patient–derived tumours in mice (known as *xenopatients*) in a manner that maintains clinical characteristics [31]. These PDX models can be used to simulate a number of therapies with the aim of identifying potentially effective strategies. This approach could be coupled with the mutational spectra models to identify bet–hedging by examining a number of matched pre– and post–therapy PDX tumours from a single patient.

In Figures 6.5(C) and 6.6(C) the change in mutational load (as opposed to the absolute number) is shown as a violin plot for the 200 *in silico* experiments. The bars inside the violin plots correspond to a barplot with the median represented by a white circle. From these plots we see that the bet–hedging population is characterised by a smaller average change in clonal load. For the cytotoxic drug, the simulations with bet–hedging exhibit a median gain of 14.0 clonal mutations (mean 51.8) whereas the simulations of genetic resistance exhibit median gains of 110.0, 167.0 and 119.5 (means 109.5, 172.0 and 204.5) for the early, middle and late arising resistance conferring mutations respectively. Similarly for the cytostatic drug, the simulation with bet–hedging exhibits a median gain of 10.0 clonal mutations (mean 41.7) whereas the simulations with genetic resistance exhibit median gains of 97.0, 110.0 and 199.5 clonal mutations (means 98.0, 167.1 and 196.8) for the early, mid and late arising mutations respectively.

These results imply that a statistical comparison of means (or distributions) of change in mutational load, for a small number of matched pre– and post–therapy samples, could serve as a simple test to identify bet–hedging in patient–derived samples or cell–lines. As for many diseases there is no known bet–hedging population with which to compare new populations, we must rely on model–derived populations for this approach. For example, pre– and post–therapy matched samples from a candidate population could be compared to a simulated bet–hedging population parametrised loosely to match the dynamics of interest. Alternatively, a theoretical distribution can be derived and used for comparison. Of course, this approach would be highly dependent on the underlying assumptions our model being correct. However, we note that many of the assumptions we have made are common in population genetic models of cancer, and have been used to characterise a number of cancers with moderate success [354, 355]. In the absence of directly observable bet–hedging, a comparison between patient derived mutational load and the predictions of a minimal model for bet–hedging may prove a useful tool, either by giving a strong indication that bet–hedging is present, or by helping us to discern where precisely our assumptions are incorrect.

Finally, we note that our model predicts bet–hedging populations can display a negative change in clonal load, indicating a reduction in the number of clonal populations. This occurs as a combination of our definition of clonal (VAF>0.35) and the binomial sampling error associated with genomic sequencing (which is simulated here) can cause sub-clonal mutations to be classified as clonal. These errors are an inherent issue in the derivation of mutational spectra from bulk sequencing, and in particular, the identification of clonal or near–clonal mutations. Our simulations indicate the noise associated with binomial sampling is sufficient to preclude the absolute classification of bet–hedging through the comparison of clonal mutational load and that, at least for the parameters explored here, bet–hedging may be misclassified as a clonal expansion by chance. Since these parameters are chosen to correspond to cell culture conditions, this suggests that an appropriate next step is the identification of bet–hedging, for example as reported by [145], in *in vitro* cancer cell–line experiments. This approach represents an avenue of research that will be continued in future, but is beyond the scope of this thesis.

6.3.3 Conclusions

By means of a simple model of population dynamics, we have demonstrated that drug resistance driven by genetic mutation and bet–hedging may result in distinct mutational spectra in the cancer genome, assuming that our minimal models of cancer growth are sufficiently accurate. This observation suggests a potentially novel method to identify bet–hedging through currently available genomic techniques. Specifically, we predict comparisons of the clonal mutations in matched pre– and post–treatment bulk tumour samples can highlight whether resistance–conferring mutations or bet–hedging are the drivers of drug resistance. To our knowledge, this is the first time that a genetic test for identifying non–genetic drivers of resistance has been suggested. Coupled with the other meta–phenotype approaches suggested here, this test could identify bet–hedging in a clinical setting.

The next step is to test whether this approach is successful in identifying bet–hedging in preclinical models. This requires that a known bet–hedging model organism is identified. Persister cells in cultured populations of *E. coli* might be one such model organism and the application of the method described here to *E. coli* may provide further evidence that the search for single genetic factors driving the emergence of persisters is futile. Conversely, bet–hedging–like populations in cancer, such as those identified by Hata et al. [145], may serve as good model systems. The high load of passenger mutations in cancers is especially useful for the identification of mutational spectra, as was noted by Williams et al. [354]. Non–small cell lung cancer cell–lines are likely the best candidate for this approach owing to a high mutation rate and previous studies that indicate potential bet–hedging mechanisms [145].

A second further direction is to extend the mutational spectra approach to the classification of other properties of the GP–mapping. For example, can the different mechanisms of bet–hedging, or different co–occurring mechanisms of phenotypic heterogeneity, be identified

from the mutational spectra measured from cancer populations? Answering these questions is beyond the scope of this thesis but is a direction we will certainly pursue in future.

6.4 Implementation Details and Computational Considerations

The implementation of the models, derivation of the results and the production of all of the figures was performed using the Python programming language. A full annotated implementation of all models, the derivation of all results, and the code used to produce the figures is available from Github (see Appendix 3). Appendix 3 contains a Jupyter notebook that provides an overview of this code.

6.4.1 Performance

The code used to generate the results in Section 6.2 executes near instantaneously. The simulations outline in Algorithms 4 and 5, coupled with the *in silico* simulation of genomic sequencing (Section 6.3.1), take between two and six hours to execute. The results presented here were generated by executing these algorithms on a scientific computing cluster.

6.4.2 Testing

The results of Section 6.2 are derived from the associated matrix models by use fo the SciPy package [167]. In Figure 6.2 the growth dynamics are shown alongside analytically derived phenotype distributions, verifying the correctness of our model implementation.

The simulations of mutational spectra were tested as follows. 1) The distribution of population sizes before treatment was compared to an analytically derived growth rate to verify that the population growth dynamics were correctly implemented. 2) The mutational spectrum after this initial growth period was compared to the characteristic '1/f' distribution for neutral growth derived by [354]. This test helps verify that the initial growth and mutational accumulation is likely neutral as intended. 3) The population size reduction associated with therapy was compared with an analytic prediction.

6.5 Discussion

By means of two different modelling approaches, we have demonstrated how properties of the GP–mapping can be manifest at the population scale: in the sizes of the different sub–populations and in the mutational spectra that arise naturally during the expansion of

the population. From these meta–phenotypic properties it may be possible to identify bet–hedging or phenotypic plasticity in clinical populations where these mechanisms of resistance have been proposed but remain elusive. The approach of mutational spectra analysis is especially promising, as the necessary data are readily collected as part of many rapid autopsy programmes or clinical trials.

We began by highlighting that a shifting distribution of phenotypes in a bet–hedging population need not be 'puzzling' as it has been previously described. Instead we identified that competition, in this case as imposed by a carrying capacity, alters the relative fitness of the two phenotypes and causes the shift in their frequency. This observation demonstrated that the manifestation of properties of the GP–mapping at the population scale need not be intuitive, as the meta–phenotype is ultimately governed by the GP–mapping, population interactions and natural selection. Critically, in our model we found that the expansion of the persister phenotype in the population is dependent on the existence of phenotypic memory, and that such an expansion did not occur with a strong–genetic bet–hedging population. This result indicates that bet–hedging in *E. coli* likely exhibits some degree of epigenetic memory, although this hypothesis must be verified empirically. More generally, this result highlighted how properties of the GP–mapping can be inferred from meta–phenotypic (population scale) properties which may be easier to measure than the molecular machinery that ultimately governs the mapping.

We next explored how bulk genomic profiling of the population might identify features of the GP–mapping. We identified that an expansion in the number of clonal mutations in the population is associated with the outgrowth of a population driven by the emergence of a resistant clone. Conversely, where bet–hedging drives drug resistance, natural selection acts to 'sub–sample' the phylogenetic tree (rather than select a specific branch associated with a resistance conferring mutation) resulting in an unchanged mutational spectra. Critically, we found that in our model the mean change in the number of clonal mutations is approximately zero where bet–hedging drives resistance. From this observation we suggested a simple statistical test to reject bet–hedging as a driver of resistance in the clinical setting. Whilst such a test will not explicitly *prove* the existence of bet–hedging, it may serve to narrow down potential candidates for further, more expensive assays. As this test relies solely on the use of matched pre– and post– therapy bulk genomic sequencing, there is a large pool of data on which it may be tested. Further, with the advent of additional rapid autopsy schemes, and the falling price of genomic sequencing, we can expect a rapid expansion in the number of appropriate data sets.

The identification of meta–phenotypic characteristics of the GP–mapping again highlights the potential of the algorithmic approach for studying evolution. By isolating the GP–mapping and modelling it abstractly we are able to distinguish properties of the population phenotype and genotype distributions that are ultimately driven by properties of the GP–mapping and not solely by the evolutionary processes of genetic mutation and natural selection. From these model predictions, we have generated testable experimental hypotheses that could help identify non–genetic drivers of drug resistance. As discussed above, a major impediment to the

study of non–genetic drivers of resistance, especially in oncology, is a relative dearth of model systems. As such, techniques to identify non–genetic heterogeneity could have considerable scientific value. The critical next step is to test our hypotheses, and in our future work this is one of the first steps we will take. Specifically, we will perform large–population–size *in vitro* experimental evolution of the EGFR+ positive lung cancer cell–line HCC827 (which was previously identified by [145] to potentially exhibit bet–hedging). Following drug exposure, both genetic sequencing and phenotypic assays will be performed to further characterise bet–hedging in this cell–line and explicitly test the mutational spectra hypothesis.

Chapter 7

Conclusions and Future Work

We began our work by highlighting the underlying algorithmic structure of evolution and emphasising that this structure is not restricted to biological systems. Instead, evolution can occur in any system consisting of imperfectly replicating entities with heritable variation in their propensity to survive and replicate. From this perspective, evolution can be observed in many domains including linguistics, economics, and critically, artificial systems [341]. Under the common lens of an algorithmic evolutionary framework, we explored a number of previously published evolutionary models, and in particular, genotype–phenotype mappings. Many of these models are still the subject of active research. Thus, applying the models introduced in Chapter 2 to new evolutionary questions represents an immediate way to extend or corroborate the results we have presented. Potential avenues of further research in other GP–mapping models were outlined in Chapter 2. Here, we present an overview of potential future work following from our modelling of the GP–mapping.

7.1 Evolutionary Steering and Fitness Landscapes

In Chapter 3, we coupled the fitness landscape model for the GP–mapping with strong assumptions on the dynamics of a disease population to generate a tractable model of evolving drug resistance. Critically, we predicted that the order in which drugs are prescribed is likely to alter the likelihood that highly–resistant disease arises through evolution. Through exhaustive enumeration of the evolutionary trajectories arising under short sequences of drugs, we found that arbitrary drug orderings may increase the likelihood of highly resistant disease whereas careful selection of drug sequences can ‘steer’ evolution such that highly resistant disease is inaccessible. Finally, we considered the present experimental approach to finding collaterally sensitive drug pairs, highlighting the potential for cross-resistance to be missed where too few replicates are performed. Taken together, our results indicate the need for improved experimental approaches for the design of sequential drug therapies. Identifying and verifying these approaches in a systematic and model–driven way represents a key extension of the work

presented here. Below, we present an overview the potential extensions of our work which could help to bring model–derived multi–drug therapies closer to the clinic.

7.1.1 Empirical Fitness Landscapes

Our modelling suggests that knowledge of (partial) empirical fitness landscapes could be used to design sequential therapies. However, as discussed in Chapter 2, the difficulty of deriving an empirical landscape grows exponentially with the number of mutations of interest. Specifically, for L mutations, 2^L strains need to be engineered and assayed for resistance to each drug of interest. Finding tractable ways to approximate empirical fitness landscapes will help accelerate the design of model–driven sequential therapies that can translate to a clinical setting.

One approach for empirically determining fitness landscapes is to infer the fitness of unseen genotypes from a training set of known genotype–fitness pairs. Where a large dataset of this kind exists, as it does for HIV [146], machine learning approaches may prove successful. This approach was taken to infer the fitness landscape of HIV under 15 drugs from a data set of approximately 70,000 patient samples, with moderate success. However, an issue with the generalised ridge regression (GRR) approach taken by Hinkley et al. [146], both in HIV or applied to other diseases, is that higher order epistasis cannot be considered. Indeed, doing so would introduce a large number of parameters, resulting in potential overfitting. Since higher order epistasis is prevalent in many biological systems [348], this may result in the prediction of incorrect landscapes.

A survey of alternative machine learning techniques for landscape inference has, to our knowledge, not been performed. It may be that an alternative approach can outperform the presently most effective algorithms. Thus, performing such a survey is a potentially valuable avenue of future work. The dataset used by Hinkley et al. [146] is published Open Access, and so anyone is free to perform such a survey. Unfortunately, no such accessible dataset presently exists for bacterial infections. However, it is common clinical practice to perform coupled genomic (targeted) sequencing and drug–sensitivity assays for many infections. If this data can be pooled, then it may prove an invaluable asset in determining empirical fitness landscapes. Any effort to pool this, presently distributed, clinical data would need to be spearheaded by a researcher with a strong clinical background, an understanding of the operation of the health care system in which the project is undertaken, and an appreciation of the potential power of mathematical modelling. As mathematical modelling becomes more mainstream we are optimistic that more clinicians working in infectious diseases will be open to working with mathematical modellers. Indeed, this is precisely what has been observed in both oncology and microbiology over the preceding fifteen years [13].

An alternative approach to deriving empirical landscapes is to perform laboratory evolution [199, 22, 175, 323, 161, 197, 196, 275, 120]. High throughput genomic sequencing allows us to determine the genotypes of any existing strains. For this reason, we are able to perform experimental evolution coupled with frequent genomic analysis to track evolutionary

trajectories. These trajectories contain information about the underlying fitness landscapes. Principally, we can determine an order relation on a subset of the genotypes by observing a trajectory. Where a large number of trajectories are measured, we may be able to infer (at least the ordering of) the local fitness landscape around a specific genotype. Further, the fact that a specific mutation is never observed to arise in a high number of replicates could also provide information, suggesting that it may be rare, deleterious or neutral. The central questions that underpin this approach are: what properties of the fitness landscapes can be learned from a collection of known trajectories? How many trajectories are we likely to need to make accurate predictions (is the problem tractable)? And finally, how does placing restrictions on model landscape properties such as epistasis change the answers to these questions? The tools of computer science are well suited to tackling these questions as they correspond naturally to related problems in learning theory and optimisation.

A key result of our work in Chapter 3 is that experimental evolution can fail owing to the non–repeatability of evolution. That evolution is non–repeatable is not a new result and has been observed previously in experimental evolution studies [199, 22, 120]. The results of Chapter 3 are novel in that they help clarify, and attempt to quantify, how this non–repeatability affects the prediction of collaterally sensitive drug pairs from experimental evolution studies. Where evolution *is* highly repeatable, *in vitro* evolution studies are potentially very effective. One way in which evolution can be rendered more repeatable is through the cultivation of populations that are very large. In such cases, the population will be genetically heterogeneous, containing many of the different resistant genotypes. This population will resemble, for example, a genetically heterogeneous pre–treatment cancer population and the selection pressure of a drug will likely result in the fittest mutants outgrowing and sweeping through the population in a predictable and repeatable manner (provided that the drug induces death in the majority of genotypes, and that the initial drug–resistant sub–populations are sufficiently small). Thus, *in vitro* experiments that aim to determine collaterally sensitive drug pairs from an initially large population may have some success in finding effective and repeatable drug sequences. However, following drug treatment, the population will shrink significantly and evolution may be rendered stochastic once more, again requiring us to account for stochasticity in evolutionary trajectories. Thus, effective drug sequences may be required to initially 'herd' a genetically heterogeneous population to a single genetic state and then 'steer' the evolution of this population in the sense of the model presented in Chapter 3. Following from these observations, we propose that future work in experimental evolution (particularly in cancers) should use large populations, specifically HYPERflasks with a capacity of approximately 200 million cells, and high drug concentrations to more closely resemble the true biology of drug resistance.

7.1.2 Extensions to the Model

The mathematical model presented in Chapter 3 makes a number of potentially biologically unreasonable assumptions on both the population size and mutation rate. In a number of

diseases, most critically many cancers, these assumptions are likely to break down. For this reason, the model is insufficient for understanding the emergence of drug resistance in large genetically heterogeneous populations. Thus, a key extension of the work we have presented here is to implement a model of sequential drug therapy that accounts for heterogeneous populations. An agent based model in which each individual is modelled explicitly and subject to stochastic mutation is the conceptually simplest such extension. Such models are common for the study of many diseases [272, 195, 12, 293]. At present, our collaborators at the Cleveland Clinic Department of Translational Hematology and Oncology are working on implementing agent based models of evolutionary steering (and herding) with the specific aim to determine whether drug sequences predicted by the SSWM model may still be effective when the strong assumptions on the population dynamics are weakened.

A second major assumption in our model is that neutral and deleterious mutations can never achieve fixation. In reality, such mutations are able to fix with non–zero probability in small populations. As drug treatment likely induces a massive reduction in population size it may be possible that neutral or deleterious mutations achieve fixation. In this case, the Markov chain model will resemble a previous model studied by Sella and Hirsch [295]. In this model, there are no absorbing states and evolution is reversible. Thus, steering in the absolute sense is impossible, although in reality the time frame of a patient's treatment is sufficiently short that we might expect no more than at most a single rare event to occur. Following from the hypothesis of adaptive therapy studied in mathematical oncology, we might ask how a drug–sequencing regimen can be administered in an alternative dosing strategy to avoid the fixation of critical deleterious or neutral mutations. For example, there may be critical genotypes along an evolutionary trajectory where the fixation of a neutral mutation may then permit the evolution to a highly resistant phenotype. Identifying such genotypes, and lowering the drug dose such that the population does not become too small, may prove effective. However, such an approach is inherently risky, as we may lower the dose in a situation where a high dose would have completely extinguished the disease population, resulting in cure. Mathematical modelling is the perfect tool to explore the risks and potential rewards of such an approach.

In our modelling, we partially justified the assumption of evolutionary convergence (which may require time scales longer than those applicable to a single patient) by the observation that the model is applicable at a whole–clinic scale, modelling the population of pathogenic bacteria present within an entire clinic, rather than in a specific patient. A potential extension of our model is to consider a *meta*–population model. Under this model, a number of distinct evolving populations (the bacteria present in distinct hospitals) are considered with periodic transfer of individuals between them (for example, through patient transfer). The dynamics of these transfers are determined by a network structure. These transfers will allow the arrival of new genotypes that are not immediate mutational neighbours of those present. Thus, the network structure is likely to have a strong impact on the likelihood of highly resistant disease emerging and we can ask what network structures can be implemented to avoid highly resistant disease. Such a study would resemble previous work that considers the likelihood of invasion

for a mutant arising in a structured population wherein each individual inhabits a unique node in a graph [6]. Our model extension would differ from these previous studies in that each graph node would correspond to a population and not an individual.

Owing to international borders, or simply indifference or misunderstanding on the part of individual hospitals, any antibiotic stewardship strategy is likely to be followed imperfectly or in only a subset of the hospitals. The key question then becomes, can we design steering strategies that are robust to imperfect adherence? Further, are there nodes of the hospital network whose adherence is critical in maintaining an effective steering strategy, even where other nodes do not follow the strategy. For a real network of hospitals, these critical nodes might represent the best place spend limited resources to educate staff about the importance of a clinical protocol. Any mathematical theory of steering on a network might also be applicable to a network of metastatic tumours should the hypothesis of inter–metastasis seeding prove true. In this case, imperfect adherence to a cycling strategy may arise due to localised differences in a the pharmacokinetic properties of a drug, owing for example due to local differences in metabolism or drug penetration.

7.2 Models of Bet–Hedging

In Chapter 4 we introduced a novel GP–mapping wherein bet–hedging arises naturally from the stochastic simulation of a chemical reaction network switch. Through this model we demonstrated that the structure of the chemical reaction (i.e. the GP-map) has important implications for the evolution of bet–hedging. Specifically, the CRN structure can inhibit the evolutionary loss of bet–hedging, maintaining a survival mechanism in the event of rare catastrophes. This model GP–mapping has, to our knowledge, not been studied before. Thus, following the development of previous GP–mapping models (Chapter 2), we can identify a number of potential avenues of further research.

Firstly, we note that we have only considered four CRN switches stemming from only two forms of a switch (DC and AM). Given the important evolutionary differences that we identified amongst these four switches, a clear extension of this work is to explore other bistable (or multistable) switches. A large number of such switches have been previously identified [42, 43] and a survey of the evolutionary implications of these switches for the robustness or persistence of bet–hedging may prove insightful.

For a number of other model GP–mappings, for example fitness landscapes, secondary structure prediction and gene regulatory models, studies have focused on characterising the neutral spaces wherein mutations do not effect fitness. In our work we have introduced an explicit distinction between mutations to the genotype and mutations to the abstracted mapping. It is (near–) neutrality between genetic mutations (those effecting the initial abundance of the molecules forming the chemical switch) that results in reduced evolvability and a slowed loss of bet–hedging in the CRN model. We have not yet studied how neutral alterations to the GP–mapping (i.e. the CRN) may affect evolution. Cardelli [42] introduced the concept of

emulation for chemical reaction networks, wherein one network can match the dynamics of a second. These emulations could be used to define a neutral mutation to a network as one resulting in a new network that can emulate the original. The duplication of chemical species within the network, driven (for example) by a gene duplication, could act as such a neutral mutation. We explored a similar phenomenon in the context of redundancy and robustness in Section 4.4. We hypothesise that the dynamics of evolution under such neutral mutations will mirror the phenomenon of punctuated equilibria in classical population genetics. Specifically, we propose that subsequent neutral mutations may expand a CRN until it can emulate the dynamics of many switching networks. A non–neutral 'jump' in the network dynamics might then occur from a single mutation. Exploring these evolutionary dynamics at the network scale is a potential extension of the results presented in Chapter 4.

7.2.1 Phenotypic Memory

In Chapter 5, we weakened our assumptions on the CRN model to explore the evolutionary implications of epigenetic inheritance and phenotypic memory. We began by noting that previous modelling of bet–hedging is predicated on an assumption of phenotypic memory [232] – with an offspring's phenotype determined stochastically from that of its parent. We sought to both clarify the potential contribution of phenotypic memory to fitness and to propose a potential mechanism for phenotypic memory in bet–hedging. Crucially, we identified that the GP–mapping, the genotype and the dynamics of memory interact non–linearly to determine organismal fitness in a specific environment. From these results we inferred that each of these components must be accounted for in order to predict the evolution of bet–hedging.

The explicit distinction between phenotypic memory and strong genetic bet–hedging has, to our knowledge, not previously been thoroughly explored. For future work, we propose that a comprehensive survey of previous bet–hedging results, many of which are derived under the assumption of unconstrained phenotypic memory, should be performed. Highlighting if and how theoretical results break down under different assumptions about phenotypic memory may serve to highlight how restricted, but mechanistically simple, modes of bet–hedging can achieve fitness on a par with unconstrained memory.

We have demonstrated the importance of understanding the dynamics of memory in predicting the evolution of bet–hedging, and thus in designing evolutionarily–informed therapies. We propose that empirical studies should now be performed to probe the dynamics of memory in biological models known to exhibit bet–hedging. The intracellular mechanisms governing the stochastic determination of phenotypes, and any associated phenotypic memory, are often not known or poorly understood in model organisms. Thus, a direct observation of any memory–associated molecular signature is impossible. Instead, memory might be identified through an approach following that outlined in Chapter 6 – by temporal observation of population–scale patterns of heterogeneity. In particular, the survey of theoretical results and memory mechanisms suggested above may highlight results that hold under some memory models and fail for others. A test of these results *in vitro* could represent an elegant, model–derived

experimental approach for determining the role of memory in bet–hedging. In our future work, we will attempt to verify and characterise bet–hedging in cancer cell lines (HCC827 specifically) grown *in vitro* through a combination of mathematical modelling and cell lineage tracking that utilises genomic barcoding technology [160].

7.2.2 Combining Phenotypic Plasticity and Bet–Hedging

In our review of phenotypic heterogeneity we highlighted a distinction between genetically–determined phenotypes, bet–hedging, wherein phenotypes are determined stochastically, and phenotypic plasticity, wherein the environment modulates organismal phenotype. We argue that each of these mechanisms can contribute to the phenotype, potentially in combination. Thus, an important extension of our work, and more broadly in studies of phenotypic heterogeneity in general, is to explore the effects of combinations of these factors. This approach has been partially undertaken previously. For example, the study of a genetically determined programme for phenotypic plasticity is precisely what is studied in the field of evolutionary developmental biology (evo–devo) [257]. The well studied ecological concept of the reaction norm (accounting for both environment and genotype) [328] is an another example of the power of distinguishing the factors that drive evolution and exploring their interactions.

We believe that the interaction between bet–hedging and phenotypic plasticity is a potentially important extension of our work on GP–mappings that has not yet seen sufficient study. One reason that the interaction of these two non–genetic mechanisms of phenotypic heterogeneity has not been studied may stem from confusion in terminology and a lack of clarity in defining the precise nature of phenotypic heterogeneity. 'Phenotypic plasticity' and 'bet–hedging' are often used interchangeably in discussions of phenotypic heterogeneity, despite referring to different ultimate causes for that heterogeneity (environments vs noise). Without this distinction, it is not clear that there really are two separate phenomena, and thus an interaction between them to study.

We propose that environmentally–mediated bet–hedging will emerge as a fruitful abstract GP–mapping to study, with potentially important biological insights to be gained. Some justification for this hypothesis can be drawn from the molecular chaperones and their role in mediating environmental stress and genetic mutation through a phenomenon known as *evolutionary capacitance* [280]. Molecular chaperones are responsible for correcting misfolded proteins within the cell cytoplasm. Such misfolded proteins often arise owing to mutations in their associated genes, and chaperone proteins ultimately render such mutations neutral. However, under certain environmental conditions (for example, heat stress), chaperone proteins are denatured and the proteins that were previously corrected by them take an aberrant form. Thus, under certain environmental stresses a potentially large number of neutral genetic mutations are suddenly rendered non–neutral, inducing phenotypic heterogeneity in the population and permitting a rapid 'search' of phenotypic space, guided by latent genetic heterogeneity, for a better adapted phenotype. This phenomenon can be considered environmentally–induced

mutagenesis and to lie at the intersection of genetic and plastic mechanisms of phenotypic heterogeneity.

We propose that environmentally–induced or modulated bet–hedging could serve a similar role to evolutionary capacitance. The evolutionary dynamics of such a mechanism of heterogeneity is well suited for study through the framework introduced in Chapter 2. In our study of epigenetic memory in bet–hedging (Chapter 5), we noted that fitness (as measured by reproductive rate) offers a proxy for environmental conditions, and highlighted that the environmentally driven variability in phenotypic memory in this model is a kind of plastic bet–hedging. More complex models should now begin to be studied. For example, in future we will extend the chemical reaction model such that the initial conditions are determined by both an inherited (epi–)genetic component and a component that is determined by an individual's environment. Such dynamics will induce spatial heterogeneity in the dynamics of bet–hedging. Owing to the multi–scale nature of such a model, with many components interacting non–linearly, it is not intuitive how such dynamics will effect evolution, and in particular, drug resistance. Mathematical modelling is likely the best way to explore the dynamics of environmentally–induced bet–hedging and, following similar approaches to those outlined in Chapter 6, design experimental approaches to test for its existence empirically.

7.3 Identifying Model Systems

In Chapter 6, we explored how features of the GP–mapping can be manifested at the population scale. As the intracellular dynamics underpinning the GP–mapping are overwhelmingly complex, we proposed that the identification of model–predicted features of a population may provide a means to identify bet–hedging or phenotypic plasticity in model systems or clinical samples. Specifically, through mathematical modelling we outlined how population–scale patterns of phenotypic heterogeneity, which we call the *meta*–phenotype, could potentially arise from the integration of genetic heterogeneity, the GP–mapping and heterogeneous environments. As population–scale properties are more readily measurable in experimental systems, especially in comparison to the complex dynamics of intracellular dynamics at the single cell scale, this meta–phenotype perspective suggests an experimental approach that integrates mathematical modelling with empirical observations of experimental systems to probe the GP–mapping.

In our future work, we will continue to develop this hypothesis with the aim to design model–driven experiments to identify bet–hedging in cancer cell–line systems or clinical samples. In particular, in addition to empirically testing the signature of bet–hedging discussed in Chapter 6, wherein drug treatment induces no change in clonal mutational load, we aim to determine additional genomic signatures through mathematical modelling. For example, many cancer populations consist of cells derived from a phenotypic hierarchy in which all cells are derived from a pool of cancer stem cells through phenotypic differentiation [251]. A potential extension of our work is to determine whether different topologies of phenotypic hierarchies

induce different genomic signatures. Such signatures could be particularly useful for studying some cancers, such as colorectal cancer or leukaemia, in which cancer stem cells likely exist but are difficult to observe directly. Alternatively, genomic signatures of phenotypic memory in bet–hedging, coupled with *in vitro* cell culture and deep genomic sequencing, could help to discern the precise dynamics of bet–hedging in cell–lines where it is hypothesised.

7.4 Summary and Final Conclusions

We began this work by introducing a philosophy of mathematical modelling that emphasises the power of abstraction over a reductionist approach attempting to simulate biological reality. Many mathematical modellers might summarise their scientific approach through the quote (falsely attributed to Albert Einstein[1]) that "everything should be made as simple as possible, but no simpler." We argue that it is possible for models to be much simpler than many people might assume. Here, we have followed this modelling philosophy to demonstrate the critical importance of the GP–mapping in understanding evolution.

Throughout our work we have attempted to keep the clinical and biological significance of our results at the forefront of our discussion. Indeed, it is the clinical question of drug resistance that motivated our study of evolution to begin with. This was not a new approach. Many studies have advocated for evolutionarily–informed approaches to managing disease, and these approaches are now making their way into clinical practice [84, 103, 100]. However, evolutionary approaches have all too often been driven heavily from a genetic perspective. It is the phenotype that is subject to natural selection and it is drug–resistant phenotypes that ultimately drive mortality in many diseases. We know that these phenotypes are determined by more than simply genetics. Both the environment (phenotypic plasticity) and chance (bet–hedging) play an important role in the determination of phenotypes. Our aim in this work was to integrate these non–genetic mechanisms of phenotypic heterogeneity with the paradigm of predictive evolutionary modelling and evolutionarily informed therapy. To do so we highlighted the fundamental algorithmic structure of evolution (Chapter 2) and demonstrated that evolution is neither simply mutation at the genetic scale, nor selection at the phenotypic scale, but rather the adaptation that arises from the intersection of these biological processes. We argue that the relationship between genotype and phenotype, the GP–mapping, is critical in understanding and predicting evolution and thus this mapping should be studied in its own right.

We can summarise our key results as follows. In Chapter 2, we introduced an algorithmic evolutionary framework and performed a literature survey of previous modelling of GP–mappings to demonstrate how they fall under this common algorithmic structure. In Chapter 3, we used a simple GP–mapping model to demonstrate the power of predictive evolutionary

[1] Einstein's actual quote is "It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience" [80], which one might argue is certainly not as simple as possible.

modelling — highlighting that if we are able to make evolutionary predictions then we may be able to design effective sequential drug therapies. However, this model failed to account for non–genetic phenotypic heterogeneity, which is known to play a key role in drug resistance. In Chapter 4, we introduced a novel GP–mapping model for bet–hedging wherein phenotypes arise stochastically from a genetically parametrised bistable switch. Through this model we resolved an apparent paradox of bet–hedging wherein phenotypic heterogeneity persists where it is necessarily deleterious. We highlighted that it is the structure of the GP–mapping that can reduce evolvability and maintain bet–hedging where it is selected against. In the context of therapy, we determined that the efficacy of drug holidays will be strongly dependent on the structure of the GP–mapping.

In Chapter 5, we extended the bistable switch model to consider different mechanisms of phenotypic memory. Specifically, we considered the cases where organismal phenotype is determined purely randomly (strong–genetic bet–hedging), dependent on the parental phenotype (constrained memory bet–hedging) or influenced by the inheritance of any residual molecules involved in the bistable switching mechanism (epigenetic memory). Critically, we identified that memory can be either advantageous or deleterious in a switch–dependent manner. We found that the genotype, molecular switch and memory mechanism all interact in a non–additive fashion, which we characterised in a fitness–landscape–like construct to demonstrate epistasis–like interactions between these aspects of the GP–mapping. Further, in the context of the epigenetic model, we found that the molecular switching mechanism determines whether selection acts to increase or decrease phenotypic memory. Finally, we explored the therapeutic potential of different drug dosing schedules and demonstrated that both metronomic and adaptive therapies will outperform a maximum–dose–density approach. This result was observed to be independent of the underlying mechanisms of memory, suggesting that the precise dynamics of bet–hedging need not be known prior to clinical intervention.

Finally, in Chapter 6, we highlighted the dearth of model systems for studying non–genetic mechanisms of phenotypic heterogeneity, particularly in cancers. We suggested that population–scale features (i.e. meta–phenotypes) can reflect properties of the underlying GP–mapping. We began by highlighting that growth–phase–dependent changes in the frequency of phenotypes can arise naturally in a bet–hedging population that exhibits memory. This simple mathematical result might help explain the *in vitro* dynamics of *E. coli* [204]. We then demonstrated that non–genetic mechanisms of phenotypic heterogeneity may leave distinct genomic signatures following drug treatment. As genomic profiling is common clinical practice, this genomic approach to identifying bet–hedging is readily applicable to existing datasets. If successful, this approach could have a significant impact on treatment decisions for any of the diseases in which bet–hedging is identified, a fact supported by our results highlighting how the efficacy of different therapeutic strategies is dependent on the underlying drivers of phenotypic differentiation.

Throughout all of our work a common theme has arisen — if we are to predict evolution, as we must to predict and counteract the emergence of resistance to our presently most effective

drugs, then we must gain a mechanistic understanding of the GP–mapping. Consider the results of Chapters 4 and 5. We demonstrated that it is often not sufficient to simply know that bet–hedging occurs. Instead, we must identify the precise mechanisms that drive this bet–hedging, for example the underlying molecular switch or dynamics of memory, if we are to predict its evolution. We hope that the results presented here will encourage both mathematical and biological researchers to probe the GP–mapping through an integrated mathematical and experimental approach that uses abstract modelling to generate biological hypotheses. Certainly, this will be the goal of our future research.

In 1973, Dobzhansky wrote that "nothing in biology makes sense except in the light of evolution". Certainly we believe this to be true, but perhaps we should add further, that nothing in evolution makes sense except in the light of the genotype–phenotype mapping.

References

[1] Acar, M., Mettetal, J. T., and van Oudenaarden, A. (2008). Stochastic switching as a survival strategy in fluctuating environments. *Nature Genetics*, 40(4):471–475.

[2] Adkins, S. and Shabbir, A. (2014). Biology, ecology and management of the invasive parthenium weed (Parthenium hysterophorus l.). *Pest Management Science*, 70(7):1023–1029.

[3] Aguilar-Rodríguez, J., Payne, J. L., and Wagner, A. (2017). A thousand empirical adaptive landscapes and their navigability. *Nature Ecology & Evolution*, 1:0045.

[4] Aguirre-Ghiso, J. A. (2007). Models, mechanisms and clinical evidence for cancer dormancy. *Nature Reviews Cancer*, 7(11):834–846.

[5] Alberch, P. (1991). From genes to phenotype: dynamical systems and evolvability. *Genetica*, 84(1):5–11.

[6] Allen, B., Lippner, G., Chen, Y.-T., Fotouhi, B., Momeni, N., Yau, S.-T., and Nowak, M. A. (2017). Evolutionary dynamics on any population structure. *Nature*, 544(7649):227–230.

[7] Altenberg, L. (1995). Genome growth and the evolution of the genotype-phenotype map. In *Evolution and Biocomputation*, pages 205–259. Springer.

[8] Altenberg, L. (1997). NK fitness landscapes. *Handbook of Evolutionary Computation*.

[9] Alto, B. W., Lampman, R. L., Kesavaraju, B., and Muturi, E. J. (2013). Pesticide-induced release from competition among competing Aedes aegypti and Aedes albopictus (Diptera: Culicidae). *Journal of Medical Entomology*, 50(6):1240–1249.

[10] Ambler, R., Coulson, A., Frère, J.-M., Ghuysen, J.-M., Joris, B., Forsman, M., Levesque, R., Tiraby, G., and Waley, S. (1991). A standard numbering scheme for the class a beta-lactamases. *Biochemical Journal*, 276(Pt 1):269.

[11] Ancel, L. W. and Fontana, W. (2000). Plasticity, evolvability, and modularity in rna. *Journal of Experimental Zoology*, 288(3):242–283.

[12] Anderson, A., Chaplain, M., Rejniak, K., and Fozard, J. (2008). Single-cell-based models in biology and medicine.

[13] Anderson, A. R. and Quaranta, V. (2008). Integrative mathematical oncology. *Nature Reviews Cancer*, 8(3):227–234.

[14] Andersson, D. I. and Hughes, D. (2010). Antibiotic resistance and its cost: is it possible to reverse resistance? *Nature Reviews Microbiology*, 8(4):260–271.

[15] Andor, N., Graham, T. A., Jansen, M., Xia, L. C., Aktipis, C. A., Petritsch, C., Ji, H. P., and Maley, C. C. (2016). Pan-cancer analysis of the extent and consequences of intratumor heterogeneity. *Nature medicine*, 22(1):105–113.

[16] Andrews, J. M. (2001). Determination of minimum inhibitory concentrations. *Journal of Antimicrobial Chemotherapy*, 48(suppl 1):5–16.

[17] Angluin, D., Aspnes, J., and Eisenstat, D. (2008). A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102.

[18] Ashlock, D. (2006). *Evolutionary Computation for Modeling and Optimization*, volume 103. Springer.

[19] Bakhoum, S. F. and Landau, D. A. (2017). Cancer evolution: No room for negative selection. *Cell*, 171(5):987–989.

[20] Balaban, N. Q., Merrin, J., Chait, R., Kowalik, L., and Leibler, S. (2004). Bacterial persistence as a phenotypic switch. *Science*, 305(5690):1622–1625.

[21] Balázsi, G., van Oudenaarden, A., and Collins, J. J. (2011). Cellular decision making and biological noise: from microbes to mammals. *Cell*, 144(6):910–925.

[22] Barrick, J. E., Yu, D. S., Yoon, S. H., Jeong, H., Oh, T. K., Schneider, D., Lenski, R. E., and Kim, J. F. (2009). Genome evolution and adaptation in a long-term experiment with escherichia coli. *Nature*, 461(7268):1243–1247.

[23] Bartlett, J. G. (2011). A call to arms: the imperative for antimicrobial stewardship. *Clinical Infectious Diseases*, 53(suppl 1):S4–S7.

[24] Barton, N. (2001). The role of hybridization in evolution. *Molecular Ecology*, 10(3):551–568.

[25] Bassler, B. L. (1999). How bacteria talk to each other: regulation of gene expression by quorum sensing. *Current Opinion in Microbiology*, 2(6):582–587.

[26] Beaumont, H. J., Gallie, J., Kost, C., Ferguson, G. C., and Rainey, P. B. (2009). Experimental evolution of bet hedging. *Nature*, 462(7269):90–93.

[27] Becskei, A., Séraphin, B., and Serrano, L. (2001). Positive feedback in eukaryotic gene networks: cell differentiation by graded to binary response conversion. *The EMBO Journal*, 20(10):2528–2535.

[28] Becskei, A. and Serrano, L. (2000). Engineering stability in gene networks by autoregulation. *Nature*, 405(6786):590–593.

[29] Bendall, D. S. (1985). *Evolution from Molecules to Men*. CUP Archive.

[30] Bergstrom, C. T., Lo, M., and Lipsitch, M. (2004). Ecological theory suggests that antimicrobial cycling will not reduce antimicrobial resistance in hospitals. *Proceedings of the National Academy of Sciences*, 101(36):13285–13290.

[31] Bertotti, A., Migliardi, G., Galimi, F., Sassi, F., Torti, D., Isella, C., Corà, D., Di Nicolantonio, F., Buscarino, M., Petti, C., et al. (2011). A molecularly annotated platform of patient-derived xenografts ("xenopatients") identifies HER2 as an effective therapeutic target in cetuximab-resistant colorectal cancer. *Cancer Discovery*, 1(6):508–523.

[32] Bigger, J. (1944). Treatment of staphylococcal infections with penicillin by intermittent sterilisation. *The Lancet*, 244(6320):497–500.

[33] Blake, W. J., Balázsi, G., Kohanski, M. A., Isaacs, F. J., Murphy, K. F., Kuang, Y., Cantor, C. R., Walt, D. R., and Collins, J. J. (2006). Phenotypic consequences of promoter-mediated transcriptional noise. *Molecular Cell*, 24(6):853–865.

[34] Bonhoeffer, S., Wiesch, P., and Kouyos, R. D. (2010). Rotating antibiotics does not minimize selection for resistance. *Mathematical Biosciences and Engineering*, 7(4):919–922.

[35] Bornberg-Bauer, E. (1997). How are model protein structures distributed in sequence space? *Biophysical Journal*, 73(5):2393.

[36] Bornberg-Bauer, E. (2002). Randomness, structural uniqueness, modularity and neutral evolution in sequence space of model proteins. *Zeitschrift für Physikalische Chemie International journal of Research in Physical Chemistry and Chemical Physics*, 216(2/2002):139.

[37] Botero, C. A., Weissing, F. J., Wright, J., and Rubenstein, D. R. (2015). Evolutionary tipping points in the capacity to adapt to environmental change. *Proceedings of the National Academy of Sciences*, 112(1):184–189.

[38] Brock, A., Chang, H., and Huang, S. (2009). Non-genetic heterogeneity—a mutation-independent driving force for the somatic evolution of tumours. *Nature Reviews Genetics*, 10(5):336–342.

[39] Buchler, N. E. and Goldstein, R. A. (2000). Surveying determinants of protein structure designability across different energy models and amino-acid alphabets: A consensus. *The Journal of Chemical Physics*, 112(5):2533–2547.

[40] Bukh, J., Miller, R. H., and Purcell, R. H. (1995). Genetic heterogeneity of hepatitis C virus: quasispecies and genotypes. In *Seminars in Liver Disease*, volume 15, pages 41–63.

[41] Bull, J. J., Meyers, L. A., and Lachmann, M. (2005). Quasispecies made simple. *PLoS Computational Biology*, 1(6):e61.

[42] Cardelli, L. (2014). Morphisms of reaction networks that couple structure to function. *BMC systems biology*, 8(1):84.

[43] Cardelli, L. and Csikász-Nagy, A. (2012). The cell cycle switch computes approximate majority. *Scientific reports*, 2:656.

[44] Carja, O. and Plotkin, J. B. (2015). The evolutionary advantage of heritable phenotypic heterogeneity. *bioRxiv*, page 028795.

[45] Chan, H. and Bornberg-Bauer, E. (2001). Perspectives on protein evolution from simple exact models. *Applied Bioinformatics*, 1(3):121–144.

[46] Charlebois, D. A., Abdennur, N., and Kaern, M. (2011). Gene expression noise facilitates adaptation and drug resistance independently of mutation. *Physical Review Letters*, 107(21):218101.

[47] Charlebois, D. A., Balázsi, G., and Kærn, M. (2014). Coherent feedforward transcriptional regulatory motifs enhance drug resistance. *Physical Review E*, 89(5):052708.

[48] Childs, D. Z., Metcalf, C. J. E., and Rees, M. (2010). Evolutionary bet-hedging in the real world: empirical evidence and challenges revealed by plants. *Proceedings of the Royal Society of London B: Biological Sciences*, 277:3055–3064.

[49] Chmielecki, J., Foo, J., Oxnard, G. R., Hutchinson, K., Ohashi, K., Somwar, R., Wang, L., Amato, K. R., Arcila, M., Sos, M. L., et al. (2011). Optimization of dosing for EGFR-mutant non–small cell lung cancer with evolutionary cancer modeling. *Science Translational Medicine*, 3(90):90ra59–90ra59.

[50] Chou, H.-H., Chiu, H.-C., Delaney, N. F., Segrè, D., and Marx, C. J. (2011). Diminishing returns epistasis among beneficial mutations decelerates adaptation. *Science*, 332(6034):1190–1192.

[51] Chouard, T. (2008). Beneath the surface. *Nature*, 456(7220):300–303.

[52] Ciliberti, S., Martin, O. C., and Wagner, A. (2007). Innovation and robustness in complex regulatory gene networks. *Proceedings of the National Academy of Sciences*, 104(34):13591–13596.

[53] Citron, M. L., Berry, D. A., Cirrincione, C., Hudis, C., Winer, E. P., Gradishar, W. J., Davidson, N. E., Martino, S., Livingston, R., Ingle, J. N., Perez, E. A., Carpenter, J., Hurd, D., Holland, J. F., Smith, B. L., Sartor, C. I., Leung, E. H., Abrams, J., Schilsky, R. L., Muss, H. B., and Norton, L. (2003). Randomized trial of dose-dense versus conventionally scheduled and sequential versus concurrent combination chemotherapy as postoperative adjuvant treatment of node-positive primary breast cancer: first report of Intergroup Trial C9741/Cancer and Leukemia Group B Trial 9741. *Journal of Clinical Oncology*, 21(8):1431–1439.

[54] Consortium, C. T. et al. (2003). The nature and identification of quantitative trait loci: a community's view. *Nature Reviews Genetics*, 4(11):911.

[55] Copley, S. D., Smith, E., and Morowitz, H. J. (2007). The origin of the RNA world: co-evolution of genes and metabolism. *Bioorganic Chemistry*, 35(6):430–443.

[56] Coulson, A. F. and Moult, J. (2002). A unifold, mesofold, and superfold model of protein fold use. *Proteins: Structure, Function, and Bioinformatics*, 46(1):61–71.

[57] Cowperthwaite, M. C., Bull, J. J., and Meyers, L. A. (2005). Distributions of beneficial fitness effects in RNA. *Genetics*, 170(4):1449–1457.

[58] Cowperthwaite, M. C. and Meyers, L. A. (2007). How mutational networks shape evolution: lessons from RNA models. *Annual Review of Ecology, Evolution, and Systematics*, pages 203–230.

[59] Croft, W. (2008). Evolutionary linguistics. *Annual Review of Anthropology*, 37(1):219.

[60] Crombach, A. and Hogeweg, P. (2008). Evolution of evolvability in gene regulatory networks. *PLoS Computational Biology*, 4(7):e1000112.

[61] Cross, W. C., Graham, T. A., and Wright, N. A. (2016). New paradigms in clonal evolution: punctuated equilibrium in cancer. *The Journal of Pathology*, 240(2):126–136.

[62] Crow, J. F. and Kimura, M. (1965). Evolution in sexual and asexual populations. *The American Naturalist*, 99(909):439–450.

[63] Cushing, J. M. (1998). *An introduction to structured population dynamics*. SIAM.

[64] da Silva, J., Coetzer, M., Nedellec, R., Pastore, C., and Mosier, D. E. (2010). Fitness epistasis and constraints on adaptation in a human immunodeficiency virus type 1 protein region. *Genetics*, 185(1):293–303.

[65] Dalerba, P., Cho, R. W., and Clarke, M. F. (2007). Cancer stem cells: models and concepts. *Annual Review of Medicine*, 58:267–284.

[66] Danforth, B. N. (1999). Emergence dynamics and bet hedging in a desert bee, Perdita portalis. *Proceedings of the Royal Society of London B: Biological Sciences*, 266(1432):1985–1994.

[67] Darwin, C. (1859). *On the Origin of Species by Means of Natural Selection*. Murray.

[68] Dawkins, R. (1999). *The Extended Phenotype: The Long Reach of the Gene*. Oxford University Press.

[69] Dawson, C. C., Intapa, C., and Jabra-Rizk, M. A. (2011). "persisters": survival at the cellular level. *PLoS Pathogens*, 7(7):e1002121.

[70] de Jong, I. G., Haccou, P., and Kuipers, O. P. (2011). Bet hedging or not? a guide to proper classification of microbial survival strategies. *Bioessays*, 33(3):215–223.

[71] De Visser, J. A. G. and Krug, J. (2014). Empirical fitness landscapes and the predictability of evolution. *Nature Reviews Genetics*, 15(7):480–490.

[72] Dean, M., Fojo, T., and Bates, S. (2005). Tumour stem cells and drug resistance. *Nature Reviews Cancer*, 5(4):275–284.

[73] Dhawan, A., Nichol, D., Kinose, F., Abazeed, M. E., Marusyk, A., Haura, E. B., and Scott, J. G. (2017). Collateral sensitivity networks reveal evolutionary instability and novel treatment strategies in alk mutated non-small cell lung cancer. *Scientific Reports*, 7:1232.

[74] Dobzhansky, T. (2013). Nothing in biology makes sense except in the light of evolution. *The American Biology Teacher*, 75(2):87–91.

[75] Dodd, I. B., Micheelsen, M. A., Sneppen, K., and Thon, G. (2007). Theoretical analysis of epigenetic cell memory by nucleosome modification. *Cell*, 129(4):813–822.

[76] Donnelly, L. and Walton, G. (2015). 25 cancer drugs to be denied on NHS. *The Daily Telegraph - http://www.telegraph.co.uk*. Published: 05/01/2015 (Online), Accessed: 09/07/2017 (Online).

[77] Donnenberg, V. S. and Donnenberg, A. D. (2005). Multiple drug resistance in cancer revisited: the cancer stem cell hypothesis. *The Journal of Clinical Pharmacology*, 45(8):872–877.

[78] Eigen, M., McCaskill, J., and Schuster, P. (1988). Molecular quasi-species. *The Journal of Physical Chemistry*, 92(24):6881–6891.

[79] Eigen, M., McCaskill, J., and Schuster, P. (1989). The molecular quasi-species. *Adv. Chem. Phys*, 75:149–263.

[80] Einstein, A. (1934). On the method of theoretical physics. *Philosophy of Science*, 1(2):163–169.

[81] Elowitz, M. B., Levine, A. J., Siggia, E. D., and Swain, P. S. (2002). Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186.

[82] Enderling, H., Hlatky, L., and Hahnfeldt, P. (2013). Cancer stem cells: a minor cancer subpopulation that redefines global cancer features. *Breast*, 11(35):200.

[83] Engelman, J. A. and Settleman, J. (2008). Acquired resistance to tyrosine kinase inhibitors during cancer therapy. *Current Opinion in Genetics & Development*, 18(1):73–79.

[84] Enriquez-Navas, P. M., Kam, Y., Das, T., Hassan, S., Silva, A., Foroutan, P., Ruiz, E., Martinez, G., Minton, S., Gillies, R. J., et al. (2016). Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer. *Science Translational Medicine*, 8(327):327ra24–327ra24.

[85] Enriquez-Navas, P. M., Wojtkowiak, J. W., and Gatenby, R. A. (2015). Application of evolutionary principles to cancer therapy. *Cancer Research*, 75(22):4675–4680.

[86] Erban, R., Chapman, J., and Maini, P. (2007). A practical guide to stochastic simulations of reaction-diffusion processes. *arXiv preprint arXiv:0704.1908*.

[87] Errington, J. (2013). L-form bacteria, cell walls and the origins of life. *Open Biology*, 3(1):120143.

[88] Ferguson, L. R. and Pearson, A. E. (1996). The clinical use of mutagenic anticancer drugs. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis*, 355(1):1–12.

[89] Fisher, R. (1930). *The genetical theory of natural selection*. The Clarendon Press.

[90] Flyvbjerg, H. and Lautrup, B. (1992). Evolution in a rugged fitness landscape. *Physical Review A*, 46(10):6714.

[91] Fontana, W. (2002). Modelling 'evo-devo'with RNA. *BioEssays*, 24(12):1164–1177.

[92] Fontana, W. and Schuster, P. (1998a). Continuity in evolution: on the nature of transitions. *Science*, 280(5368):1451–1455.

[93] Fontana, W. and Schuster, P. (1998b). Shaping space: the possible and the attainable in RNA genotype–phenotype mapping. *Journal of Theoretical Biology*, 194(4):491–515.

[94] Fontana, W., Stadler, P. F., Bornberg-Bauer, E. G., Griesmacher, T., Hofacker, I. L., Tacker, M., Tarazona, P., Weinberger, E. D., and Schuster, P. (1993). RNA folding and combinatory landscapes. *Physical Review E*, 47(3):2083.

[95] Foo, J., Chmielecki, J., Pao, W., and Michor, F. (2012). Effects of pharmacokinetic processes and varied dosing schedules on the dynamics of acquired resistance to erlotinib in EGFR-mutant lung cancer. *Journal of Thoracic Oncology*, 7(10):1583–1593.

[96] Franke, J., Klözer, A., De Visser, J., and Krug, J. (2011). Evolutionary accessibility of mutational pathways. *PLoS Computational Biology*, 7(8):e1002134.

[97] Frei, E., Elias, A., Wheeler, C., Richardson, P., and Hryniuk, W. (1998). The relationship between high-dose treatment and combination chemotherapy: the concept of summation dose intensity. *Clinical Cancer Research*, 4(9):2027–2037.

[98] Fuqua, C., Parsek, M. R., and Greenberg, E. P. (2001). Regulation of gene expression by cell-to-cell communication: acyl-homoserine lactone quorum sensing. *Annual Review of Genetics*, 35(1):439–468.

[99] Gallie, J., Libby, E., Bertels, F., Remigi, P., Jendresen, C. B., Ferguson, G. C., Desprat, N., Buffing, M. F., Sauer, U., Beaumont, H. J., et al. (2015). Bistability in a metabolic network underpins the de novo evolution of colony switching in pseudomonas fluorescens. *PLoS Biology*, 13(3):e1002109.

[100] Gatenby, R. A. (2009). A change of strategy in the war on cancer. *Nature*, 459(7246):508–509.

[101] Gatenby, R. A., Brown, J., and Vincent, T. (2009a). Lessons from applied ecology: cancer control using an evolutionary double bind. *Cancer Research*, 69(19):7499–7502.

[102] Gatenby, R. A., Gillies, R. J., and Brown, J. S. (2011). Of cancer and cave fish. *Nature Reviews Cancer*, 11(4):237–238.

[103] Gatenby, R. A., Silva, A. S., Gillies, R. J., and Frieden, B. R. (2009b). Adaptive therapy. *Cancer Research*, 69(11):4894–4903.

[104] Gay, L., Baker, A.-M., and Graham, T. A. (2016). Tumour cell heterogeneity. *F1000Research*, 5.

[105] Gerlee, P. (2015). Directional variation in evolution: consequences for the fitness landscape metaphor. *bioRxiv*, page 015529.

[106] Gerlee, P. and Anderson, A. R. (2007a). An evolutionary hybrid cellular automaton model of solid tumour growth. *Journal of Theoretical Biology*, 246(4):583–603.

[107] Gerlee, P. and Anderson, A. R. (2007b). Stability analysis of a hybrid cellular automaton model of cell colony growth. *Physical Review E*, 75(5):051911.

[108] Gerlee, P. and Anderson, A. R. (2008). A hybrid cellular automaton model of clonal evolution in cancer: the emergence of the glycolytic phenotype. *Journal of Theoretical Biology*, 250(4):705–722.

[109] Gerlee, P. and Anderson, A. R. (2009). Modelling evolutionary cell behaviour using neural networks: application to tumour growth. *Biosystems*, 95(2):166–174.

[110] Gerlee, P., Kim, E., and Anderson, A. R. (2015). Bridging scales in cancer progression: Mapping genotype to phenotype using neural networks. In *Seminars in Cancer Biology*, volume 30, pages 30–41. Elsevier.

[111] Gerlinger, M., Rowan, A. J., Horswell, S., Larkin, J., Endesfelder, D., Gronroos, E., Martinez, P., Matthews, N., Stewart, A., Tarpey, P., Varela, I., Phillimore, B., Begum, S., McDonald, N. Q., Butler, A., Jones, D., Raine, K., Latimer, C., Santos, C. R., Nohadani, M., Eklund, A. C., Spencer-Dene, B., Clark, G., Pickering, L., Stamp, G., Gore, M., Szallasi, Z., Downward, J., Futreal, P. A., and Swanton, C. (2012). Intratumor heterogeneity and branched evolution revealed by multiregion sequencing. *New England Journal of Medicine*, 366(10):883–892.

[112] Gerlinger, M. and Swanton, C. (2010). How Darwinian models inform therapeutic failure initiated by clonal heterogeneity in cancer medicine. *British Journal of Cancer*, 103(8):1139–1143.

[113] Gillespie, D. T. (1977). Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361.

[114] Gillespie, D. T. (2000). The chemical Langevin equation. *The Journal of Chemical Physics*, 113(1):297–306.

[115] Gillespie, D. T. (2001). Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733.

[116] Gillespie, J. H. (1983). A simple stochastic gene substitution model. *Theoretical Population Biology*, 23(2):202–215.

[117] Gillespie, J. H. (1984). Molecular evolution over the mutational landscape. *Evolution*, 38(5):1116–1129.

[118] Gillespie, J. H. (1991). *The causes of molecular evolution*. Oxford University Press.

[119] Gillies, R. J., Verduzco, D., and Gatenby, R. A. (2012). Evolutionary dynamics of carcinogenesis and why targeted therapy does not work. *Nature Reviews Cancer*, 12(7):487–493.

[120] Good, B. H., McDonald, M. J., Barrick, J. E., Lenski, R. E., and Desai, M. M. (2017). The dynamics of molecular evolution over 60,000 generations. *Nature*, 551(7678):45.

[121] Gottesman, M. M. (2002). Mechanisms of cancer drug resistance. *Annual Review of Medicine*, 53(1):615–627.

[122] Goulart, C. P., Mahmudi, M., Crona, K. A., Jacobs, S. D., Kallmann, M., Hall, B. G., Greene, D. C., and Barlow, M. (2013). Designing antibiotic cycling strategies by determining and understanding local adaptive landscapes. *PloS One*, 8(2):e56040.

[123] Gould, S. J. (2002). *The structure of evolutionary theory*. Harvard University Press.

[124] Gould, S. J. and Eldredge, N. (1977). Punctuated equilibria: the tempo and mode of evolution reconsidered. *Paleobiology*, 3(2):115–151.

[125] Gould, S. J. and Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society of London B: Biological Sciences*, 205(1161):581–598.

[126] Graham, T. A. and Sottoriva, A. (2016). Measuring cancer evolution from the genome. *The Journal of Pathology*.

[127] Gravner, J., Pitman, D., and Gavrilets, S. (2007). Percolation on fitness landscapes: Effects of correlation, phenotype, and incompatibilities. *Journal of Theoretical Biology*, 248(4):627–645.

[128] Greaves, M. and Maley, C. C. (2012). Clonal evolution in cancer. *Nature*, 481(7381):306–313.

[129] Greenbury, S. F., Johnston, I. G., Louis, A. A., and Ahnert, S. E. (2014). A tractable genotype–phenotype map modelling the self-assembly of protein quaternary structure. *Journal of The Royal Society Interface*, 11(95):20140249.

[130] Greenbury, S. F., Schaper, S., Ahnert, S. E., and Louis, A. A. (2016). Genetic correlations greatly increase mutational robustness and can both reduce and enhance evolvability. *PLoS Computational Biology*, 12(3):e1004773.

[131] Grinstead, C. M. and Snell, J. L. (1998). *Introduction to Probability*. American Mathematical Soc.

[132] Grüner, W., Giegerich, R., Strothmann, D., Reidys, C., Weber, J., Hofacker, I. L., Stadler, P. F., and Schuster, P. (1996a). Analysis of RNA sequence structure maps by exhaustive enumeration i. neutral networks. *Monatshefte für Chemie/Chemical Monthly*, 127(4):355–374.

[133] Grüner, W., Giegerich, R., Strothmann, D., Reidys, C., Weber, J., Hofacker, I. L., Stadler, P. F., and Schuster, P. (1996b). Analysis of RNA sequence structure maps by exhaustive enumeration ii. structures of neutral networks and shape space covering. *Monatshefte für Chemie/Chemical Monthly*, 127(4):375–389.

[134] Gupta, P. B., Fillmore, C. M., Jiang, G., Shapira, S. D., Tao, K., Kuperwasser, C., and Lander, E. S. (2011). Stochastic state transitions give rise to phenotypic equilibrium in populations of cancer cells. *Cell*, 146(4):633–644.

[135] Hall, B. G., Acar, H., Nandipati, A., and Barlow, M. (2014). Growth rates made easy. *Molecular Biology and Evolution*, 31(1):232–238.

[136] Hall, D. W., Agan, M., and Pope, S. C. (2010). Fitness epistasis among 6 biosynthetic loci in the budding yeast Saccharomyces cerevisiae. *Journal of Heredity*, 101(suppl 1):S75–S84.

[137] Hamming, R. W. (1950). Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2):147–160.

[138] Hanahan, D. and Weinberg, R. A. (2000). The hallmarks of cancer. *Cell*, 100(1):57–70.

[139] Hanahan, D. and Weinberg, R. A. (2011). Hallmarks of cancer: the next generation. *cell*, 144(5):646–674.

[140] Harris, T. E. (2002). *The Theory of Branching Processes*. Courier Corporation.

[141] Hartl, D. L., Clark, A. G., and Clark, A. G. (1997). *Principles of population genetics*, volume 116. Sinauer associates Sunderland.

[142] Hartl, D. L. and Taubes, C. H. (1996). Compensatory nearly neutral mutations: selection without adaptation. *Journal of Theoretical Biology*, 182(3):303–309.

[143] Hartl, D. L. and Taubes, C. H. (1998). Towards a theory of evolutionary adaptation. *Genetica*, 102:525–533.

[144] Hasty, J., Pradines, J., Dolnik, M., and Collins, J. J. (2000). Noise-based switches and amplifiers for gene expression. *Proceedings of the National Academy of Sciences*, 97(5):2075–2080.

[145] Hata, A. N., Niederst, M. J., Archibald, H. L., Gomez-Caraballo, M., Siddiqui, F. M., Mulvey, H. E., Maruvka, Y. E., Ji, F., Hyo-eun, C. B., Radhakrishna, V. K., Siravegna, G., Hu, H., Raoof, S., Lockerman, E., Kalsy, A., Lee, D., Keating, C. L., Ruddy, D. A., Damon, L. J., Crystal, A. S., Costa, C., Piotrowska, Z., Bardelli, A., Iafrate, A. J., Sadreyev, R. I., Stegmeier, F., Getz, G., Sequist, L. V., Faber, A. C., and Engelman, J. A. (2016). Tumor cells can follow distinct evolutionary paths to become resistant to epidermal growth factor receptor inhibition. *Nature medicine*, 22(3):262–269.

[146] Hinkley, T., Martins, J., Chappey, C., Haddad, M., Stawiski, E., Whitcomb, J. M., Petropoulos, C. J., and Bonhoeffer, S. (2011). A systems analysis of mutational effects in hiv-1 protease and reverse transcriptase. *Nature Genetics*, 43(5):487–489.

[147] Hobert, O. (2008). Gene regulation by transcription factors and microRNAs. *Science*, 319(5871):1785–1786.

[148] Hofacker, I. L., Fontana, W., Stadler, P. F., Bonhoeffer, L. S., Tacker, M., and Schuster, P. (1994). Fast folding and comparison of RNA secondary structures. *Monatshefte für Chemie/Chemical Monthly*, 125(2):167–188.

[149] Holder, K. K. and Bull, J. (2001). Profiles of adaptation in two similar viruses. *Genetics*, 159(4):1393–1404.

[150] Hood, L., Heath, J. R., Phelps, M. E., and Lin, B. (2004). Systems biology and new technologies enable predictive and preventative medicine. *Science*, 306(5696):640–643.

[151] Hopper, K. R. (1999). Risk-spreading and bet-hedging in insect population biology. *Annual Review of Entomology*, 44(1):535–560.

[152] Horn, F. and Jackson, R. (1972). General mass action kinetics. *Archive for Rational Mechanics and Analysis*, 47(2):81–116.

[153] Hornberg, J. J., Bruggeman, F. J., Westerhoff, H. V., and Lankelma, J. (2006). Cancer: a systems biology disease. *Biosystems*, 83(2):81–90.

[154] Huang, S. (2009). Non-genetic heterogeneity of cells in development: more than just noise. *Development*, 136(23):3853–3862.

[155] Huang, S. (2012). The molecular and mathematical basis of Waddington's epigenetic landscape: A framework for post-Darwinian biology? *Bioessays*, 34(2):149–157.

[156] Huang, S., Eichler, G., Bar-Yam, Y., and Ingber, D. E. (2005). Cell fates as high-dimensional attractor states of a complex gene regulatory network. *Physical Review Letters*, 94(12):128701.

[157] Huether, A., Höpfner, M., Sutter, A. P., Schuppan, D., and Scherübl, H. (2005). Erlotinib induces cell cycle arrest and apoptosis in hepatocellular cancer cells and enhances chemosensitivity towards cytostatics. *Journal of hepatology*, 43(4):661–669.

[158] Husbands, P., Harvey, I., Cliff, D., and Miller, G. (1997). Artificial evolution: a new path for artificial intelligence? *Brain and Cognition*, 34(1):130–159.

[159] Huynen, M. A., Stadler, P. F., and Fontana, W. (1996). Smoothness within ruggedness: the role of neutrality in adaptation. *Proceedings of the National Academy of Sciences*, 93(1):397–401.

[160] Hyo-eun, C. B., Ruddy, D. A., Radhakrishna, V. K., Caushi, J. X., Zhao, R., Hims, M. M., Singh, A. P., Kao, I., Rakiec, D., Shaw, P., et al. (2015). Studying clonal dynamics in response to cancer therapy using high-complexity barcoding. *Nature Medicine*, 21(5):440–448.

[161] Imamovic, L. and Sommer, M. O. (2013). Use of collateral sensitivity networks to design drug cycling protocols that avoid resistance development. *Science Translational Medicine*, 5(204):204ra132–204ra132.

[162] Isaacs, F. J., Hasty, J., Cantor, C. R., and Collins, J. J. (2003). Prediction and measurement of an autoregulatory genetic module. *Proceedings of the National Academy of Sciences*, 100(13):7714–7719.

[163] Iwasa, Y., Michor, F., and Nowak, M. A. (2004). Stochastic tunnels in evolutionary dynamics. *Genetics*, 166(3):1571–9.

[164] Jablonka, E., Oborny, B., Molnar, I., Kisdi, E., Hofbauer, J., and Czaran, T. (1995). The adaptive advantage of phenotypic memory in changing environments. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 350(1332):133–141.

[165] Jensen, S. O. and Lyon, B. R. (2009). Genetics of antimicrobial resistance in Staphylococcus aureus. *Future Microbiology*, 4(5):565–582.

[166] Johannsen, W. (1911). The genotype concept of heredity. *The American Naturalist*, 45:129–159.

[167] Jones, E., Oliphant, T., Peterson, P., et al. (2001–). SciPy: Open source scientific tools for Python.

[168] Joyce, A. R. and Palsson, B. Ø. (2006). The model organism as a system: integrating 'omics' data sets. *Nature Reviews Molecular Cell Biology*, 7(3):198–210.

[169] Kaern, M., Elston, T. C., Blake, W. J., and Collins, J. J. (2005). Stochasticity in gene expression: from theories to phenotypes. *Nature Reviews Genetics*, 6(6):451–464.

[170] Kalluri, R. and Weinberg, R. A. (2009). The basics of epithelial-mesenchymal transition. *The Journal of Clinical Investigation*, 119(6):1420–1428.

[171] Kauffman, S. (1971). Gene regulation networks: A theory for their global structure and behaviors. *Current topics in developmental biology*, 6:145–182.

[172] Kauffman, S. and Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45.

[173] Kauffman, S. A. (1993). *The origins of order: Self organization and selection in evolution*. Oxford university press.

[174] Kauffman, S. A. and Weinberger, E. D. (1989). The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology*, 141(2):211–245.

[175] Kawecki, T. J., Lenski, R. E., Ebert, D., Hollis, B., Olivieri, I., and Whitlock, M. C. (2012). Experimental evolution. *Trends in ecology & evolution*, 27(10):547–560.

[176] Kaznatcheev, A. (2013). Complexity of evolutionary equilibria in static fitness landscapes. *arXiv preprint arXiv:1308.5094*.

[177] Keren, I., Kaldalu, N., Spoering, A., Wang, Y., and Lewis, K. (2004a). Persister cells and tolerance to antimicrobials. *FEMS Microbiology Letters*, 230(1):13–18.

[178] Keren, I., Shah, D., Spoering, A., Kaldalu, N., and Lewis, K. (2004b). Specialized persister cells and the mechanism of multidrug tolerance in Escherichia coli. *Journal of Bacteriology*, 186(24):8172–8180.

[179] Khan, A. I., Dinh, D. M., Schneider, D., Lenski, R. E., and Cooper, T. F. (2011). Negative epistasis between beneficial mutations in an evolving bacterial population. *Science*, 332(6034):1193–1196.

[180] Kim, E., Rebecca, V. W., Smalley, K. S., and Anderson, A. R. (2016). Phase i trials in melanoma: A framework to translate preclinical findings to the clinic. *European Journal of Cancer*, 67:213–222.

[181] Kimura, M. (1969). The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations. *Genetics*, 61(4):893.

[182] Kimura, M. (1980). A simple method for estimating evolutionary rates of base substitutions through comparative studies of nucleotide sequences. *Journal of Molecular Evolution*, 16(2):111–120.

[183] Kimura, M. (1984). *The Neutral Theory of Molecular Evolution*. Cambridge University Press.

[184] Komarova, N. L., Sengupta, A., and Nowak, M. A. (2003). Mutation–selection networks of cancer initiation: tumor suppressor genes and chromosomal instability. *Journal of Theoretical Biology*, 223(4):433–450.

[185] Koonin, E. V., Wolf, Y. I., and Karev, G. P. (2002). The structure of the protein universe and genome evolution. *Nature*, 420(6912):218–223.

[186] Kounios, L., Clune, J., Kouvaris, K., Wagner, G. P., Pavlicev, M., Weinreich, D. M., and Watson, R. A. (2016). Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes. *arXiv preprint arXiv:1612.05955*.

[187] Kouvaris, K., Clune, J., Kounios, L., Brede, M., and Watson, R. A. (2017). How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. *PLoS Computational Biology*, 13(4):e1005358.

[188] Kouyos, R. D., Leventhal, G. E., Hinkley, T., Haddad, M., Whitcomb, J. M., Petropoulos, C. J., and Bonhoeffer, S. (2012). Exploring the complexity of the HIV-1 fitness landscape. *PLoS Genetics*, 8(3):e1002551.

[189] Kreeger, P. K. and Lauffenburger, D. A. (2010). Cancer systems biology: a network modeling perspective. *Carcinogenesis*, 31(1):2–8.

[190] Kryazhimskiy, S. and Plotkin, J. B. (2008). The population genetics of dN/dS. *PLoS Genetics*, 4(12):e1000304.

[191] Kumar, S. (2005). Molecular clocks: four decades of evolution. *Nature Reviews Genetics*, 6(8):654–662.

[192] Kurata, T., Tamura, K., Kaneda, H., Nogami, T., Uejima, H., Asai, G. o., Nakagawa, K., and Fukuoka, M. (2004). Effect of re-treatment with gefitinib ('Iressa', ZD1839) after acquisition of resistance. *Annals of Oncology*, 15(1):173–174.

[193] Kussell, E. and Leibler, S. (2005). Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, 309(5743):2075–2078.

[194] Land, M. F. and Fernald, R. D. (1992). The evolution of eyes. *Annual review of neuroscience*, 15(1):1–29.

[195] Laubenbacher, R., Hinkelmann, F., and Oremland, M. (2013). Agent-based models and optimal control in biology: a discrete approach. In Robeva, R. and Hodge, T., editors, *Mathematical concepts and methods in modern biology*, pages 143–178. Academic Press, Boston.

[196] Lázár, V., Nagy, I., Spohn, R., Csörgő, B., Györkei, Á., Nyerges, Á., Horváth, B., Vörös, A., Busa-Fekete, R., Hrtyan, M., et al. (2014). Genome-wide analysis captures the determinants of the antibiotic cross-resistance interaction network. *Nature communications*, 5.

[197] Lázár, V., Singh, G. P., Spohn, R., Nagy, I., Horváth, B., Hrtyan, M., Busa-Fekete, R., Bogos, B., Méhi, O., Csörgő, B., et al. (2013). Bacterial evolution of antibiotic hypersensitivity. *Molecular systems biology*, 9(1):700.

[198] Lemaire, P., Smith, W. C., and Nishida, H. (2008). Ascidians and the plasticity of the chordate developmental program. *Current Biology*, 18(14):R620–R631.

[199] Lenski, R. E., Rose, M. R., Simpson, S. C., and Tadler, S. C. (1991). Long-term experimental evolution in escherichia coli. i. adaptation and divergence during 2,000 generations. *The American Naturalist*, 138(6):1315–1341.

[200] Leuthner, K. D. and Doern, G. V. (2013). Antimicrobial stewardship programs. *Journal of Clinical Microbiology*, 51(12):3916–3920.

[201] Lévi, F. (2001). Circadian chronotherapy for human cancers. *The lancet oncology*, 2(5):307–315.

[202] Lévi, F., Zidani, R., Misset, J.-L., et al. (1997). Randomised multicentre trial of chronotherapy with oxaliplatin, fluorouracil, and folinic acid in metastatic colorectal cancer. *The Lancet*, 350(9079):681–686.

[203] Levy, S. F., Ziv, N., and Siegal, M. L. (2012). Bet hedging in yeast by heterogeneous, age-correlated expression of a stress protectant. *PLoS Biology*, 10(5):952.

[204] Lewis, K. (2006). Persister cells, dormancy and infectious disease. *Nature Reviews Microbiology*, 5(1):48–56.

[205] Li, C.-K. and Schneider, H. (2002). Applications of Perron–Frobenius theory to population dynamics. *Journal of Mathematical Biology*, 44(5):450–462.

[206] Lito, P., Rosen, N., and Solit, D. B. (2013). Tumor adaptation and resistance to RAF inhibitors. *Nature Medicine*, 19(11):1401–1409.

[207] Lord, C. J. and Ashworth, A. (2013). Mechanisms of resistance to therapies targeting BRCA-mutant cancers. *Nature Medicine*, 19(11):1381–1388.

[208] Lorenz, R., Bernhart, S. H., Zu Siederdissen, C. H., Tafer, H., Flamm, C., Stadler, P. F., and Hofacker, I. L. (2011). ViennaRNA package 2.0. *Algorithms for Molecular Biology*, 6(1):1.

[209] Lozovsky, E. R., Chookajorn, T., Brown, K. M., Imwong, M., Shaw, P. J., Kamchonwongpaisan, S., Neafsey, D. E., Weinreich, D. M., and Hartl, D. L. (2009). Stepwise acquisition of pyrimethamine resistance in the malaria parasite. *Proceedings of the National Academy of Sciences*, 106(29):12025–12030.

[210] Luidalepp, H., Jõers, A., Kaldalu, N., and Tenson, T. (2011). Age of inoculum strongly influences persister frequency and can mask effects of mutations implicated in altered persistence. *Journal of Bacteriology*, 193(14):3598–3605.

[211] Lv, W., Xu, Y., Guo, Y., Yu, Z., Feng, G., Liu, P., Luan, M., Zhu, H., Liu, G., Zhang, M., et al. (2016). The drug target genes show higher evolutionary conservation than non-target genes. *Oncotarget*, 7(4):4961.

[212] Ma, J., Ratan, A., Raney, B. J., Suh, B. B., Miller, W., and Haussler, D. (2008). The infinite sites model of genome evolution. *Proceedings of the National Academy of Sciences*, 105(38):14254–14261.

[213] Macken, C. A., Hagan, P. S., and Perelson, A. S. (1991). Evolutionary walks on rugged landscapes. *SIAM Journal on Applied Mathematics*, 51(3):799–827.

[214] Macken, C. A. and Perelson, A. S. (1989). Protein evolution on rugged landscapes. *Proceedings of the National Academy of Sciences*, 86(16):6191–6195.

[215] Martincorena, I., Raine, K. M., Gerstung, M., Dawson, K. J., Haase, K., Van Loo, P., Davies, H., Stratton, M. R., and Campbell, P. J. (2017). Universal patterns of selection in cancer and somatic tissues. *bioRxiv*, page 132324.

[216] Marusyk, A., Almendro, V., and Polyak, K. (2012). Intra-tumour heterogeneity: a looking glass for cancer? *Nature Reviews Cancer*, 12(5):323–334.

[217] Marusyk, A. and Polyak, K. (2010). Tumor heterogeneity: causes and consequences. *Biochimica et Biophysica Acta (BBA)-Reviews on Cancer*, 1805(1):105–117.

[218] Mason, O. and Verwoerd, M. (2007). Graph theory and networks in biology. *Systems Biology, IET*, 1(2):89–119.

[219] Mathew, A. G., Cissell, R., and Liamthong, S. (2007). Antibiotic resistance in bacteria associated with food animals: a United States perspective of livestock production. *Foodborne Pathogens and Disease*, 4(2):115–133.

[220] Maxwell, P., Dachs, G., Gleadle, J., Nicholls, L., Harris, A., Stratford, I., Hankinson, O., Pugh, C., , and Ratcliffe, P. (1997). Hypoxia-inducible factor-1 modulates gene expression in solid tumors and influences both angiogenesis and tumor growth. *Proceedings of the National Academy of Sciences*, 94(15):8104–8109.

[221] McCaskill, J. S. (1990). The equilibrium partition function and base pair binding probabilities for RNA secondary structure. *Biopolymers*, 29(6-7):1105–1119.

[222] McFarland, C. D., Mirny, L. A., and Korolev, K. S. (2014). Tug-of-war between driver and passenger mutations in cancer and other adaptive processes. *Proceedings of the National Academy of Sciences*, 111(42):15138–15143.

[223] Meacham, C. E. and Morrison, S. J. (2013). Tumour heterogeneity and cancer cell plasticity. *Nature*, 501(7467):328–337.

[224] Mendel, G. (1866). Versuche über pflanzenhybriden. *Verhandlungen des naturforschenden Vereines in Brunn 4: 3*, 44.

[225] Merlo, L. M., Pepper, J. W., Reid, B. J., and Maley, C. C. (2006). Cancer as an evolutionary and ecological process. *Nature Reviews Cancer*, 6(12):924–935.

[226] Metz, J. A. J. (2008). Fitness. *Encyclopedia of Ecology*, 2:1599–1612.

[227] Mira, P. M., Crona, K., Greene, D., Meza, J. C., Sturmfels, B., and Barlow, M. (2015a). Correction: Rational design of antibiotic treatment plans: A treatment strategy for managing evolution and reversing resistance. *PLoS One*, 10(9).

[228] Mira, P. M., Crona, K., Greene, D., Meza, J. C., Sturmfels, B., and Barlow, M. (2015b). Rational design of antibiotic treatment plans: a treatment strategy for managing evolution and reversing resistance. *PloS one*, 10(5):e0122283.

[229] Moody, J., Cosgrove, S. E., Olmsted, R., Septimus, E., Aureden, K., Oriola, S., Patel, G. W., and Trivedi, K. K. (2012). Antimicrobial stewardship: a collaborative partnership between infection preventionists and healthcare epidemiologists. *Infection Control and Hospital Epidemiology*, 33(4):328–330.

[230] Moran, P. A. P. et al. (1962). *The statistical processes of evolutionary theory*. Clarendon Press; Oxford University Press.

[231] Mormont, M.-C. and Levi, F. (2003). Cancer chronotherapy: principles, applications, and perspectives. *Cancer*, 97(1):155–169.

[232] Müller, J., Hense, B., Fuchs, T., Utz, M., and Pötzsche, C. (2013). Bet-hedging in stochastically switching environments. *Journal of Theoretical Biology*, 336:144–157.

[233] Munck, C., Gumpert, H. K., Wallin, A. I. N., Wang, H. H., and Sommer, M. O. (2014). Prediction of resistance development against drug combinations by collateral responses to component drugs. *Science Translational Medicine*, 6(262):262ra156–262ra156.

[234] Nichol, D., Jeavons, P., Fletcher, A. G., Bonomo, R. A., Maini, P. K., Paul, J. L., Gatenby, R. A., Anderson, A. R., and Scott, J. G. (2015). Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS Computational Biology*, 11(9):e1004493.

[235] Nichol, D., Robertson-Tessi, M., Jeavons, P., and Anderson, A. R. (2016). Stochasticity in the genotype-phenotype map: Implications for the robustness and persistence of bet-hedging. *Genetics*, 204(4):1523–1539.

[236] Nichol, D., Rutter, J., Bryant, C., Jeavons, P., Anderson, A., Bonomo, R., and Scott, J. (2017). Collateral sensitivity is contingent on the repeatability of evolution. *bioRxiv*, page 185892.

[237] Nikaido, H. (2009). Multidrug resistance in bacteria. *Annual Review of Biochemistry*, 78:119.

[238] Noble, D. (2002). Modeling the heart–from genes to cells to the whole organ. *Science*, 295(5560):1678–1682.

[239] Noble, D. (2008). Genes and causation. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 366(1878):3001–3015.

[240] Nowak, M. A. (1992). What is a quasispecies? *Trends in Ecology & Evolution*, 7(4):118–121.

[241] Nowak, M. A., Komarova, N. L., Sengupta, A., Jallepalli, P. V., Shih, I.-M., Vogelstein, B., and Lengauer, C. (2002). The role of chromosomal instability in tumor initiation. *Proceedings of the National Academy of Sciences*, 99(25):16226–16231.

[242] Nowell, P. C. (1976). The clonal evolution of tumor cell populations. *Science*, 194(4260):23–28.

[243] Nyström, T. (2007). A bacterial kind of aging. *PLoS Genetics*, 3(12):e224.

[244] Orlando, D. A., Lin, C. Y., Bernard, A., Wang, J. Y., Socolar, J. E., Iversen, E. S., Hartemink, A. J., and Haase, S. B. (2008). Global control of cell-cycle transcription by coupled cdk and network oscillators. *Nature*, 453(7197):944–947.

[245] Orr, H. A. (1998). The population genetics of adaptation: the distribution of factors fixed during adaptive evolution. *Evolution*, 52(4):935–949.

[246] Orr, H. A. (1999). The evolutionary genetics of adaptation: a simulation study. *Genetical Research*, 74(03):207–214.

[247] Orr, H. A. (2000). Adaptation and the cost of complexity. *Evolution*, 54(1):13–20.

[248] Orr, H. A. (2005). The genetic theory of adaptation: a brief history. *Nature Reviews Genetics*, 6(2):119–127.

[249] Orr, H. A. (2009). Fitness and its role in evolutionary genetics. *Nature Reviews Genetics*, 10(8):531–539.

[250] Papoulis, A. and Pillai, S. U. (2002). *Probability, random variables, and stochastic processes*. Tata McGraw-Hill Education.

[251] Pardal, R., Clarke, M. F., and Morrison, S. J. (2003). Applying the principles of stem-cell biology to cancer. *Nature Reviews Cancer*, 3(12):895–902.

[252] Patti, G. J., Yanes, O., and Siuzdak, G. (2012). Innovation: Metabolomics: the apogee of the omics trilogy. *Nature Reviews Molecular Cell Biology*, 13(4):263–269.

[253] Pavlopoulos, G. A., Secrier, M., Moschopoulos, C. N., Soldatos, T. G., Kossida, S., Aerts, J., Schneider, R., and Bagos, P. G. (2011). Using graph theory to analyze biological networks. *BioData Mining*, 4(1):1.

[254] Peck, J. R., Barreau, G., and Heath, S. C. (1997). Imperfect genes, Fisherian mutation and the evolution of sex. *Genetics*, 145(4):1171–1199.

[255] Pena-Miller, R. and Beardmore, R. (2010). Antibiotic cycling versus mixing: the difficulty of using mathematical models to definitively quantify their relative merits. *Mathematical Biosciences and Engineering*, 7(4):923–933.

[256] Phillips, P. C. (2008). Epistasis—the essential role of gene interactions in the structure and evolution of genetic systems. *Nature reviews. Genetics*, 9(11):855.

[257] Pigliucci, M. (2001). *Phenotypic Plasticity: Beyond Nature and Nurture*. JHU Press.

[258] Pigliucci, M. (2010). Genotype–phenotype mapping and the end of the 'genes as blueprint' metaphor. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1540):557–566.

[259] Pisco, A. and Huang, S. (2015). Non-genetic cancer cell plasticity and therapy-induced stemness in tumour relapse:'what does not kill me strengthens me'. *British Journal of Cancer*, 112(11):1725.

[260] Poelwijk, F. J., Kiviet, D. J., Weinreich, D. M., and Tans, S. J. (2007). Empirical fitness landscapes reveal accessible evolutionary paths. *Nature*, 445(7126):383–386.

[261] Poelwijk, F. J., Tănase-Nicola, S., Kiviet, D. J., and Tans, S. J. (2011). Reciprocal sign epistasis is a necessary condition for multi-peaked fitness landscapes. *Journal of Theoretical Biology*, 272(1):141–144.

[262] Poon, A. and Otto, S. P. (2000). Compensating for our load of mutations: freezing the meltdown of small populations. *Evolution*, 54(5):1467–1479.

[263] Quintana, E., Shackleton, M., Foster, H. R., Fullen, D. R., Sabel, M. S., Johnson, T. M., and Morrison, S. J. (2010). Phenotypic heterogeneity among tumorigenic melanoma cells from patients that is reversible and not hierarchically organized. *Cancer Cell*, 18(5):510–523.

[264] Raj, A. and van Oudenaarden, A. (2008). Nature, nurture, or chance: stochastic gene expression and its consequences. *Cell*, 135(2):216–226.

[265] Ramirez, M., Rajaram, S., Steininger, R. J., Osipchuk, D., Roth, M. A., Morinishi, L. S., Evans, L., Ji, W., Hsu, C.-H., Thurley, K., et al. (2016). Diverse drug-resistance mechanisms can emerge from drug-tolerant cancer persister cells. *Nature Communications*, 7.

[266] Reya, T., Morrison, S. J., Clarke, M. F., and Weissman, I. L. (2001). Stem cells, cancer, and cancer stem cells. *Nature*, 414(6859):105–111.

[267] Rezvani, A. R. and Maloney, D. G. (2011). Rituximab resistance. *Best practice & research Clinical haematology*, 24(2):203–216.

[268] Ricci-Vitiani, L., Lombardi, D. G., Pilozzi, E., Biffoni, M., Todaro, M., Peschle, C., and De Maria, R. (2007). Identification and expansion of human colon-cancer-initiating cells. *Nature*, 445(7123):111–115.

[269] Rice, S. H. (1990). A geometric model for the evolution of development. *Journal of Theoretical Biology*, 143(3):319–342.

[270] Robbins, G. K., De Gruttola, V., Shafer, R. W., Smeaton, L. M., Snyder, S. W., Pettinelli, C., Dubé, M. P., Fischl, M. A., Pollard, R. B., Delapenha, R., et al. (2003). Comparison of sequential three-drug regimens as initial therapy for HIV-1 infection. *New England Journal of Medicine*, 349(24):2293–2303.

[271] Robertson-Tessi, M. and Anderson, A. (2015). Big Bang and context-driven collapse. *Nature Genetics*, 47(3):196–197.

[272] Robertson-Tessi, M., Gillies, R. J., Gatenby, R. A., and Anderson, A. R. (2015). Impact of metabolic heterogeneity on tumor growth, invasion, and treatment outcomes. *Cancer Research*, 75(8):1567–1579.

[273] Rockwell, S. (1980). In vivo-in vitro tumour cell lines: characteristics and limitations as models for human cancer. *The British journal of cancer (Supplement)*, 4:118.

[274] Rodriguez, B., Burrage, K., Gavaghan, D., Grau, V., Kohl, P., and Noble, D. (2010). The systems biology approach to drug development: application to toxicity assessment of cardiac drugs. *Clinical Pharmacology & Therapeutics*, 88(1):130–134.

[275] Rodriguez de Evgrafov, M., Gumpert, H., Munck, C., Thomsen, T. T., and Sommer, M. O. (2015). Collateral resistance and sensitivity modulate evolution of high-level resistance to drug combination treatment in staphylococcus aureus. *Molecular Biology and Evolution*, 32(5):1175–1185.

[276] Roninson, I. B. (2003). Tumor cell senescence in cancer treatment. *Cancer Research*, 63(11):2705–2715.

[277] Rosell, R., Carcereny, E., Gervais, R., Vergnenegre, A., Massuti, B., Felip, E., Palmero, R., Garcia-Gomez, R., Pallares, C., Sanchez, J. M., et al. (2012). Erlotinib versus standard chemotherapy as first-line treatment for European patients with advanced EGFR mutation-positive non-small-cell lung cancer (EURTAC): a multicentre, open-label, randomised phase 3 trial. *The Lancet Oncology*, 13(3):239–246.

[278] Rosen, R. (1991). *Life itself: a comprehensive inquiry into the nature, origin, and fabrication of life*. Columbia University Press.

[279] Russell, G. (2016). Antibiotic resistance in children with E coli urinary tract infection. *BMJ*, 352:i1399.

[280] Rutherford, S. L. (2003). Between genotype and phenotype: protein chaperones and evolvability. *Nature Reviews Genetics*, 4(4):263–274.

[281] Sampieri, K. and Fodde, R. (2012). Cancer stem cells and metastasis. In *Seminars in Cancer Biology*, volume 22, pages 187–193. Elsevier.

[282] Sanderson, B. J., Ferguson, L. R., and Denny, W. A. (1996). Mutagenic and carcinogenic properties of platinum-based anticancer drugs. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis*, 355(1):59–70.

[283] Sauer, U., Heinemann, M., and Zamboni, N. (2007). Getting closer to the whole picture. *Science*, 316(5824):550–551.

[284] Sawyers, C. (2004). Targeted cancer therapy. *Nature*, 432(7015):294–297.

[285] Schaper, S., Johnston, I. G., and Louis, A. A. (2011). Epistasis can lead to fragmented neutral spaces and contingency in evolution. *Proceedings of the Royal Society of London B: Biological Sciences*, page rspb20112183.

[286] Schaper, S. and Louis, A. A. (2014). The arrival of the frequent: How bias in genotype-phenotype maps can steer populations to local optima. *PLoS One*, 9(2):e86635.

[287] Schenk, M. F., Szendro, I. G., Salverda, M. L., Krug, J., and de Visser, J. A. G. (2013). Patterns of epistasis between beneficial mutations in an antibiotic resistance gene. *Molecular Biology and Evolution*, 30(8):1779–1787.

[288] Schrödinger, E. (1992). *What is life?: With mind and matter and autobiographical sketches*. Cambridge University Press.

[289] Schuster, P. (1997). Genotypes with phenotypes: Adventures in an RNA toy world. *Biophysical Chemistry*, 66(2):75–110.

[290] Schuster, P., Fontana, W., Stadler, P. F., and Hofacker, I. L. (1994). From sequences to shapes and back: a case study in rna secondary structures. *Proceedings of the Royal Society of London B: Biological Sciences*, 255(1344):279–284.

[291] Scott, J. (2012). Phase i trialist. *Lancet Oncology*, 13(3):236.

[292] Seger, J. (1987). What is bet-hedging? *Oxford Surveys in Evolutionary Biology*, 4:182–211.

[293] Segovia-Juarez, J. L., Ganguli, S., and Kirschner, D. (2004). Identifying control mechanisms of granuloma formation during M. tuberculosis infection using an agent-based model. *Journal of Theoretical Biology*, 231(3):357–376.

[294] Seidman, A. D. (2005). Current status of dose-dense chemotherapy for breast cancer. *Cancer Chemotherapy and Pharmacology*, 56(1):78–83.

[295] Sella, G. and Hirsh, A. E. (2005). The application of statistical physics to evolutionary biology. *Proceedings of the National Academy of Sciences*, 102(27):9541–9546.

[296] Sharma, S. V., Lee, D. Y., Li, B., Quinlan, M. P., Takahashi, F., Maheswaran, S., McDermott, U., Azizian, N., Zou, L., Fischbach, M. A., et al. (2010). A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell*, 141(1):69–80.

[297] Sheen, J. (1994). Feedback control of gene expression. *Photosynthesis Research*, 39(3):427–438.

[298] Shendure, J. and Ji, H. (2008). Next-generation dna sequencing. *Nature biotechnology*, 26(10):1135–1145.

[299] Shibata, D. (2011). Mutation and epigenetic molecular clocks in cancer. *Carcinogenesis*, 32(2):123–128.

[300] Shipitsin, M., Campbell, L. L., Argani, P., Weremowicz, S., Bloushtain-Qimron, N., Yao, J., Nikolskaya, T., Serebryiskaya, T., Beroukhim, R., Hu, M., et al. (2007). Molecular definition of breast tumor heterogeneity. *Cancer cell*, 11(3):259–273.

[301] Shlaes, D. M., Gerding, D. N., John, J. F., Craig, W. A., Bornstein, D. L., Duncan, R. A., Eckman, M. R., Farrer, W. E., Greene, W. H., Lorian, V., et al. (1997). Society for Healthcare Epidemiology of America and Infectious Diseases Society of America Joint Committee on the Prevention of Antimicrobial Resistance: Guidelines for the Prevention of Antimicrobial Resistance in Hospitals. *Clinical Infectious Diseases*, 25(3):584–599.

[302] Shoval, O., Sheftel, H., Shinar, G., Hart, Y., Ramote, O., Mayo, A., Dekel, E., Kavanagh, K., and Alon, U. (2012). Evolutionary trade-offs, Pareto optimality, and the geometry of phenotype space. *Science*, 336(6085):1157–1160.

[303] Sigmund, K. (1994). *Games of Life*. Oxford University Press.

[304] Silva, A. S., Kam, Y., Khin, Z. P., Minton, S. E., Gillies, R. J., and Gatenby, R. A. (2012). Evolutionary approaches to prolong progression-free survival in breast cancer. *Cancer Research*, 72(24):6362–6370.

[305] Silver, L. L. (2011). Challenges of antibacterial discovery. *Clinical Microbiology Reviews*, 24(1):71–109.

[306] Smith, J. M. (1970). Natural selection and the concept of a protein space. *Nature*, 225(5232):563–564.

[307] Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge university press.

[308] Sniegowski, P. D. and Gerrish, P. J. (2010). Beneficial mutations and the dynamics of adaptation in asexual populations. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 365(1544):1255–1263.

[309] Sottoriva, A., Kang, H., Ma, Z., Graham, T. A., Salomon, M. P., Zhao, J., Marjoram, P., Siegmund, K., Press, M. F., Shibata, D., et al. (2015). A Big Bang model of human colorectal tumor growth. *Nature Genetics*, 47(3):209–216.

[310] Stadler, B. M., Stadler, P. F., Wagner, G. P., and Fontana, W. (2001). The topology of the possible: Formal spaces underlying patterns of evolutionary change. *Journal of Theoretical Biology*, 213(2):241–274.

[311] Stadler, P. F. (2006). Genotype-phenotype maps. *Biological Theory*, 2.

[312] Stern, D. L. and Orgogozo, V. (2009). Is genetic evolution predictable? *Science*, 323(5915):746–751.

[313] Stewart, E. J., Madden, R., Paul, G., and Taddei, F. (2005). Aging and death in an organism that reproduces by morphologically symmetric division. *PLoS Biology*, 3(2):e45.

[314] Strebhardt, K. and Ullrich, A. (2008). Paul ehrlich's magic bullet concept: 100 years of progress. *Nature Reviews Cancer*, 8(6):473–480.

[315] Sunnåker, M., Busetto, A. G., Numminen, E., Corander, J., Foll, M., and Dessimoz, C. (2013). Approximate Bayesian computation. *PLoS Computational Biology*, 9(1):e1002803.

[316] Szendro, I. G., Schenk, M. F., Franke, J., Krug, J., and De Visser, J. A. G. (2013). Quantitative analyses of empirical fitness landscapes. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01):P01005.

[317] Tan, B. T., Park, C. Y., Ailles, L. E., and Weissman, I. L. (2006). The cancer stem cell hypothesis: a work in progress. *Laboratory Investigation*, 86(12):1203–1207.

[318] Tan, L., Serene, S., Chao, H. X., and Gore, J. (2011). Hidden randomness between fitness landscapes limits reverse evolution. *Physical Review Letters*, 106(19):198102.

[319] Tenover, F. C. (2006). Mechanisms of antimicrobial resistance in bacteria. *The American Journal of Medicine*, 119(6):S3–S10.

[320] Thattai, M. and Van Oudenaarden, A. (2004). Stochastic gene expression in fluctuating environments. *Genetics*, 167(1):523–530.

[321] Thompson, E. W. and Newgreen, D. F. (2005). Carcinoma invasion and metastasis: a role for epithelial-mesenchymal transition? *Cancer Research*, 65(14):5991–5995.

[322] Tonegawa, S. (1983). Somatic generation of antibody diversity. *Nature*, 302(5909):575–581.

[323] Toprak, E., Veres, A., Michel, J.-B., Chait, R., Hartl, D. L., and Kishony, R. (2012). Evolutionary paths to antibiotic resistance under dynamically sustained drug selection. *Nature genetics*, 44(1):101–105.

[324] Trewavas, A. (2006). A brief history of systems biology "every object that biology studies is a system of systems." francois jacob (1974). *The Plant Cell*, 18(10):2420–2430.

[325] Vaira, D., Zullo, A., Vakil, N., Gatta, L., Ricci, C., Perna, F., Hassan, C., Bernabucci, V., Tampieri, A., and Morini, S. (2007). Sequential therapy versus standard triple-drug therapy for Helicobacter pylori eradication: a randomized trial. *Annals of Internal Medicine*, 146(8):556–563.

[326] Van Nimwegen, E., Crutchfield, J. P., and Huynen, M. (1999). Neutral evolution of mutational robustness. *Proceedings of the National Academy of Sciences*, 96(17):9716–9720.

[327] Veening, J.-W., Smits, W. K., and Kuipers, O. P. (2008). Bistability, epigenetics, and bet-hedging in bacteria. *Annual Review of Microbiology*, 62:193–210.

[328] Via, S. and Lande, R. (1985). Genotype-environment interaction and the evolution of phenotypic plasticity. *Evolution*, pages 505–522.

[329] Vogelstein, B. and Kinzler, K. W. (2004). Cancer genes and the pathways they control. *Nature Medicine*, 10(8):789–799.

[330] Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150(3811):563–565.

[331] Wagner, A. (1994). Evolution of gene networks by gene duplications: a mathematical model and its implications on genome organization. *Proceedings of the National Academy of Sciences*, 91(10):4387–4391.

[332] Wagner, A. (1996). Does evolutionary plasticity evolve? *Evolution*, 50(3):1008–1023.

[333] Wagner, A. (2008). Robustness and evolvability: a paradox resolved. *Proceedings of the Royal Society of London B: Biological Sciences*, 275(1630):91–100.

[334] Wagner, A. (2013). *Robustness and evolvability in living systems*. Princeton University Press.

[335] Wagner, G. P. (1988). The influence of variation and of developmental constraints on the rate of multivariate phenotypic evolution. *Journal of Evolutionary Biology*, 1(1):45–66.

[336] Wagner, G. P. and Altenberg, L. (1996). Perspective: complex adaptations and the evolution of evolvability. *Evolution*, pages 967–976.

[337] Wagner, G. P., Kenney-Hunt, J. P., Pavlicev, M., Peck, J. R., Waxman, D., and Cheverud, J. M. (2008). Pleiotropic scaling of gene effects and the 'cost of complexity'. *Nature*, 452(7186):470–472.

[338] Wang, P., Robert, L., Pelletier, J., Dang, W. L., Taddei, F., Wright, A., and Jun, S. (2010). Robust growth of Escherichia coli. *Current Biology*, 20(12):1099–1103.

[339] Watson, J. D., Crick, F. H., et al. (1953). Molecular structure of nucleic acids. *Nature*, 171(4356):737–738.

[340] Watson, N. (1988). A new revision of the sequence of plasmid pBR322. *Gene*, 70(2):399–403.

[341] Watson, R. A. and Szathmáry, E. (2016). How can evolution learn? *Trends in ecology & evolution*, 31(2):147–157.

[342] Weghorn, D. and Sunyaev, S. (2017). Bayesian inference of negative and positive selection in human cancers. *Nature genetics*, 49(12):1785.

[343] Weinberg, R. (2013). *The biology of cancer*. Garland Science.

[344] Weinberger, E. (1990). Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63(5):325–336.

[345] Weinberger, E. D. (1991). Local properties of kauffman's N-k model: A tunably rugged energy landscape. *Physical Review A*, 44(10):6399.

[346] Weinreich, D. M., Delaney, N. F., DePristo, M. A., and Hartl, D. L. (2006). Darwinian evolution can follow only very few mutational paths to fitter proteins. *Science*, 312(5770):111–114.

[347] Weinreich, D. M., Lan, Y., Jaffe, J., and Heckendorn, R. B. (2017). The influence of higher-order epistasis on biological fitness landscape topography. *bioRxiv*, page 164798.

[348] Weinreich, D. M., Lan, Y., Wylie, C. S., and Heckendorn, R. B. (2013). Should evolutionary geneticists worry about higher-order epistasis? *Current Opinion in Genetics & Development*, 23(6):700–707.

[349] Weinreich, D. M., Watson, R. A., and Chao, L. (2005). Perspective: sign epistasis and genetic costraint on evolutionary trajectories. *Evolution*, 59(6):1165–1174.

[350] Weinstein, J. N., Collisson, E. A., Mills, G. B., Shaw, K. R. M., Ozenberger, B. A., Ellrott, K., Shmulevich, I., Sander, C., Stuart, J. M., Network, C. G. A. R., et al. (2013). The cancer genome atlas pan-cancer analysis project. *Nature Genetics*, 45(10):1113–1120.

[351] Weissman, D. B., Desai, M. M., Fisher, D. S., and Feldman, M. W. (2009). The rate at which asexual populations cross fitness valleys. *Theoretical Population Biology*, 75(4):286–300.

[352] Wilke, C. O. (2003). Probability of fixation of an advantageous mutant in a viral quasispecies. *Genetics*, 163(2):467–474.

[353] Wilke, C. O. (2005). Quasispecies theory in the context of population genetics. *BMC Evolutionary Biology*, 5(1):44.

[354] Williams, M. J., Werner, B., Barnes, C. P., Graham, T. A., and Sottoriva, A. (2016a). Identification of neutral tumor evolution across cancer types. *Nature Genetics*.

[355] Williams, M. J., Werner, B., Curtis, C., Barnes, C., Sottoriva, A., and Graham, T. A. (2016b). Quantification of subclonal selection in cancer from bulk sequencing data. *bioRxiv*, page 096305.

[356] Wolf, D. M., Vazirani, V. V., and Arkin, A. P. (2005). Diversity in times of adversity: probabilistic strategies in microbial survival games. *Journal of Theoretical Biology*, 234(2):227–253.

[357] Wood, T. K., Knabel, S. J., and Kwan, B. W. (2013). Bacterial persister cell formation and dormancy. *Applied and Environmental Microbiology*, 79(23):7116–7121.

[358] Woodford, N. and Ellington, M. J. (2007). The emergence of antibiotic resistance by mutation. *Clinical Microbiology and Infection*, 13(1):5–18.

[359] Wright, S. (1932). The Roles of Mutation, Inbreeding, Crossbreeding, and Selection in Evolution. *Proceeding of the 6th International Congress of Genetics*, 1:356–366.

[360] Wright, S. (1967). “surfaces" of selective value. *Proceedings of the National Academy of Sciences*, 58(1):165.

[361] Wuchty, S., Fontana, W., Hofacker, I. L., Schuster, P., et al. (1999). Complete suboptimal folding of RNA and the stability of secondary structures. *Biopolymers*, 49(2):145–165.

[362] Yang, J. and Weinberg, R. A. (2008). Epithelial-mesenchymal transition: at the crossroads of development and tumor metastasis. *Developmental Cell*, 14(6):818–829.

[363] Yano, S., Nakataki, E., Ohtsuka, S., Inayama, M., Tomimoto, H., Edakuni, N., Kakiuchi, S., Nishikubo, N., Muguruma, H., and Sone, S. (2005). Retreatment of lung adenocarcinoma patients with gefitinib who had experienced favorable results from their initial treatment with this selective epidermal growth factor receptor inhibitor: a report of three cases. *Oncology Research Featuring Preclinical and Clinical Cancer Therapeutics*, 15(2):107–111.

[364] Yap, T. A., Gerlinger, M., Futreal, P. A., Pusztai, L., and Swanton, C. (2012). Intratumor heterogeneity: seeing the wood for the trees. *Science Translational Medicine*, 4(127):127ps10–127ps10.

[365] Zhang, Y., Yew, W. W., and Barer, M. R. (2012). Targeting persisters for tuberculosis control. *Antimicrobial Agents and Chemotherapy*, 56(5):2223–2230.

[366] Zhao, B., Sedlak, J. C., Srinivas, R., Creixell, P., Pritchard, J. R., Tidor, B., Lauffenburger, D. A., and Hemann, M. T. (2016). Exploiting temporal collateral sensitivity in tumor clonal evolution. *Cell*, 165(1):234–246.

[367] Zuker, M. and Sankoff, D. (1984). RNA secondary structures and their prediction. *Bulletin of Mathematical Biology*, 46(4):591–621.

[368] zur Wiesch, P. A., Kouyos, R., Abel, S., Viechtbauer, W., and Bonhoeffer, S. (2014). Cycling empirical antibiotic therapy in hospitals: meta-analysis and models. *PLoS Pathogens*, 10(6):e1004225.

Appendix A

Additional Figures

A.1 Scatterplots of Fitness Between Landscapes

Here, we present the scatter plots used to determine the correlation coefficients in Table 3.2.

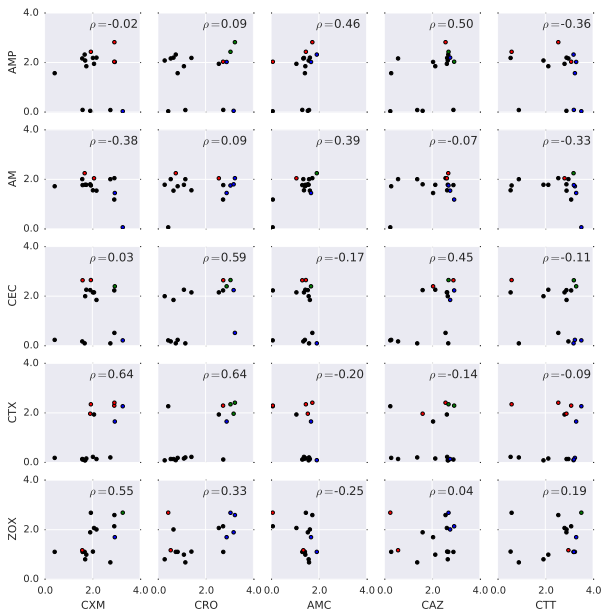

Fig. A.1

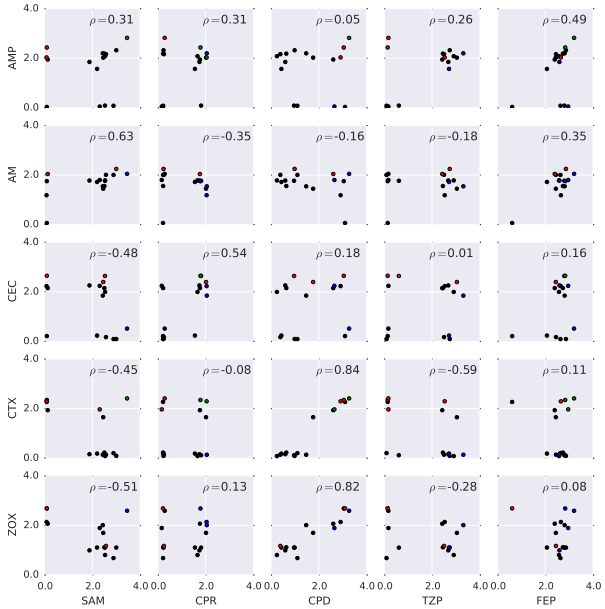
AMP
AM
CEC
CTX
ZOX
SAM
CPR
CPD
TZP
FEP
ρ=0.31 ρ=0.31 ρ=0.05 ρ=0.26 ρ=0.49
ρ=0.63 ρ=-0.35 ρ=-0.16 ρ=-0.18 ρ=0.35
ρ=-0.48 ρ=0.54 ρ=0.18 ρ=0.01 ρ=0.16
ρ=-0.45 ρ=-0.08 ρ=0.84 ρ=-0.59 ρ=0.11
ρ=-0.51 ρ=0.13 ρ=0.82 ρ=-0.28 ρ=0.08

Fig. A.2

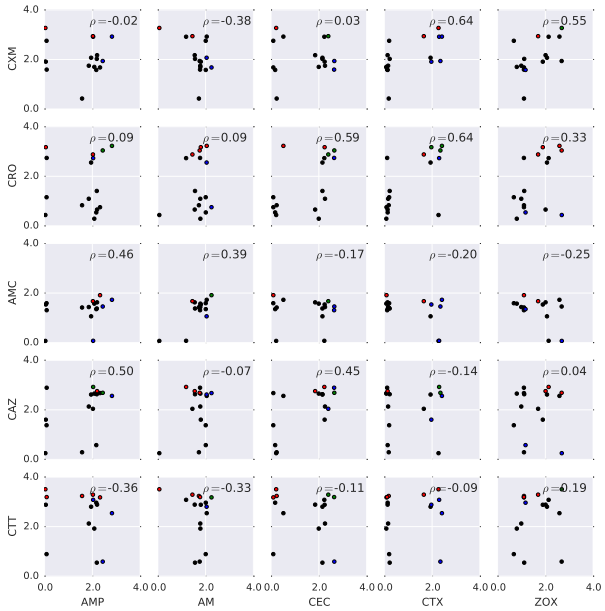
CXM
CRO
AMC
CAZ
CTT
AMP
AM
CEC
CTX
ZOX
ρ=-0.02 ρ=-0.38 ρ=0.03 ρ=0.64 ρ=0.55
ρ=0.09 ρ=0.09 ρ=0.59 ρ=0.64 ρ=0.33
ρ=0.46 ρ=0.39 ρ=-0.17 ρ=-0.20 ρ=-0.25
ρ=0.50 ρ=-0.07 ρ=0.45 ρ=-0.14 ρ=0.04
ρ=-0.36 ρ=-0.33 ρ=-0.11 ρ=-0.09 ρ=0.19

Fig. A.3

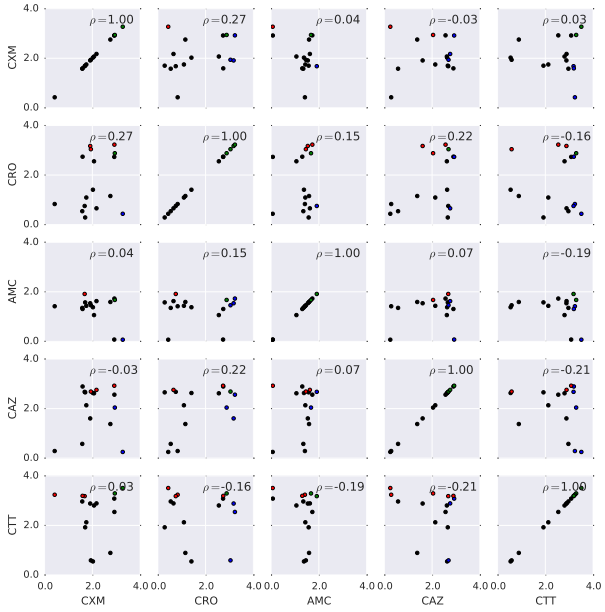
ρ=1.00 ρ=0.27 ρ=0.04 ρ=-0.03 ρ=0.03
ρ=0.27 ρ=1.00 ρ=0.15 ρ=0.22 ρ=-0.16
ρ=0.04 ρ=0.15 ρ=1.00 ρ=0.07 ρ=-0.19
ρ=-0.03 ρ=0.22 ρ=0.07 ρ=1.00 ρ=-0.21
ρ=0.03 ρ=-0.16 ρ=-0.19 ρ=-0.21 ρ=1.00
CXM
CRO
AMC
CAZ
CTT
0.0
2.0
4.0

Fig. A.4

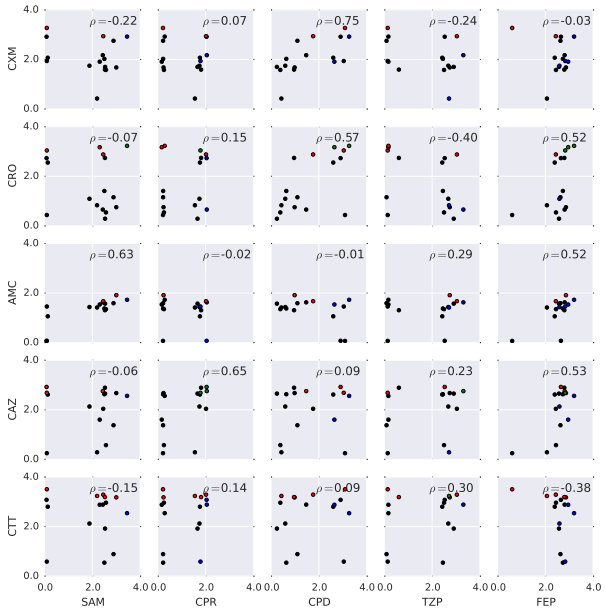
ρ=-0.22 ρ=0.07 ρ=0.75 ρ=-0.24 ρ=-0.03
ρ=-0.07 ρ=0.15 ρ=0.57 ρ=-0.40 ρ=0.52
ρ=0.63 ρ=-0.02 ρ=-0.01 ρ=0.29 ρ=0.52
ρ=-0.06 ρ=0.65 ρ=0.09 ρ=0.23 ρ=0.53
ρ=-0.15 ρ=0.14 ρ=0.09 ρ=0.30 ρ=-0.38
CXM
CRO
AMC
CAZ
CTT
SAM
CPR
CPD
TZP
FEP
0.0
2.0
4.0

Fig. A.5

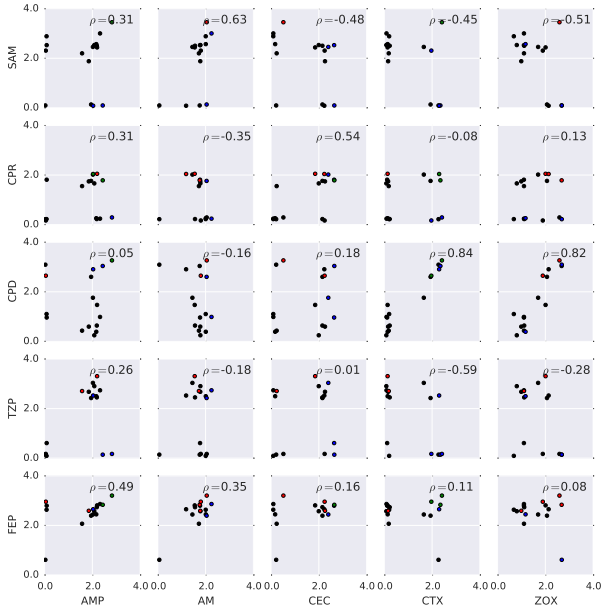
SAM
CPR
CPD
TZP
FEP
AMP
AM
CEC
CTX
ZOX
ρ=0.31
ρ=0.63
ρ=-0.48
ρ=-0.45
ρ=-0.51
ρ=0.31
ρ=-0.35
ρ=0.54
ρ=-0.08
ρ=0.13
ρ=0.05
ρ=-0.16
ρ=0.18
ρ=0.84
ρ=0.82
ρ=0.26
ρ=-0.18
ρ=0.01
ρ=-0.59
ρ=-0.28
ρ=0.49
ρ=0.35
ρ=0.16
ρ=0.11
ρ=0.08

Fig. A.6

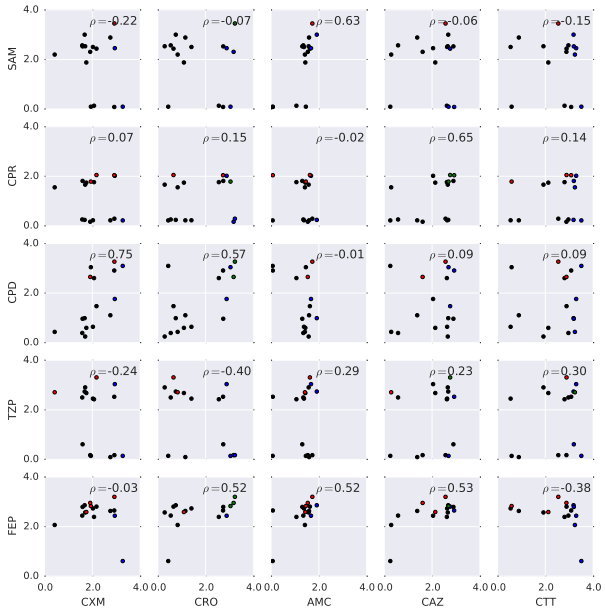
SAM
CPR
CPD
TZP
FEP
CXM
CRO
AMC
CAZ
CTT
ρ=-0.22
ρ=-0.07
ρ=0.63
ρ=-0.06
ρ=-0.15
ρ=0.07
ρ=0.15
ρ=-0.02
ρ=0.65
ρ=0.14
ρ=0.75
ρ=0.57
ρ=-0.01
ρ=0.09
ρ=0.09
ρ=-0.24
ρ=-0.40
ρ=0.29
ρ=0.23
ρ=0.30
ρ=-0.03
ρ=0.52
ρ=0.52
ρ=0.53
ρ=-0.38

Fig. A.7

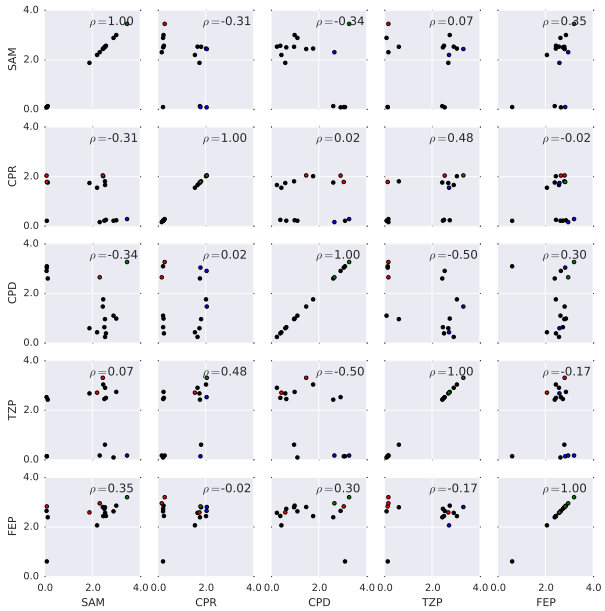
SAM
CPR
CPD
TZP
FEP
ρ=1.00
ρ=-0.31
ρ=-0.34
ρ=0.07
ρ=0.35
ρ=-0.31
ρ=1.00
ρ=0.02
ρ=0.48
ρ=-0.02
ρ=-0.34
ρ=0.02
ρ=1.00
ρ=-0.50
ρ=0.30
ρ=0.07
ρ=0.48
ρ=-0.50
ρ=1.00
ρ=-0.17
ρ=0.35
ρ=-0.02
ρ=0.30
ρ=-0.17
ρ=1.00
0.0
2.0
4.0

Fig. A.8

Appendix B

Complementary Analysis for Chapter 4 and Chapter 5

Deriving probabilities for the AM Switch

Explicit simulation of the Gillespie algorithm is computationally expensive. Instead, we derive the probability for switching for all initial conditions (x_0, y_0). For the DC, DCx and DCy switches, a closed form solution for these probabilities exist 4.4. For the AM switch, we determine the switching probabilities by constructing the Markov chain on the state space as follows.

Assume that there are $n = x_0 + y_0$ molecules in the system. Denote a configuration of the CRN by (r,s) for $r,s \in \mathbb{N}$ with $0 < r+s \leq n$ where r and s denote the number of molecules x and y respectively. Call a configuration stable if no more reactions can occur. As each reaction in the AM network conserves the total number of molecules we know that the number of b molecules is equal to $n-r-s$. Denote

$$p(r,s) = \mathbb{P}[\text{The AM–network with inital config. } (r,s) \text{ ends in the stable config. } (n,0)]. \tag{B.1}$$

In the absence of y molecules, either the system has reached the stable configuration $(n,0)$ or the only reaction that can occur is $x+b \longrightarrow x+x$ and the system must reach $(n,0)$ eventually. Hence

$$p(r,0) = 1 \text{ for all } r \in [1,n]\ . \tag{B.2}$$

Further by the symmetry of the reaction network we must have that

$$p(r,s) = 1 - p(s,r) \text{ for all } s,r \in [0,n] \text{ with } 0 < r+s \leq n \tag{B.3}$$

and in particular we have

$$p(r,s) = \frac{1}{2} \text{ for all } r = s \text{ with } 0 < r+s \leq n. \tag{B.4}$$

By this symmetry we need only determine $p(r,s)$ on the set

$$T = \{(r,s) : r,s \in [0,n], 0 < r+s \leq n \text{ and } r \geq s\} \tag{B.5}$$

Denote the subset of T for which we know $p(r,s)$ by T^*. Within $T \setminus T^*$ we have a recurrence relation for $p(r,s)$.

$$\begin{aligned} p(r,s) = &\mathbb{P}[x+y \longrightarrow x+b : r,s]p(r,s-1) \\ &+\mathbb{P}[y+x \longrightarrow y+b : r,s]p(r-1,s) \\ &+\mathbb{P}[x+b \longrightarrow x+x : r,s]p(r+1,s) \\ &+\mathbb{P}[y+b \longrightarrow y+y : r,s]p(r,s+1) \end{aligned} \tag{B.6}$$

Where the probabilities of each reaction are given by the propensity functions,

$$\begin{aligned} \mathbb{P}[x+y \longrightarrow x+b : r,s] &= \frac{rs}{2rs+r(n-r-s)+s(n-r-s)} \\ \mathbb{P}[y+x \longrightarrow y+b : r,s] &= \frac{rs}{2rs+r(n-r-s)+s(n-r-s)} \\ \mathbb{P}[x+b \longrightarrow x+x : r,s] &= \frac{r(n-r-s)}{2rs+r(n-r-s)+s(n-r-s)} \\ \mathbb{P}[y+b \longrightarrow y+y : r,s] &= \frac{s(n-r-s)}{2rs+r(n-r-s)+s(n-r-s)}. \end{aligned} \tag{B.7}$$

Now, as we know the values of $p(r,s)$ along two boundaries we can solve this recurrence numerically with a Markov chain. We can view the chemical reaction simulation as a random walk inside the configuration space T with transition probabilities given by the likelihoods of each reaction. We define the $|T| \times |T|$ transition matrix for a Markov chain as

$$M((r,s),(r',s')) = \begin{cases} 1 & \text{if } (r,s) \in Y^* \text{ and } r = r', s = s' \\ 0 & \text{if } (r,s) \in T^* \text{ and } r \neq r' \text{ or } s \neq s' \\ \mathbb{P}[R : r,s] & \text{if } (r,s) \notin T^* \text{ and there exists a reaction } R \text{ taking } (r,s) \text{ to } (r',s') \\ 0 & \text{otherwise} \end{cases} \tag{B.8}$$

in which we have taken those values (r,s) for which we know $p(r,s)$ to be absorbing states. If we define the limit matrix

$$M^* = \lim_{k \to \infty} M^k \tag{B.9}$$

then $M^*((r,s),(r',s'))$ is only non-zero when (r',s') is a state for which we know $p(r',s')$. In this case $M^*((r,s),(r',s'))$ denotes the probability that the random walk in S which simulates a series of chemical reactions will reach (r',s') before any other absorbing state. As such we can determine $p(r,s)$ by summing over all of the absorbing states of the Markov chain

$$p(r,s) = \sum_{\substack{r',s' \in [0,n], 0<r'+s' \leq n \\ s'=0 \text{ or } r'=s'}} p(r',s') M^*((r,s),(r',s')). \tag{B.10}$$

This is the technique used to determine the probability of phenotype A for an inidividual with genotype $g = (x_0, y_0)$ and to generate the heatmaps shown in Figure 4.2.

B.1 Parameter Robustness for the Results of Chapter 5

B.1.1 Fixed Environments

To explore the effects of parameter variation on the results of Chapter 5 we re–derived the genotype–growth rate curves for each memory mechanism and molecular switch. We held the parameters for phenotype B constant and for phenotype A we considered each pair of parameters sampled from $b \in \{0.1b_A, b_A, 5b_A, 10b_A\}$ and $d \in \{0.1d_A, d_A, 2d_A, 5d_A\}$ (subject to the condition that the growth rate of phenotype A is greater than one). These parameters were chosen to represent a broad range scenarios for the relationship between the fitness of persister and non–persister cells, ranging from non–persister exhibiting a weekly growth rate only $2.9\times$ greater than non–persisters at the lower extreme and $141.1\times$ greater than non–persisters at the upper extreme.

Figures B.1, B.2, B.3 and B.4 show the genotype-growth rate curves for the DC, DCx, DCy and AM switches respectively for each parameter set (for the drug free environment). These figures mirror Figure 5.8 in the main text. Note that parameter sets for which phenotype A does not exhibit growth rate greater than one are omitted (shaded red subplots). These figures clearly indicate fitness is dependent on each of the components of bet–hedging: the genotype, memory mechanism and molecular switch. This confirms the key result of Chapter 5 - that each of these components is critical to the determination of fitness and must be explicitly considered if we are to understand or predict the evolution of bet–hedging. Interestingly, we find that for a fixed genotype and molecular mechanism, the ordering of fitness amongst the memory mechanisms can change as the parameter set changes (for example, Figure B.1, the ordering of fitness between constrained genetic and epigenetic memory bet–hedging is switched in the second column). This observation indicates that the fitness landscapes outlined in Figures 5.9 and 5.10 are not necessarily invariant under changes to the parameter set. This does not invalidate our results, which state that epistasis and neutrality can be present in the landscapes, but do not claim to derive the precise fitness landscape underpinning the evolution of bet–hedging.

B.1.2 Fluctuating Environments

For fluctuating environments we explored the parameter range $b \in \{0.1b_A, b_A, 5b_A\}$ and $d \in \{0.1d_A, d_A, 2d_A\}$ (subject to the condition that the growth rate of phenotype A is greater than one). Mirroring Figure 5.11 we generated genotype-growth rate curves for each of the parameter sets outlined above in environments fluctuating between 14 weeks of drug exposure and 14 weeks of drug–free growth. The drug treated growth parameters for phenotype A are taken as $b^{\text{drug}} = b$, $d^{\text{drug}} = 8d$, as in Chapter 5.

Figures B.5, B.6, B.7 and B.8 shows these growth curves for the DC, DCx, DCy and AM switches respectively. Here, we see that bet–hedging is optimal for only certain parameter sets, whereas for others (the final column in each of the figures) the fitness optimum is attained for a monomorphic population of phenotype A. Following previously published theoretical work exploring bet–hedging [232] we predict that an optimum bet–hedging strategy is in fact not monomorphic, but that the associated likelihood of phenotype A is sufficiently close to 1 that is cannot be captured by our model with $g_{\max} = 60$. From these observations we conclude that our results will hold for a large space of parameter values encoding the relationship between persister and non–persister cells arising from bet–hedging, but that they are not universal.

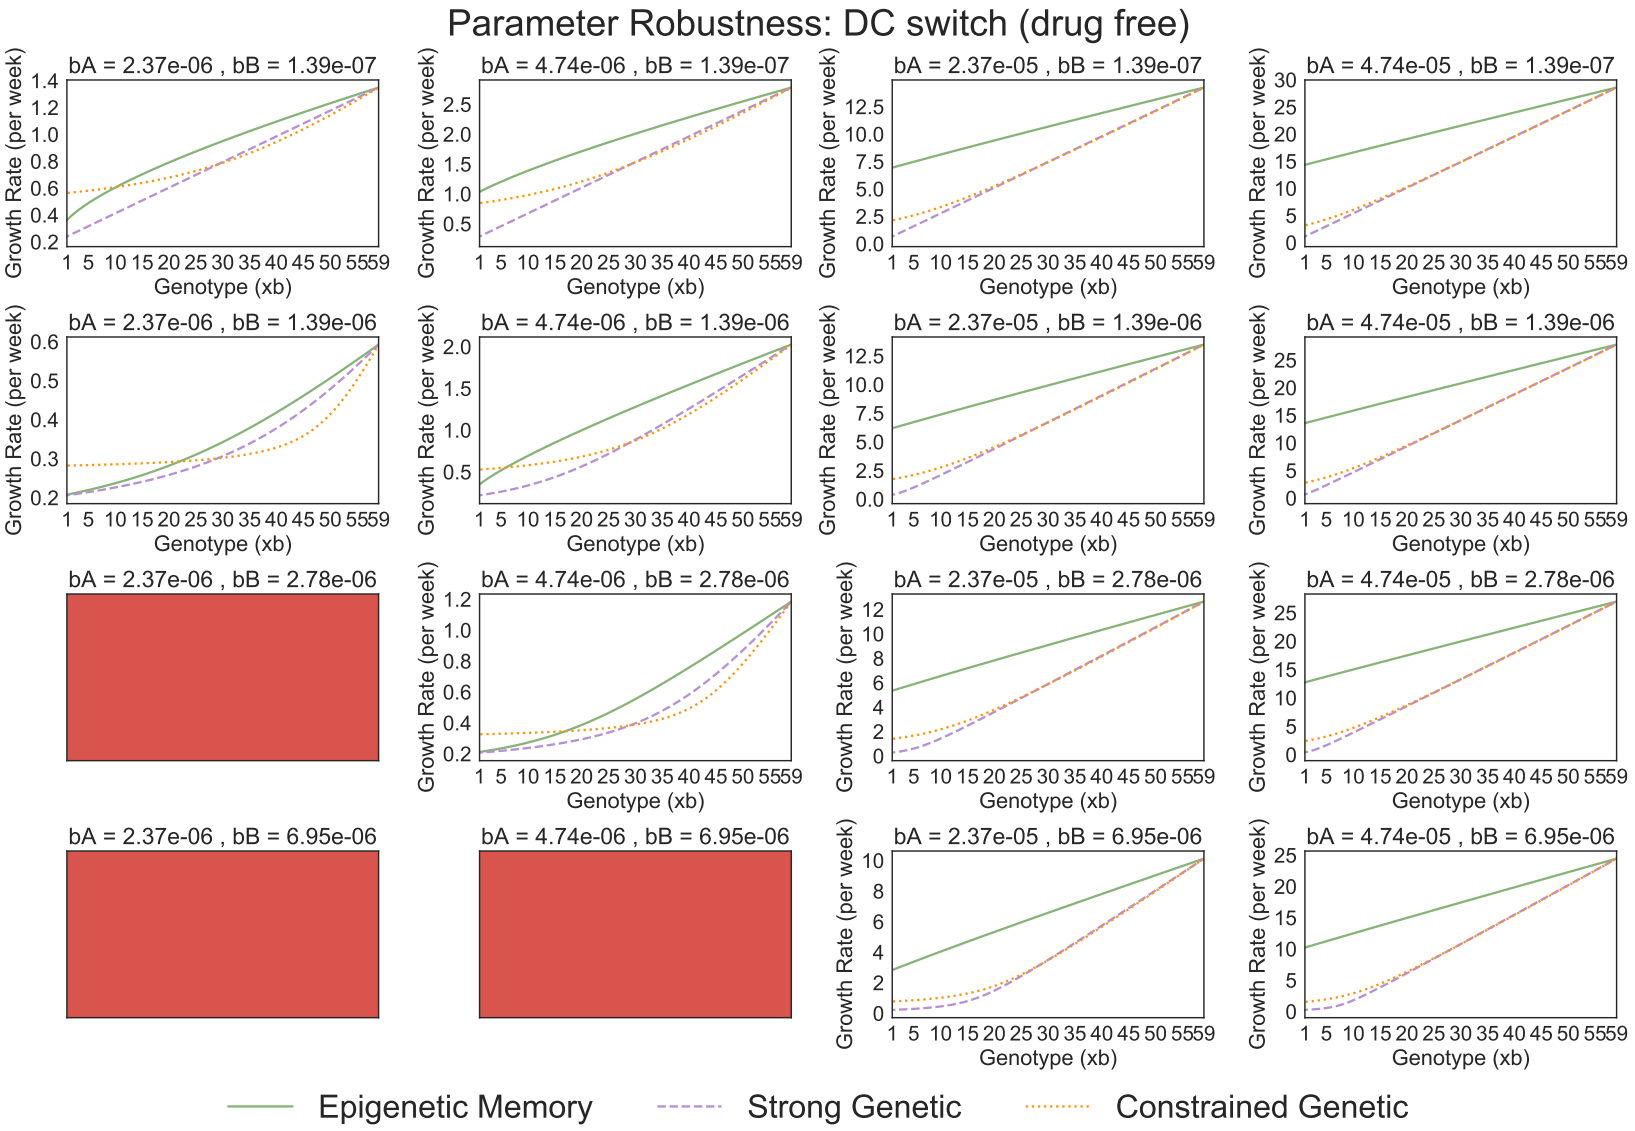

Fig. B.1 **Genotype-growth rate curves for varying parameter sets with the DC switch.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles). Subplots shaded red indicate that the phenotype *A* growth rate is not greater than 1. In all cases phenotype *B* has parameters as in the main text (Chapter 5). Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism.

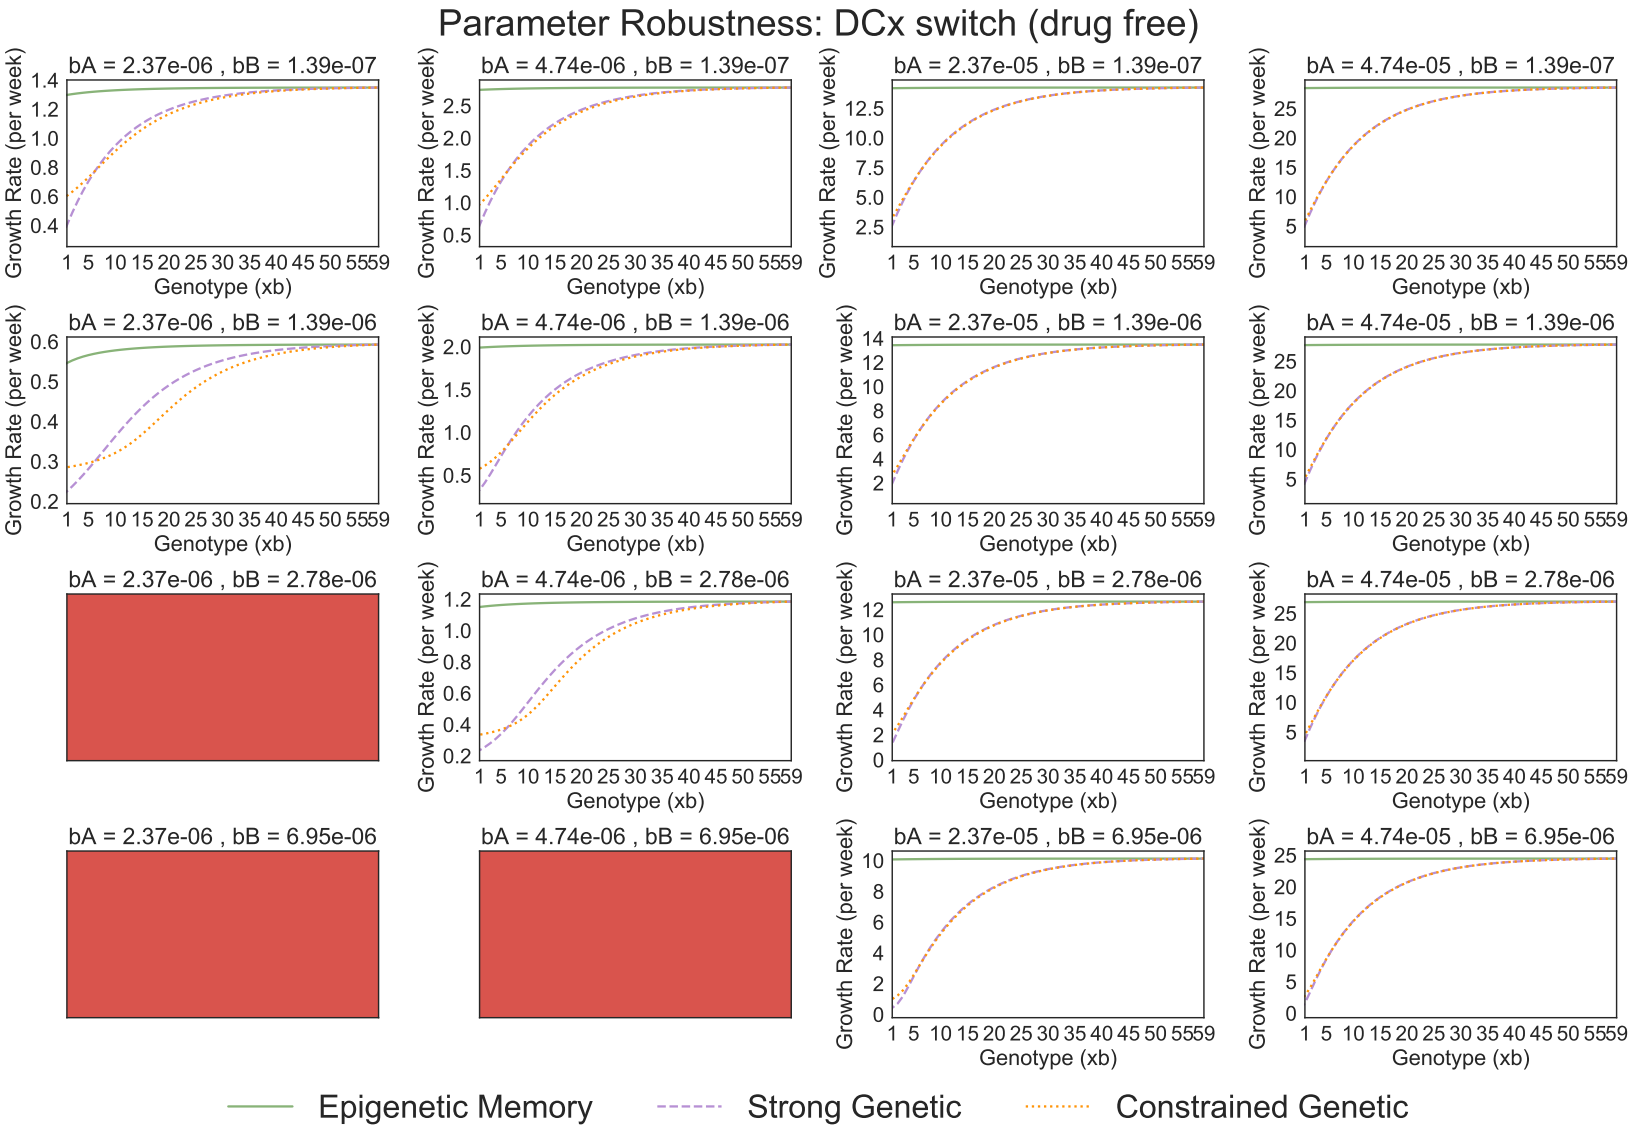

Fig. B.2 **Genotype-growth rate curves for varying parameter sets with the DCx switch.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles). Subplots shaded red indicate that the phenotype *A* growth rate is not greater than 1. In all cases phenotype *B* has parameters as in the main text (Chapter 5). Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism.

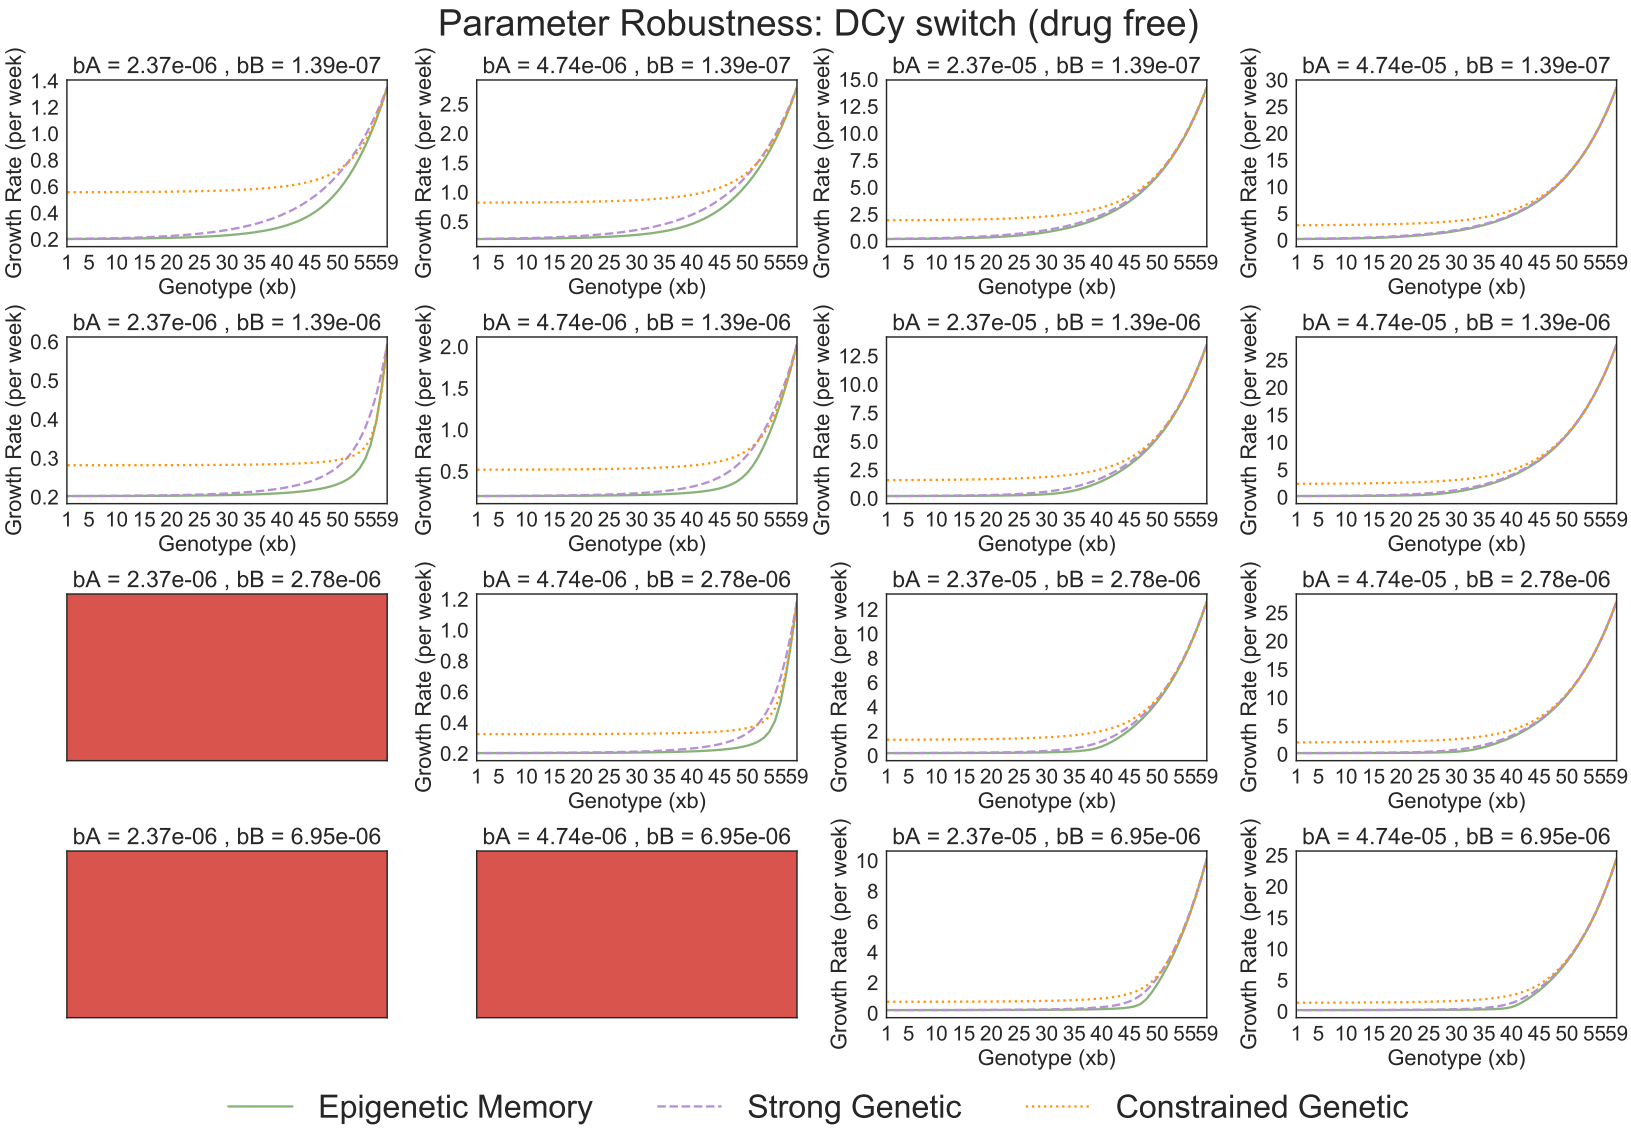

Fig. B.3 **Genotype-growth rate curves for varying parameter sets with the DCy switch.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles). Subplots shaded red indicate that the phenotype *A* growth rate is not greater than 1. In all cases phenotype *B* has parameters as in the main text (Chapter 5). Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism.

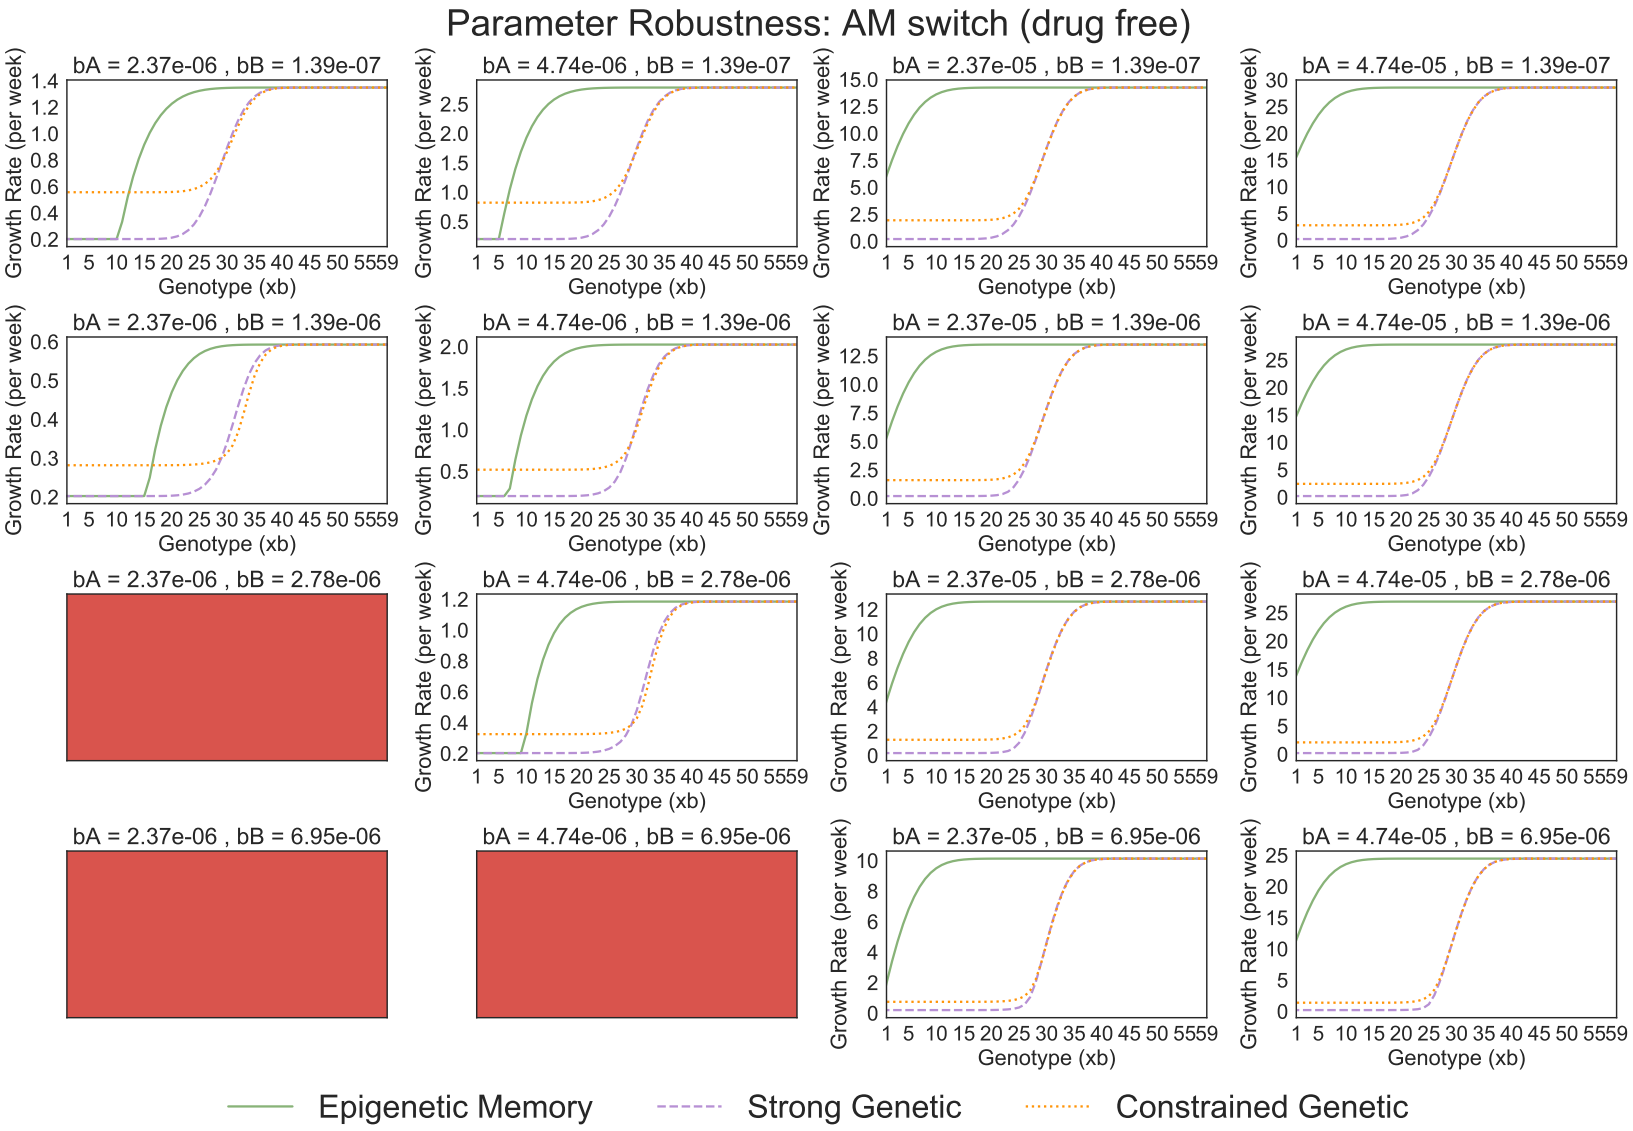

Fig. B.4 **Genotype-growth rate curves for varying parameter sets with the AM switch.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles). Subplots shaded red indicate that the phenotype *A* growth rate is not greater than 1. In all cases phenotype *B* has parameters as in the main text (Chapter 5). Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism.

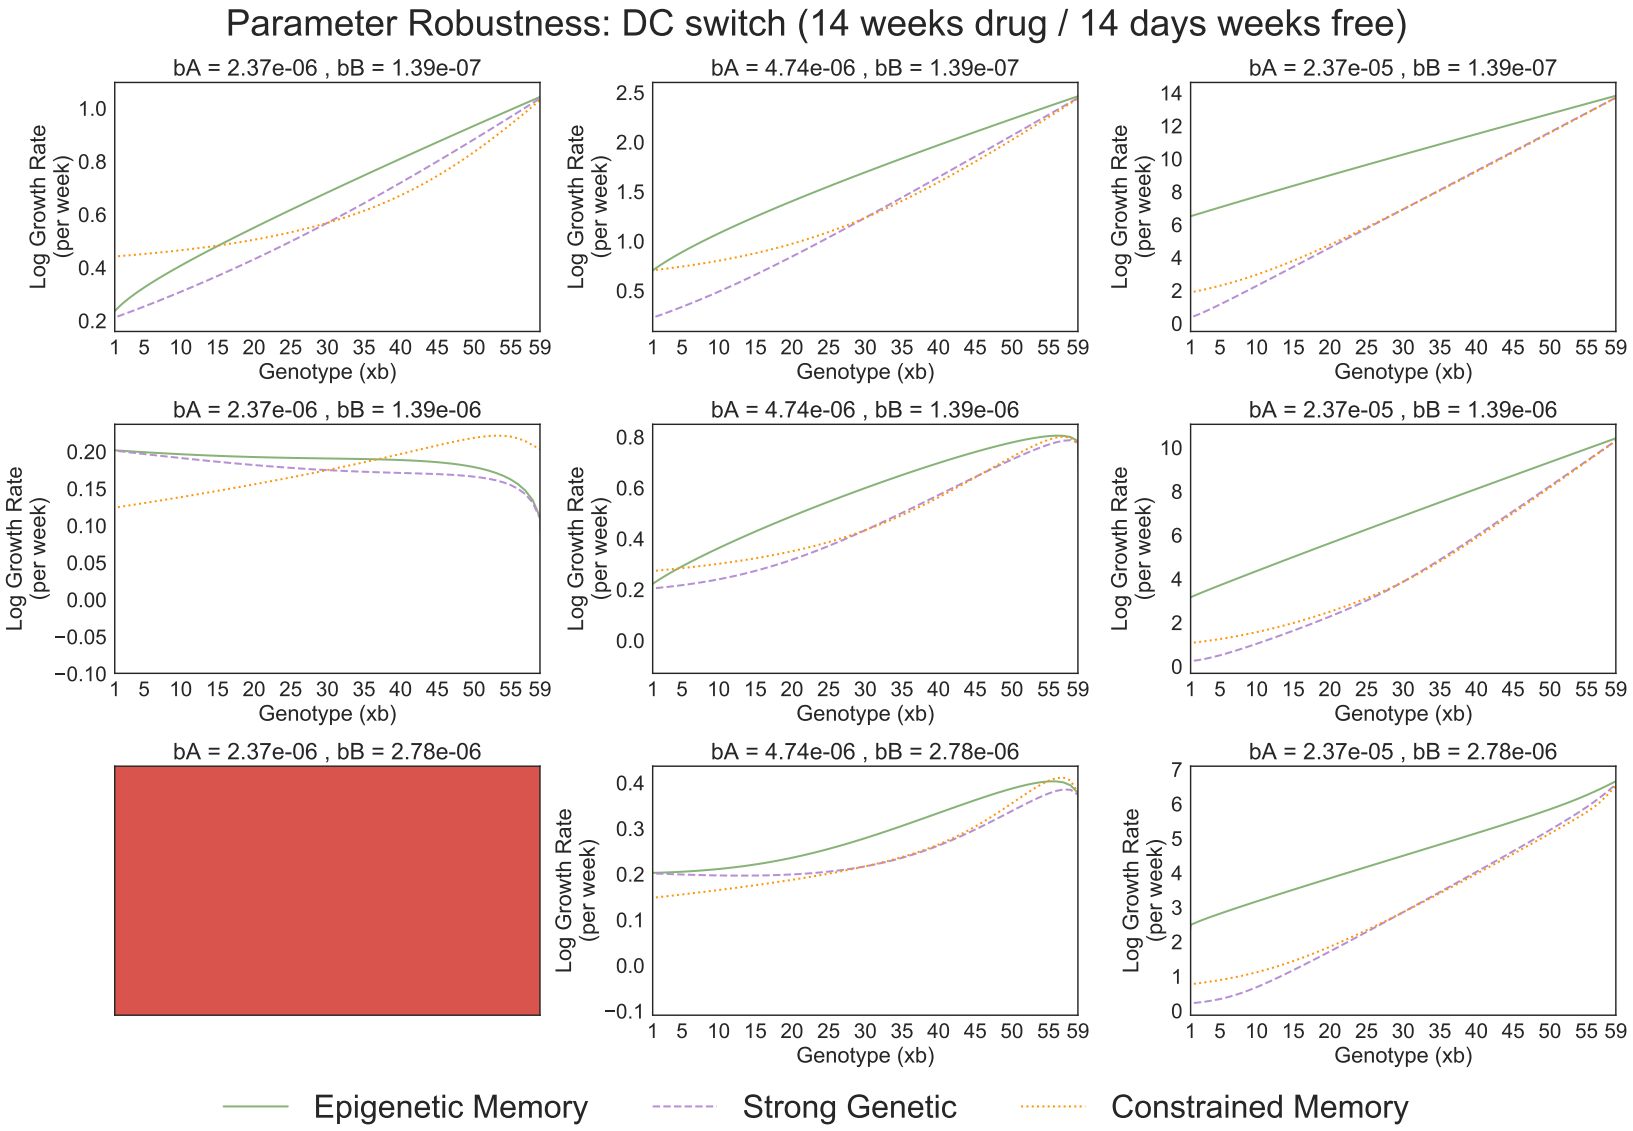

Fig. B.5 **Genotype-growth rate curves for varying parameter sets with the DC switch under a fluctuating environment.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles) under a fluctuating environment comprising 14 weeks of drug exposure followed by 14 weeks of drug–free growth. Subplots shaded red indicate that the phenotype A growth rate is not greater than 1. In all cases phenotype B has parameters as in the main text (Chapter 5) and $b^{\text{drug}} = b, d^{\text{drug}} = 8d$. Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism but that bet–hedging is not universally beneficial.

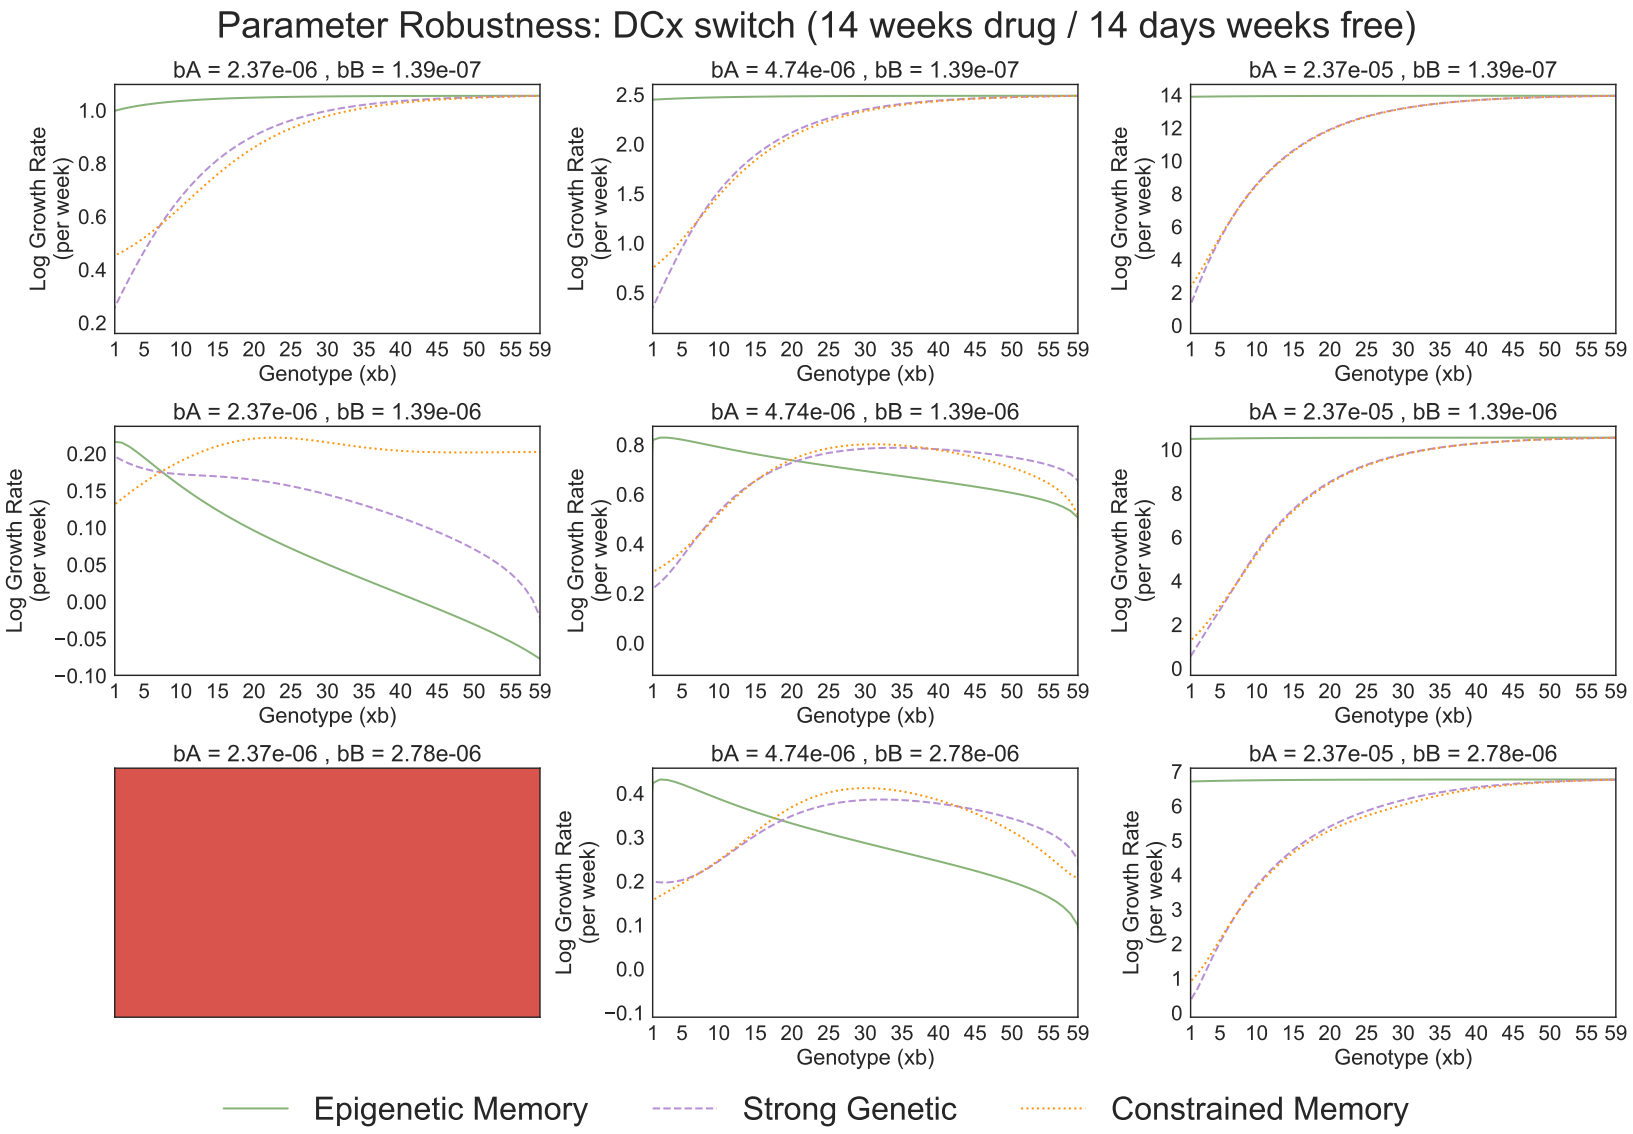

Fig. B.6 **Genotype-growth rate curves for varying parameter sets with the DCx switch under a fluctuating environment.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles) under a fluctuating environment comprising 14 weeks of drug exposure followed by 14 weeks of drug–free growth. Subplots shaded red indicate that the phenotype A growth rate is not greater than 1. In all cases phenotype B has parameters as in the main text (Chapter 5) and $b^{\text{drug}} = b, d^{\text{drug}} = 8d$. Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism but that bet–hedging is not universally beneficial.

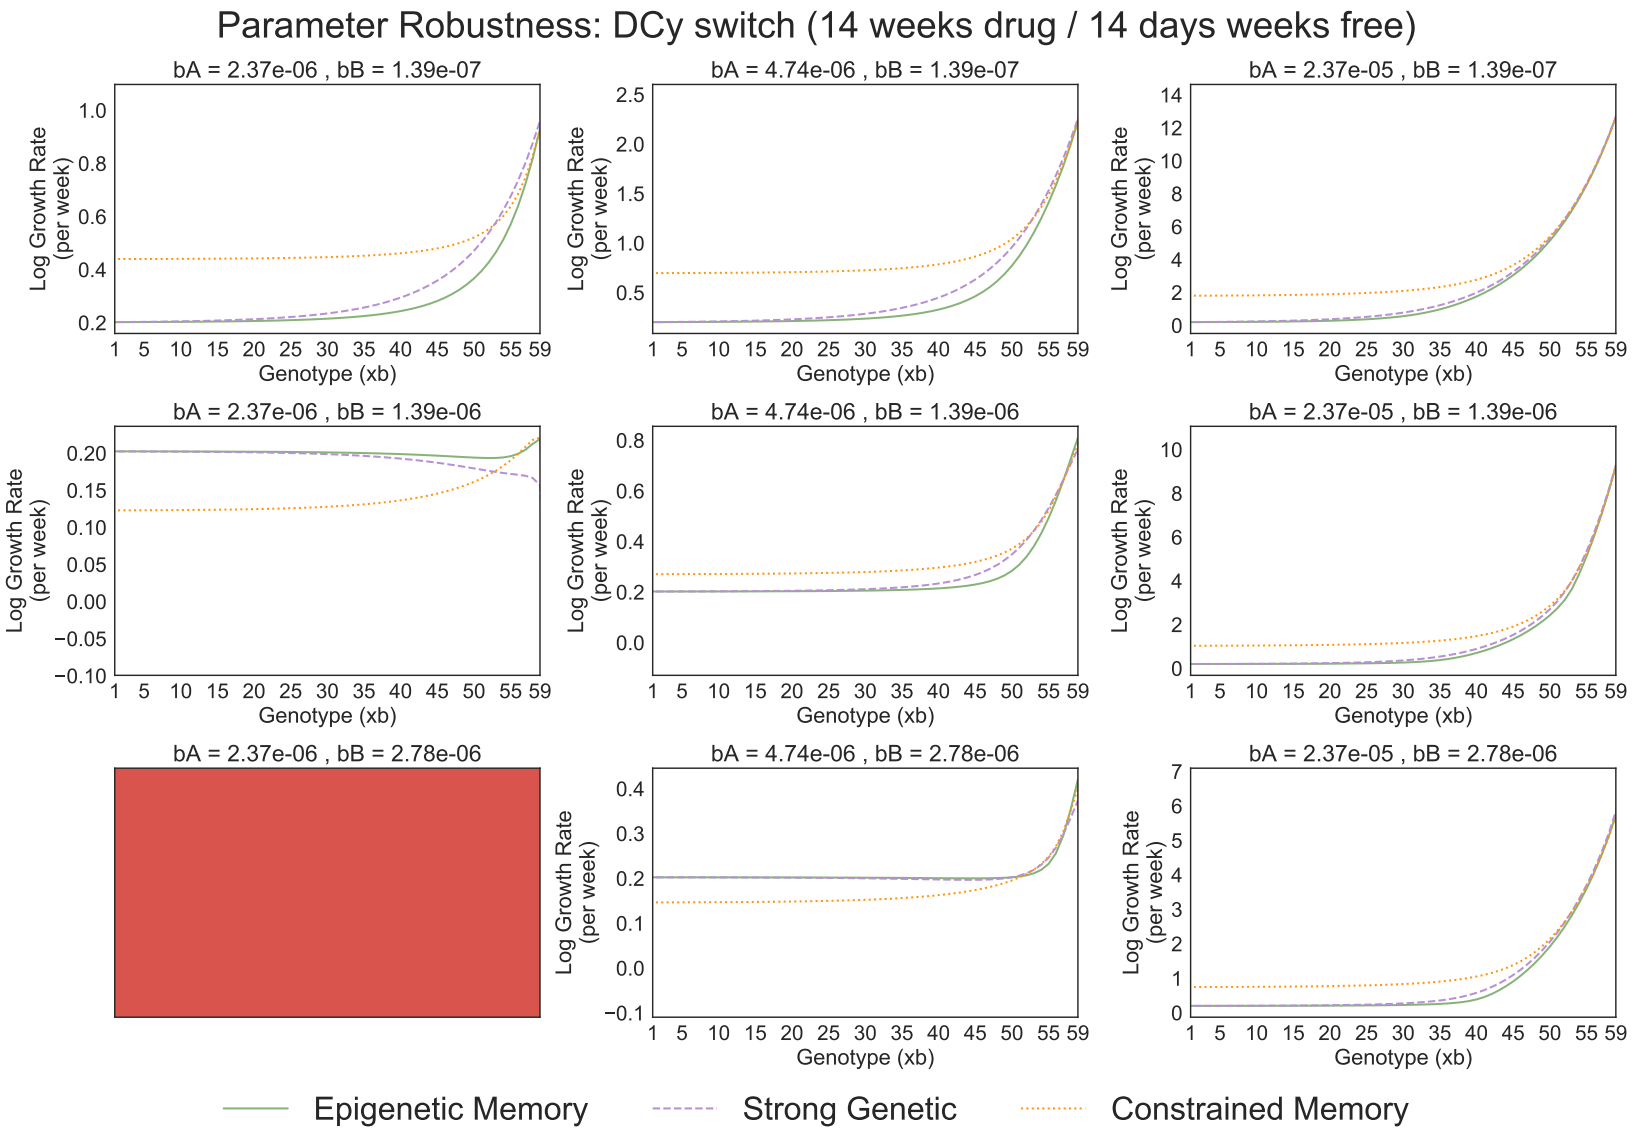

Fig. B.7 **Genotype-growth rate curves for varying parameter sets with the DCy switch under a fluctuating environment.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles) under a fluctuating environment comprising 14 weeks of drug exposure followed by 14 weeks of drug–free growth. Subplots shaded red indicate that the phenotype A growth rate is not greater than 1. In all cases phenotype B has parameters as in the main text (Chapter 5) and $b^{\text{drug}} = b, d^{\text{drug}} = 8d$. Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism but that bet–hedging is not universally beneficial.

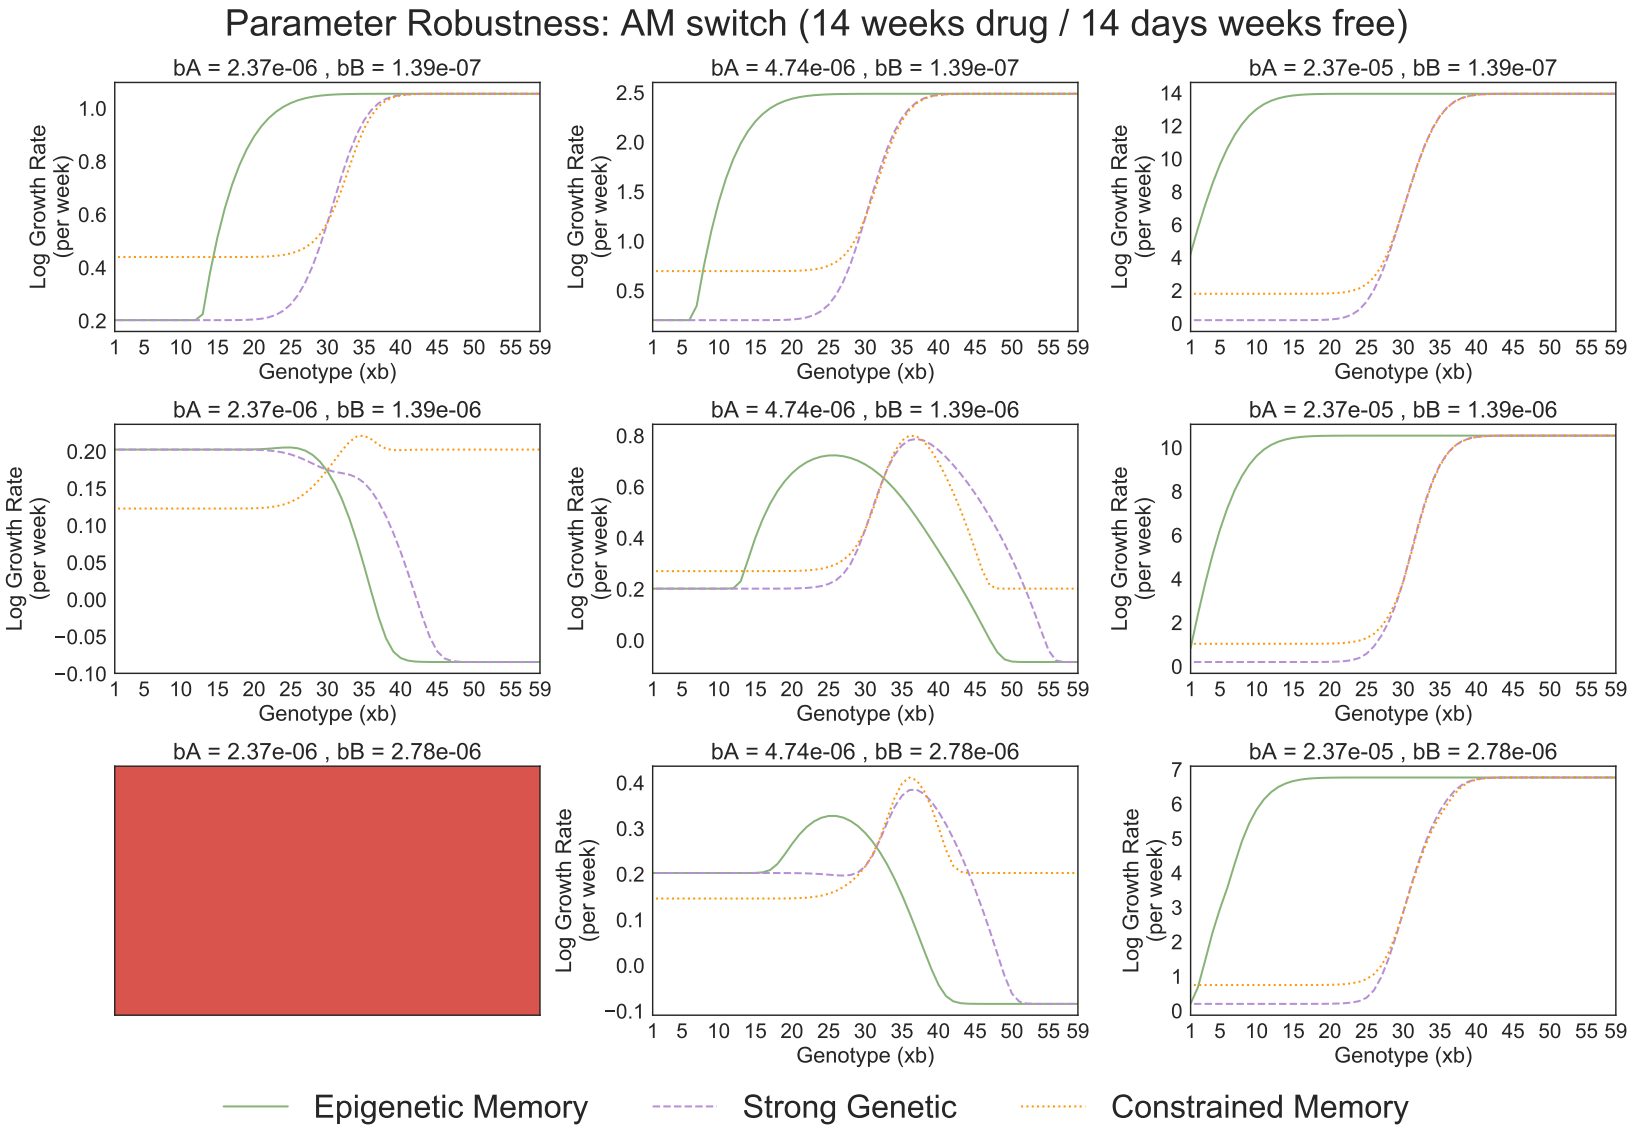

Fig. B.8 **Genotype-growth rate curves for varying parameter sets with the AM switch under a fluctuating environment.** Each subplot shows the relationship between genotype and growth rate for the three memory mechanisms (green - epigenetic memory, purple - strong genetic, orange - constrained memory) for specific parameter set (subplot titles) under a fluctuating environment comprising 14 weeks of drug exposure followed by 14 weeks of drug–free growth. Subplots shaded red indicate that the phenotype A growth rate is not greater than 1. In all cases phenotype B has parameters as in the main text (Chapter 5) and $b^{\text{drug}} = b, d^{\text{drug}} = 8d$. Note that the in all cases the fitness is dependent on both the genotype and the switching mechanism but that bet–hedging is not universally beneficial.

Appendix C

Program Code Listing

A full code listing is available on the Github repository: https://github.com/Daniel-Nichol/GPmappings.

Chapter3

January 19, 2018

1 Chapter 3 Code - Steering Evolution with Sequential Therapy

Presented below is the Python code used to generate the results presented in Chapter 3. This code is partially adapted from the Github repository with doi:10.5281/zenodo.27481 published alongside the following paper:

Nichol, Daniel, et al. "Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance" PLoS computational biology 11.9 (2015): e1004493

and is also used to generate the results in the following preprint:

Nichol, Daniel, et al. "Collateral sensitivity is contingent on the repeatability of evolution." bioRxiv (2017): 185892.

Note that where a sample size is specified, the number of samples has been reduced (from the value used in the thesis) to ensure that this notebook can be evaluated without long pauses.

1.1 Prerequisites

In [5]:
```
from copy import deepcopy
import glob
import math
import random
import sys

import matplotlib
import matplotlib.pyplot as plt
from matplotlib import rc
from mpl_toolkits.axes_grid1 import Grid
import numpy as np
import seaborn as sns

print "Python version: ", sys.version
print "numpy version: ", np.__version__
print "matplotlib version: ", matplotlib.__version__

%matplotlib inline
```

```
Python version:  2.7.13 |Continuum Analytics, Inc.| (default, Dec 20 2016, 23:05:08)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.57)]
numpy version:  1.11.3
```

```
matplotlib version:  2.0.2
```

1.2 Helper functions

Here we represent genotypes as binary strings and landscapes as lists of values. As such, it is useful to define a collection of helper functions to translate between a binary string its equivalent index in the fitness landscape.

```
In [6]: #================================================================================#
        # Helper functions
        #================================================================================#
        # Computes the hamming distance between two genotypes.
        def hammingDistance(s1, s2):
            assert len(s1) == len(s2)
            return sum(ch1 != ch2 for ch1, ch2 in zip(s1, s2))

        # Takes a genotype and converts it to an integer use indexing
        # the fitness landscape list
        def convertGenotypeToInt(genotype):
            out = 0
            for bit in genotype:
                out = (out << 1) | bit
            return out

        # Converts an integer to a genotype by taking the binary value and
        # padding to the left by 0s
        def convertIntToGenotype(anInt, pad):
            offset = 2**pad
            return [int(x) for x in bin(offset+anInt)[3:]]

        # Function which returns all genotypes at Hamming distance 1 from a
        # specified genotype
        def getOneStepNeighbours(genotype):
            neighbours = []
            for x in range(0, len(genotype)):
                temp = deepcopy(genotype)
                temp[x] = (genotype[x]+1) %2 #There is some inefficiency here.
                neighbours.append(temp)
            return neighbours
```

1.3 A Fitness Landscape Class

```
In [7]: #================================================================================#
        # Defining a fitness landscape class
        #
        # This class represents a fitness landscapes as a list of fitness values
        # (self.landscape) and provdes a collection of useful methods of
```

```
# querying the landscape
#=============================================================================#
class FitnessLandscape:
    def __init__(self, landscapeValues, name=None):
        self.landscape = landscapeValues
        self.name = name

    def getFitness(self, genotype):
        fitness = self.landscape[convertGenotypeToInt(genotype)]
        return fitness

    def genotypeLength(self):
        return int(math.log(len(self.landscape), 2))

    def numGenotypes(self):
        return len(self.landscape)

    def isPeak(self, g):
        peak = True
        for h in getOneStepNeighbours(g):
            if self.getFitness(g) < self.getFitness(h):
                peak = False
                break
        return peak

    def getPeaks(self):
        peaks = []

        allGenotypes = []
        N =self.genotypeLength()
        for x in range(0, 2**N):
            allGenotypes.append(convertIntToGenotype(x, self.genotypeLength()))

        for g in allGenotypes:
            if self.isPeak(g):
                peaks.append(g)

        return peaks

    def getGlobalPeak(self):
        return convertIntToGenotype(np.argmax(self.landscape), self.genotypeLength())

    def getLowestFitnessPeak(self):
        # Finds the peaks of the landscape
        peak_genotypes = self.getPeaks()
        lowest_peak_genotype = peak_genotypes[
            np.argmin([self.getFitness(g) for g in peak_genotypes])]
        return lowest_peak_genotype
```

1.4 Kauffman's *NK* Model

We implement Kauffman's *NK* model for generating tuneable rugged fitness landscapes. In the text of the thesis we refer to this model as the *LK* model to avoid confusion with the use of *N* to represent population size in the population genetics literature. *NK* landscapes were introduced in the following publication:

Kauffman, Stuart A., and Edward D. Weinberger. "The NK model of rugged fitness landscapes and its application to maturation of the immune response." Journal of theoretical biology 141.2 (1989): 211-245.

In [13]:
```
# For each possible value of (x_i;x_{i_1}, ... , x_{i_K})
# we have a random number sampled from [0,1)
def geneWeights(K,N):
    return [[random.random() for x in range(2**(K+1))] for y in range(N)]

# Given a genotype length (N) and the number of alleles (K)
# this function randomly choses K alleles from positions {1,...N}\{i}
# which interact epistatically with the ith position
def buildInfringersTable(N,K):
    return [random.sample(range(i)+range(i+1,N),K) for i in range(N)]

# Builds a tuple for look up in the fitness table from the infringers list
def buildTuple(allele, i, infringers):
    tp = [allele[i]]
    for j in infringers:
        tp += [allele[j]]
    return tp

# Given an allele computes the fitness by building a tuple of revelant infringers
# and looks them up in the gene weights table
def alleleFitness(allele, gw, infrs, N, K):
    s = 0.
    for i in range(N):
        index = buildTuple(allele,i,infrs[i])
        s += gw[i][convertGenotypeToInt(index)]
    s = s / N
    return s

#Generates an N-K landscape from Kauffman's method.
def generateNKLandscape(N,K):
    gw = geneWeights(K,N)
    infrs = buildInfringersTable(N,K)
    landscape = [alleleFitness(convertIntToGenotype(a,N), gw, infrs, N, K) \
                 for a in range(2**N)]
    return FitnessLandscape(landscape)
```

In [14]:
```
# Example use:
plt.figure(figsize=(5,3))
f1 = generateNKLandscape(5,3)
```

```
plt.bar(range(1,33), f1.landscape)
plt.xlim(0,33)
plt.yticks(size=15)
plt.xticks(size=15)
plt.xlabel('Genotype Index (Integer)', size=20)
plt.ylabel('Fitness', size=20)
plt.title('Random N=5, K=3 Landscape', size=22);
```

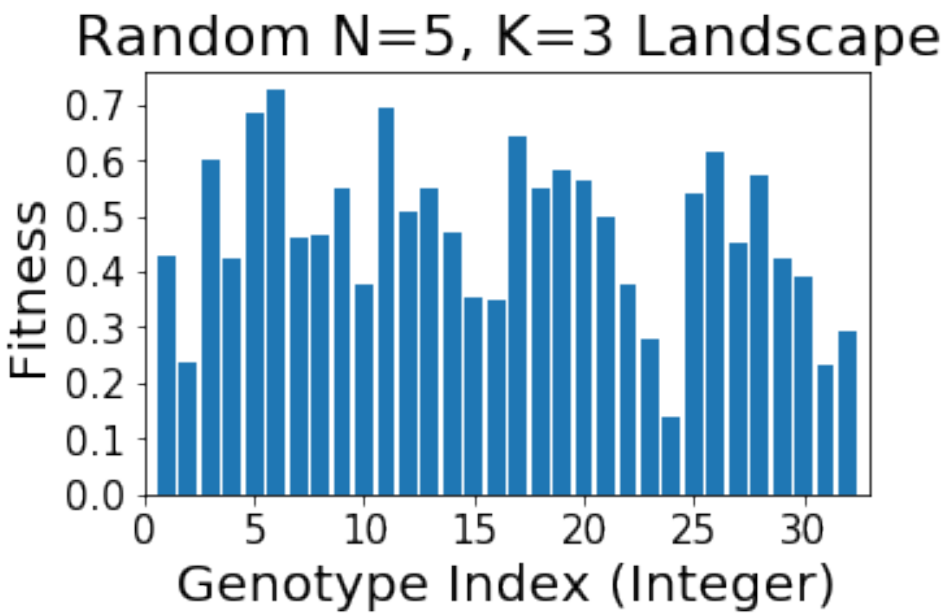

1.5 Markov Model

Here we implement the Markov chain model defined in Section 3.1

```
In [15]: #=============================================================================#
         # Equations 3.1, 3.2, 3.3 and 3.4
         #
         # Given two genotypes and a landscape, computes the transition probability
         # Pr(g1->g2) in the markov chain transition matrix
         #=============================================================================#
         def transProbSSWM(g1, g2, landscape, r=0):
             # Equation 3.1
             # If the genotypes are more than one mutation apart, then 0
             if hammingDistance(g1,g2) > 1:
                 return 0

             #Else compute Pr(g1->g2) from Equation 3.2 and 3.3
             elif hammingDistance(g1,g2) == 1:
```

```
            #Equation 3.2
            if landscape.getFitness(g1) >= landscape.getFitness(g2):
                return 0
            #Equation 3.3
            else:
                num = (landscape.getFitness(g2) - landscape.getFitness(g1))**r
                den = 0.
                for genotype in getOneStepNeighbours(g1):
                    fitDiff = (landscape.getFitness(genotype) - landscape.getFitness(g1))
                    if fitDiff > 0:
                        den += fitDiff**r
                return num / den

        #Finally add in those Pr(g1->g1)=1 for g1 a local optima (Equation 3.4)
        else:
            isPeak = landscape.isPeak(g1)
            return int(isPeak)
```

```
In [16]: #============================================================================#
         # Equation 3.5
         #
         # Builds the transition matrix for a given landscape
         #============================================================================#
         def buildTransitionMatrix(landscape, r=0):
             genomeLen = landscape.genotypeLength()
             matList = [[transProbSSWM(convertIntToGenotype(i,genomeLen),
                                       convertIntToGenotype(j, genomeLen), landscape, r) \
                         for j in range(0, 2**genomeLen)] for i in range(0, 2**genomeLen)]
             return np.matrix(matList)
```

```
In [17]: #============================================================================#
         # Computes P* from P (Equation 3.8)
         #
         # Given a stochastic matrix P, finds the limit matrix
         #============================================================================#
         def limitMatrix(P):
             Q = np.identity(len(P))
             while not np.array_equal(Q,P):
                 Q = deepcopy(P)
                 P = P*P #Square P until it no longer changes.
             return P
```

1.6 Commutativity Test

Here we demonstrate the non-commutativity of evolution, these results are listed in Section 3.2.1. This test serves as a simple example demonstrating the use of the Markov chain model defined above.

```
In [18]: %%time
         num_samples = 10**1 #10**6 #In the paper/thesis
         commuting_pairs = []
         N = 5
         for i in range(num_samples):
             K_1 = random.randint(0,N-1)
             K_2 = random.randint(0,N-1)

             l_1 = generateNKLandscape(N,K_1)
             l_2 = generateNKLandscape(N,K_2)

             limit_1 = limitMatrix(buildTransitionMatrix(l_1))
             limit_2 = limitMatrix(buildTransitionMatrix(l_2))

             if (limit_1*limit_2 == limit_2*limit_1).all():
                 commuting_pairs.append((l_1, l_2))

         print len(commuting_pairs) / float(num_samples) #Approx 0.132% for a large sample.
```

```
0.0
CPU times: user 416 ms, sys: 7.95 ms, total: 424 ms
Wall time: 429 ms
```

1.7 Model Parametrisation

The following as the landscapes derived by Mira et al in the following paper.

Mira, Portia M., et al. "Rational design of antibiotic treatment plans: a treatment strategy for managing evolution and reversing resistance." PLoS ONE 10.5 (2015): e0122283.

Below we generate Figures 3.3 and 3.5 which provide an initial exploration of the landscapes

```
In [19]: #=============================================================================#
         # The landscapes reported by Mira et. al. (2015)
         #=============================================================================#
         AMP = FitnessLandscape([1.851, 2.082, 1.948, 2.434, 2.024, 2.198, 2.033, 0.034,
                                 1.57, 2.165, 0.051, 0.083, 2.186, 2.322, 0.088, 2.821],
                                "Ampicillin")
         AM  = FitnessLandscape([1.778, 1.782, 2.042, 1.752, 1.448, 1.544, 1.184, 0.063,
                                 1.72, 2.008, 1.799, 2.005, 1.557, 2.247, 1.768, 2.047],
                                "Amoxicillin")
         CEC        = FitnessLandscape([2.258, 1.996, 2.151, 2.648, 2.396, 1.846, 2.23, 0.214,
                                 0.234, 0.172, 2.242, 0.093, 2.15, 0.095, 2.64, 0.516],
                                "Cefaclor")
         CTX = FitnessLandscape([0.16, 0.085, 1.936, 2.348, 1.653, 0.138, 2.295, 2.269,
                                 0.185, 0.14, 1.969, 0.203, 0.225, 0.092, 0.119, 2.412],
                                "Cefotaxime")
         ZOX = FitnessLandscape([0.993, 0.805, 2.069, 2.683, 1.698, 2.01, 2.138, 2.688,
                                 1.106, 1.171, 1.894, 0.681, 1.116, 1.105, 1.103, 2.591],
```

```
                        "Ceftizoxime")
CXM = FitnessLandscape([1.748, 1.7, 2.07, 1.938, 2.94, 2.173, 2.918, 3.272, 0.423,
                        1.578, 1.911, 2.754, 2.024, 1.678, 1.591, 2.923],
                        "Cefuroxime")
CRO = FitnessLandscape([1.092, 0.287, 2.554, 3.042, 2.88, 0.656, 2.732, 0.436,
                        0.83, 0.54, 3.173, 1.153, 1.407, 0.751, 2.74, 3.227],
                        "Ceftriaxone")
AMC = FitnessLandscape([1.435, 1.573, 1.061, 1.457, 1.672, 1.625, 0.073, 0.068,
                        1.417, 1.351, 1.538, 1.59, 1.377, 1.914, 1.307, 1.728],
                        "Amoxicillin+Clav")
CAZ = FitnessLandscape([2.134, 2.656, 2.618, 2.688, 2.042, 2.756, 2.924, 0.251,
                        0.288, 0.576, 1.604, 1.378, 2.63, 2.677, 2.893, 2.563],
                        "Ceftazidime")
CTT = FitnessLandscape([2.125, 1.922, 2.804, 0.588, 3.291, 2.888, 3.082, 3.508,
                        3.238, 2.966, 2.883, 0.89, 0.546, 3.181, 3.193, 2.543],
                        "Cefotetan")
SAM = FitnessLandscape([1.879, 2.533, 0.133, 0.094, 2.456, 2.437, 0.083, 0.094,
                        2.198, 2.57, 2.308, 2.886, 2.504, 3.002, 2.528, 3.453],
                        "Ampicillin+Sulbactam")
CPR = FitnessLandscape([1.743, 1.662, 1.763, 1.785, 2.018, 2.05, 2.042, 0.218,
                        1.553, 0.256, 0.165, 0.221, 0.223, 0.239, 1.811, 0.288],
                        "Cefprozil")
CPD = FitnessLandscape([0.595, 0.245, 2.604, 3.043, 1.761, 1.471, 2.91, 3.096,
                        0.432, 0.388, 2.651, 1.103, 0.638, 0.986, 0.963, 3.268],
                        "Cefpodoxime")
TZP = FitnessLandscape([2.679, 2.906, 2.427, 0.141, 3.038, 3.309, 2.528, 0.143,
                        2.709, 2.5, 0.172, 0.093, 2.453, 2.739, 0.609, 0.171],
                        "Piperacillin+Tazobactam")
FEP = FitnessLandscape([2.59, 2.572, 2.393, 2.832, 2.44, 2.808, 2.652, 0.611,
                        2.067, 2.446, 2.957, 2.633, 2.735, 2.863, 2.796, 3.203],
                        "Cefepime")

#=============================================================================#
# The limits of the Markov chain matrices corresponding to these landscapes (r=0)
#=============================================================================#
L_AMP = limitMatrix(buildTransitionMatrix(AMP))
L_AM  = limitMatrix(buildTransitionMatrix(AM))
L_CEC = limitMatrix(buildTransitionMatrix(CEC))
L_CTX = limitMatrix(buildTransitionMatrix(CTX))
L_ZOX = limitMatrix(buildTransitionMatrix(ZOX))
L_CXM = limitMatrix(buildTransitionMatrix(CXM))
L_CRO = limitMatrix(buildTransitionMatrix(CRO))
L_AMC = limitMatrix(buildTransitionMatrix(AMC))
L_CAZ = limitMatrix(buildTransitionMatrix(CAZ))
L_CTT = limitMatrix(buildTransitionMatrix(CTT))
L_SAM = limitMatrix(buildTransitionMatrix(SAM))
L_CPR = limitMatrix(buildTransitionMatrix(CPR))
L_CPD = limitMatrix(buildTransitionMatrix(CPD))
```

```
L_TZP = limitMatrix(buildTransitionMatrix(TZP))
L_FEP = limitMatrix(buildTransitionMatrix(FEP))

landscapes = [AMP, AM, CEC, CTX, ZOX, CXM, CRO, AMC,
              CAZ, CTT, SAM, CPR, CPD, TZP, FEP]
limit_matrices = [L_AMP, L_AM, L_CEC, L_CTX, L_ZOX, L_CXM,
                  L_CRO, L_AMC, L_CAZ, L_CTT, L_SAM, L_CPR, L_CPD, L_TZP, L_FEP]
labs = ["AMP", "AM", "CEC", "CTX", "ZOX", "CXM",
        "CRO", "AMC", "CAZ", "CTT", "SAM", "CPR", "CPD", "TZP", "FEP"]
```

Plotting the fitness landscapes as bar plots, highlighting fitness peaks:

```
In [20]: #============================================================================#
         # Figure 3.3
         #============================================================================#
         def lab(k):
             gt = convertIntToGenotype(k, 4)
             l = "".join(map(str,gt))
             return l

         def col(ls, k):
             r = sns.xkcd_rgb['pale red']
             b = sns.xkcd_rgb['denim blue']
             if ls.isPeak(convertIntToGenotype(k, 4)):
                 return sns.xkcd_rgb['pale red']
             else:
                 return sns.xkcd_rgb['denim blue']

         fig = plt.figure(figsize=(18,12))
         for k, ls in enumerate(landscapes):
             plt.subplot(3,5,k+1)
             cols = map(lambda k : col(ls, k), range(2**4))
             barlist = plt.bar(range(2**4), ls.landscape, color=cols)
             plt.ylim(0.0,3.0)
             plt.title(labs[k], size=16)
             if labs[k] in ['AMP', 'SAM', 'CPR']:
                 plt.text(-1,3.5,'*',size=30)
             plt.xticks(range(2**4), map(lab, range(2**4)), size=9, rotation='vertical')
             plt.yticks(np.arange(0.0,4.1 ,0.5))

         plt.subplots_adjust(hspace=0.6)
         fig.add_subplot(111, frameon=False)
         # hide tick and tick label of the big axes
         plt.tick_params(labelcolor='none', top='off', bottom='off', left='off', right='off')
         plt.ylabel("Avg. Growth Rate (x10e-3 cells per minute)", size=18)
         plt.suptitle("Fitness Landscapes", size=28)
         plt.tight_layout(rect=[0, 0.03, 1, 0.95])
```

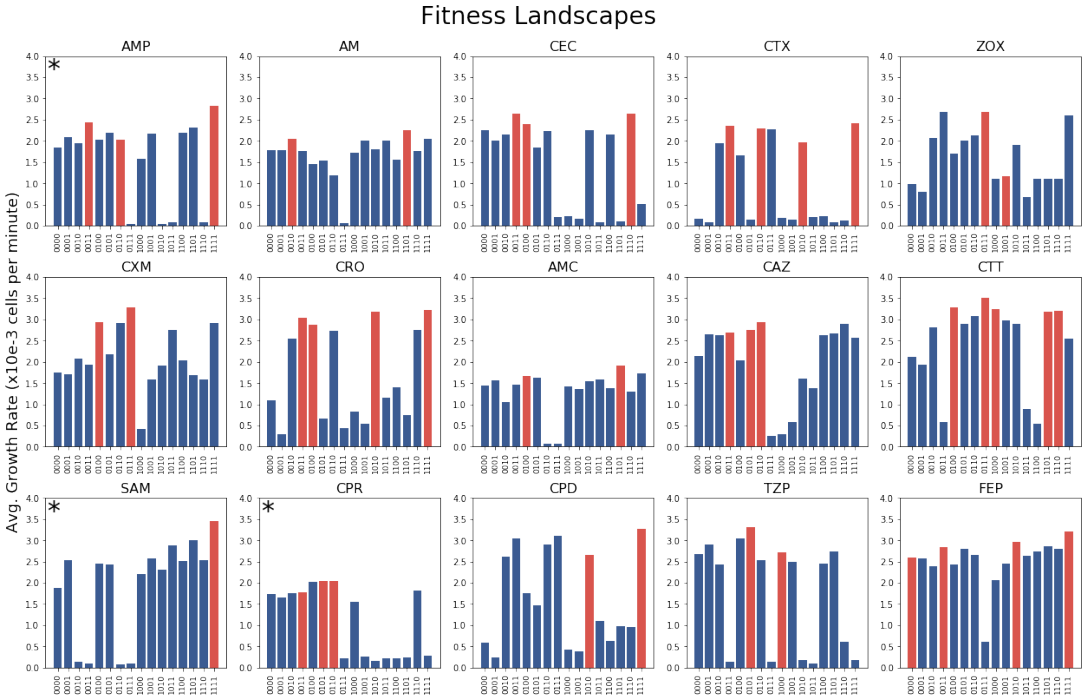

The following code produces scatter plots showing the pairwise fitness between the different landscapes, along with the correlation coefficient. This generates Figure 3.5 as well as the supplementary figures provided in Appendix A.

```
In [21]: #============================================================================#
         # Figure 3.5
         #============================================================================#
         def scatter_fitness(l1,l2):
             xsnp,ysnp = [],[] #neither a peak in l1 or l2
             xsxp,ysxp = [],[] #peak in l1, not in l2
             xsyp,ysyp = [],[] #peak in l2, not in l1
             xsbp,ysbp = [],[] #peak in l1 & l2
             for g in map(lambda x : convertIntToGenotype(x,4), range(16)):
                 if l1.isPeak(g):
                     if l2.isPeak(g):
                         xsbp.append(l1.getFitness(g))
                         ysbp.append(l2.getFitness(g))
                     else:
                         xsxp.append(l1.getFitness(g))
                         ysxp.append(l2.getFitness(g))

                 else:
                     if l2.isPeak(g):
                         xsyp.append(l1.getFitness(g))
                         ysyp.append(l2.getFitness(g))
```

```
                else:
                    xsnp.append(l1.getFitness(g))
                    ysnp.append(l2.getFitness(g))

        return xsnp, ysnp, xsxp, ysxp, xsyp, ysyp, xsbp, ysbp

    def scatter_plot(ax, xl, yl, fontsize=12):
        import matplotlib.ticker as ticker
        from scipy.stats import spearmanr
        xsnp,ysnp,xsxp,ysxp,xsyp,ysyp,xsbp,ysbp = scatter_fitness(landscapes[xl],
                                                                  landscapes[yl])

        fullxs = xsnp + xsxp + xsyp + xsbp
        fullys = ysnp + ysxp + ysyp + ysbp

        #scatter plot
        ax.scatter(xsnp,ysnp, s=50, marker=".", c='k')
        ax.scatter(xsxp,ysxp, s=50, marker=".", c='b')
        ax.scatter(xsyp,ysyp, s=50, marker=".", c='r')
        ax.scatter(xsbp,ysbp, s=50, marker=".", c='g')

        #axis ticks
        ax.get_xaxis().set_visible(True)
        ax.get_yaxis().set_visible(True)
        ax.set_xlim((0.0,4.0))
        ax.set_ylim((0.0,4.0))
        ax.xaxis.set_ticks(np.arange(0.0, 4.5, 2.0))
        ax.xaxis.set_major_formatter(ticker.FormatStrFormatter('%0.1f'))
        ax.yaxis.set_ticks(np.arange(0.0, 4.5, 2.0))
        ax.yaxis.set_major_formatter(ticker.FormatStrFormatter('%0.1f'))

        #axis labels
        ax.set_xlabel(labs[xl])
        ax.set_ylabel(labs[yl])


        #Adding the spearman r
        rho = spearmanr(fullxs, fullys)[0]
        ax.annotate(r'$\rho =$'+"{:.2f}".format(rho), xy=(1.0, 1.0),
            xytext=(.47, .85), textcoords='axes fraction', zorder=100)

    width, height=5, 5
    rc('font',**{'family':'sans-serif','sans-serif':['Helvetica']})
    fig = plt.figure(figsize=(10,10))
    sns.set_style('darkgrid')

    grid = Grid(fig, rect=111, nrows_ncols=(width,height), axes_pad=0.25, label_mode='L')
    i,j = 0,0 #Change these values to generate figures A1-A8
```

```
for ax in grid:
    scatter_plot(ax,i,j)
    ax.title.set_visible(False)
    i+=1
    if i == width:
        i = 0
        j+=1
```

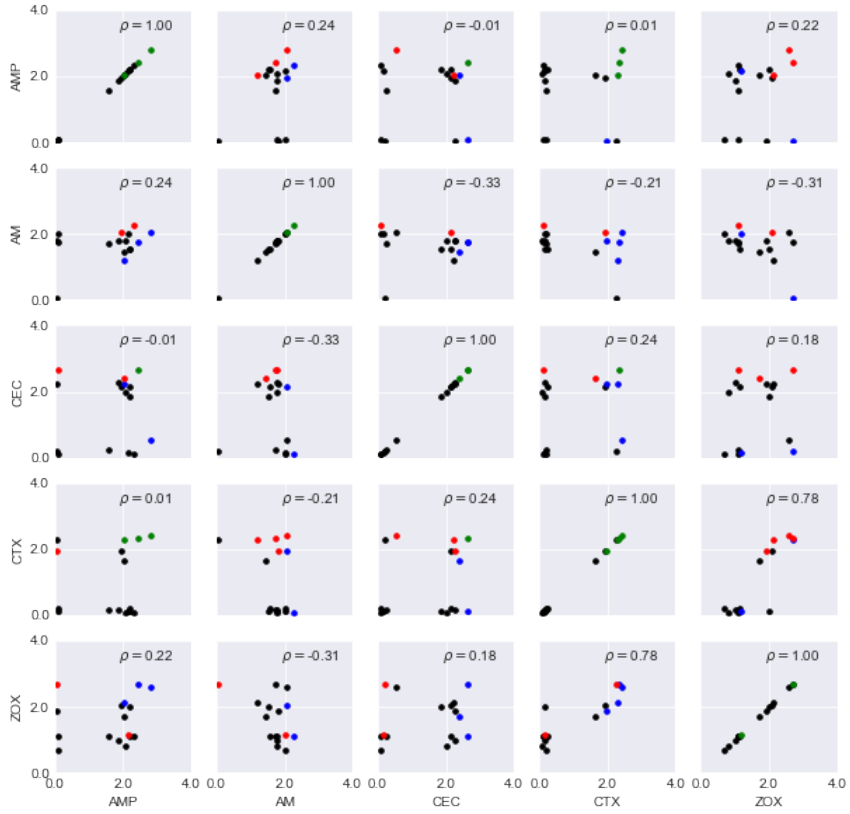

1.8 Collateral Sensitivity

The following code generates the results and figures outlined in Section 3.3. In our exploration of collateral sensitivity, we assume that the initial population is a "wild-type" population, mirroring experimental evolution of drug resistance. As such, we specify an initial population as follows

```
In [22]: # g=0000 is the initial genotype
         init_wt = np.array([1.]+[0. for i in range(15)])
```

To determine the collateral response under a second drug y, first we randomly select a fitness peak genotype (g_x^*) in the first landscape, x, arising from evolution from the initial genotype g_0. Next, the collateral response is calculated as

Collateral response of Y to $X = \log_2\left(\frac{f_y(g_x^*)}{f_y(g_0)}\right)$

In [23]:
```
#=============================================================================#
# Simulates a single instance of determining collateral sensitivity
#
# As evolution is not necessarily repeatable, this is not guaranteed to generate
# the same matrix of collateral sensitivity on each instance.
#=============================================================================#
def col_sensitivity(d1, d2, init_pop, relative=True):
    limit_d1 = limit_matrices[d1]
    pop_dist = init_pop * limit_d1
    pop_dist = np.array(pop_dist)[0]
    peak = np.random.choice(np.array([i for i in range(16)]), p=pop_dist)
    if relative:
        col_sens = np.log2(landscapes[d2].getFitness(convertIntToGenotype(peak,4)) \
                           / landscapes[d2].getFitness(convertIntToGenotype(0,4)))
    else:
        col_sens = landscapes[d2].getFitness(convertIntToGenotype(peak,4)) \
                   / landscapes[d2].getFitness(convertIntToGenotype(0,4))

    return col_sens

def generateCSN(init_pop):
    network = [[col_sensitivity(d1,d2,init_pop) for d1 in range(len(landscapes))]
               for d2 in range(len(landscapes))]
    return network
```

In [24]:
```
#=============================================================================#
# Plotting functions used to display a table of collateral response
#=============================================================================#
# Displays the matrix/heatmap
def show_CSN(network):
    fig = plt.figure(figsize=(5.3,4.5))
    ax = fig.add_subplot(111)
    cax=ax.matshow(network, cmap = plt.cm.RdBu_r, vmin = -5, vmax = 5)
    cbar = fig.colorbar(cax)
    cbar.ax.set_yticklabels(['32x','16x','8x','4x','2x','1x',
                             '2x','4x','8x','16x','32x'])
    cbar.set_label('Change From Wild-Type Sensitivity')

    ax.set_xticks(range(len(labs)))
    ax.set_yticks(range(len(labs)))
    ax.set_xticklabels(labs, rotation=45)
    ax.set_yticklabels(labs)
    ax.axis('image')

    ax.set_xlabel('First Drug')
    ax.xaxis.set_label_position('top')
    ax.set_ylabel('Second Drug')
```

```
            plt.show()

        # Displays a bar chart of showing number of CS/CR drug pairs
        def bar_comparison(network):
            bets, wors = [],[]
            for d2 in range(len(landscapes)):
                row = np.array(network)[d2]
                bets.append(sum(map(lambda x : x<0.0, row)))
                wors.append(sum(map(lambda x : x>0.0, row)))

            bets = map(lambda x : -x, bets)
            cmap = plt.cm.RdBu

            res_bar_col = cmap(0.1)
            sen_bar_col = cmap(0.9)

            fig, ax = plt.subplots(figsize=(5,5))
            ax = plt.subplot(1,2,1)
            rects1 = ax.barh(np.arange(len(landscapes)), wors, color=res_bar_col)
            rects2 = ax.barh(np.arange(len(landscapes)), bets, color=sen_bar_col)
            plt.xlim(-15,15)
            plt.ylim(-1,15)
            ax.set_xticklabels([15,10,5,0,5,10,15])
            plt.xlabel('Number of Second Drugs')
            plt.ylabel('First Drug')
            plt.yticks(np.arange(0.0,len(landscapes)+0.0,1.0), labs)
            plt.gca().invert_yaxis()
            plt.title('Collaterally Sensitive/Resistant')
```

In [26]:
```
        #================================================================================#
        # Example use.
        #================================================================================#
        sns.set_style('white')
        plt.figure(figsize=(5,5))
        example_csn = generateCSN(init_wt)
        show_CSN(example_csn)
        # bar_comparison(example_csn) #Make the bar chart...
```

```
<matplotlib.figure.Figure at 0x11cf18690>
```

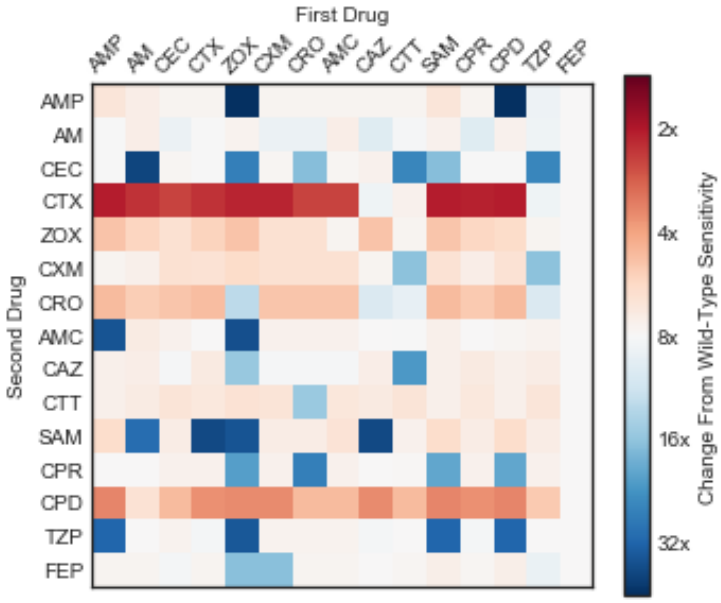

1.8.1 Exploring the space of collateral sensitivity tables

The following methods are used to explore the space of collateral sensitivity tables to find the most likely outcome along with the best-case, worst-case and average-case outcome for a table of collateral sensitivity determined from a single round of experimental evolution.

```
In [29]: #============================================================================#
         # Determines the most likely CSM and the associated probability
         #============================================================================#
         def most_likely(init_pop):
             gs,ps = [],[]
             for d in range(15):
                 limit = limit_matrices[d]
                 pop_dist = np.array(init_pop * limit)[0]
                 g = np.argmax(pop_dist)
                 gs.append(g)
                 ps.append(pop_dist[g])

             network = [[np.log2(landscapes[d2].getFitness(convertIntToGenotype(g,4)) \
                                 /landscapes[d2].getFitness(convertIntToGenotype(0,4))) \
                         for g in gs] \
                         for d2 in range(15)]
             p = np.product(ps)
             return p,network

         #============================================================================#
         # Determines the (almost certain) best and worst case outcome by means of multiple
```

```
        # trials.
        #=========================================================================#
        def best_and_worst(trials, init_pop):
            networks = []
            for t in range(trials):
                network =  np.array(generateCSN(init_pop))
                networks.append(network)

            mins,maxs = networks[0], networks[0]
            for net in networks[1:]:
                mins = np.minimum(mins,net)
                maxs = np.maximum(maxs,net)
            return mins,maxs

        #=========================================================================#
        # Determines the average collateral response for each drug pair
        #=========================================================================#
        def expected_network(init_pop):

            def expected_sensitivity(d1,d2,init_pop):
                limit_d1 = limit_matrices[d1]
                pop_dist = init_pop * limit_d1
                pop_dist = np.array(pop_dist)[0]
                expected_fitness = np.log2(sum([pop_dist[i] * \
                                        landscapes[d2].getFitness(convertIntToGenotype(i,4))
                                        for i in range(16)])
                 / landscapes[d2].getFitness(convertIntToGenotype(0,4)))
                return expected_fitness

            network =[[expected_sensitivity(d1,d2,init_pop) \
                        for d1 in range(len(landscapes))] \
                    for d2 in range(len(landscapes))]
            return network
```

In [30]:
```
        #=========================================================================#
        # Example use, generating Figure 3.7
        #=========================================================================#
        exp  = expected_network(init_wt)
        best, worst = best_and_worst(1000, init_wt)
        p,net = most_likely(init_wt)
        print "The most likely CSM occurs with probability: ", p

        # # Uncomment all of the following generate Figure 3.7 (in parts):
        # show_CSN(exp)
        # bar_comparison(exp)
        # show_CSN(best)
        # bar_comparison(best)
        # show_CSN(worst)
```

```
        # bar_comparison(worst)
        # show_CSN(net)
        # bar_comparison(net)
```

```
The most likely CSM occurs with probability:  0.00229219059595
```

1.8.2 Simulating experimental evolution to find CS drugs

In the thesis (Section 3.3.1), we asked the following two questions: 1. What is the probability that a randomly chosen the drug pair exhibits cross resistance?
2. Given that a drug pair exhibits collateral sensitivity in a single simulation of experimental evolution, what is it the likelihood that it exhibits cross resistance in second such experiment?

The following functions determine the answers to these questions.

```
In [31]: #============================================================================#
         # Computes the likelihood of CLR from a random drug pair
         #
         # Here we assume that the drug pair is chosen uniformly at random. As such,
         # the probability of choosing a CR pair at random can be determined analytically by
         # summing, over all possible pairs, the total probability mass of collateral
         # resistance and then normalising appropriately.
         #============================================================================#
         def random_cl():
             tot = 0.
             for d1 in range(15):
                 for d2 in range(15):
                     if d1!=d2:
                         temp = 0
                         wt_f = landscapes[d1].getFitness([0,0,0,0])
                         ev_pop = init_wt * limit_matrices[d1]
                         ev_pop = np.array(ev_pop)[0]
                         for k in range(len(ev_pop)):
                             ev_f = landscapes[d2].getFitness(convertIntToGenotype(k,4))
                             if ev_f > wt_f:
                                 temp+=ev_pop[k]
                         tot+=temp
             tot = tot / (15. * 14.)
             return tot

         #============================================================================#
         # Determines the probability that a pair identified as CS in a single simulation of
         # experimental evolution is CR in a second simulation.
         #
         # This is determined by repeatedly sampling a random matrix of collateral response,
         # choosing a CS pair at random and then determining the likelihood that the first of
         # that pair induces CR in the second.
         #============================================================================#
```

```
        def prob_clr():
            #generate a random csn
            csn = generateCSN(init_wt)
            #Pick a pair with CS are random
            flag = True
            while flag:
                d1 = random.randint(0,14)
                d2 = random.randint(0,14)
                if csn[d1][d2] < 0.0:
                    flag = False

            #find the probability of clr.
            prob_res = 0.0
            wt_f = landscapes[d1].getFitness([0,0,0,0])
            ev_pop = init_wt * limit_matrices[d1]
            ev_pop = np.array(ev_pop)[0]
            for k in range(len(ev_pop)):
                ev_f = landscapes[d2].getFitness(convertIntToGenotype(k,4))
                if ev_f > wt_f:
                    prob_res += ev_pop[k]

            return prob_res

        #10000 samples:
        S = 10 #10000 #In the paper/thesis
        cr_ps =  np.mean([prob_clr() for _ in range(S)])

        print "The probability of CR in a random pair is: ", random_cl()
        print "The probability of CR in a random pair identified as CS is: ", cr_ps
```

```
The probability of CR in a random pair is:  0.57671957672
The probability of CR in a random pair identified as CS is:  0.475
```

1.9 Collateral Sensitivity Likelihoods

Here we derive likelihood of collateral sensitivity between pairs of drugs. This code produces the results in Section 3.3.2 of the thesis.

```
In [35]: #============================================================================#
         # Given two drugs, returns the probability that evolution
         # under the first results in a genotype with fitness under the
         # second that is lower than the WT fitness.
         #============================================================================#
         def prob_better(d1,d2,init_pop):
             limit_d1 = limit_matrices[d1]
             pop_dist = init_pop * limit_d1
             pop_dist = np.array(pop_dist)[0]
```

```
        prob = 0
        wt_f = landscapes[d2].getFitness([0,0,0,0])
        for i in range(len(pop_dist)):
            if landscapes[d2].getFitness(convertIntToGenotype(i,4)) < wt_f:
                prob+=pop_dist[i]

        return prob
    #=============================================================================#
    # Collateral Sensitivity Likelihood matrix
    #=============================================================================#
    def prob_matrix(init_pop):
        mat = [[prob_better(d1,d2,init_pop) for d1 in range(len(landscapes))] \
               for d2 in range(len(landscapes))]
        return mat

    #=============================================================================#
    # Displays a CSL
    #=============================================================================#
    def show_CSL(mat):
        fig = plt.figure(figsize  = (4,3))
        ax = fig.add_subplot(111)
        cax=ax.matshow(mat, cmap = plt.cm.Greens, vmin = 0.0, vmax = 1.0)

        #cbar = fig.colorbar(cax)
        #cbar.set_label('Probability of Collateral Sensitivity')

        ax.set_xticks(range(len(labs)))
        ax.set_yticks(range(len(labs)))
        ax.set_xticklabels(labs, rotation=45, size='14')
        ax.set_yticklabels(labs, size='14')
        ax.axis('image')

        ax.set_xlabel('First Drug Applied', size='18')
        ax.xaxis.set_label_position('top')
        ax.set_ylabel('Second Drug (To Be Applied)', size='18')

    #=============================================================================#
    # Returns a CSL in which all entries <p are set to zero
    #=============================================================================#
    def cut_off_probs_CSL(network, p):
        new_net = [[x if p-x<=10**(-8) else 0.0 for x in y] for y in network]
        return new_net
```

In [36]:
```
    #=============================================================================#
    # Example use. Figure 3.8 (in parts)
    #=============================================================================#
    sns.set_style('white')
```

```
csl = prob_matrix(init_wt)
# show_CSL(csl)  #Uncomment for other figures.
show_CSL(cut_off_probs_CSL(csl, 1.0))
# show_CSL(cut_off_probs_CSL(csl, 0.75))
```

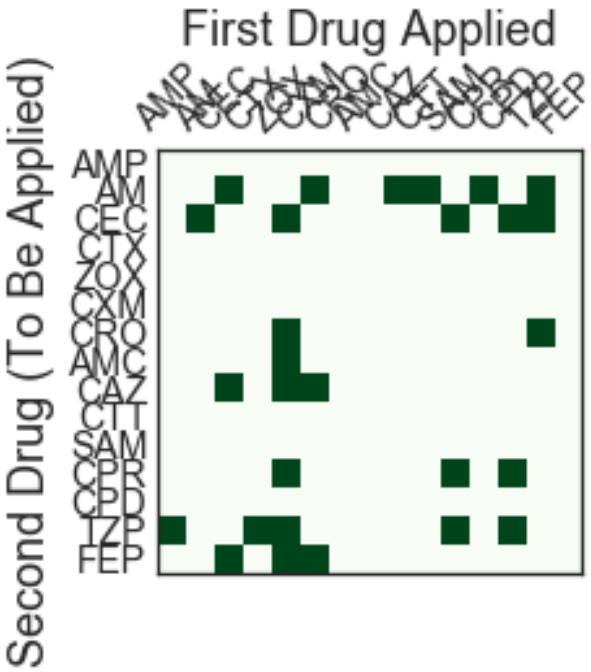

1.10 Worst Case Collateral Response

The following code identified those drug pairs where the worst case collateral response lies below a specified threshold for fold change in drug resistance. These results are presented in Figure 3.9

```
In [37]: def thresh_network(init_pop,thresh):

             # Thresholded probabilistic CSNs. What is the likelihood that your fitness is belou
             def prob_thresh(d1,d2,init_pop, thresh):
                 limit_d1 = limit_matrices[d1]
                 pop_dist = init_pop * limit_d1
                 pop_dist = np.array(pop_dist)[0]
                 prob = 0.0

                 wt_f = landscapes[d2].getFitness([0,0,0,0])

                 for i in range(len(pop_dist)):
                     if (landscapes[d2].getFitness(convertIntToGenotype(i,4))/wt_f)  < thresh:
                         prob+=pop_dist[i]
```

```
            #return prob

            if (prob - 1.0)**2 < 10**-7:
                return 1.0
            else:
                return 0.0


        mat = [[prob_thresh(d1,d2,init_pop,thresh) for d1 in range(len(landscapes))] \
               for d2 in range(len(landscapes))]
        fig = plt.figure()
        ax = fig.add_subplot(111)
        cax=ax.matshow(mat, cmap = plt.cm.Blues, vmin = 0.0, vmax = 1.0)

        #cbar = fig.colorbar(cax)
        #cbar.set_label('Probability of Collateral Sensitivity')

        ax.set_xticks(range(len(labs)))
        ax.set_yticks(range(len(labs)))
        ax.set_xticklabels(labs, rotation=45)
        ax.set_yticklabels(labs)
        ax.axis('image')

        ax.set_xlabel('First Drug Applied')
        ax.xaxis.set_label_position('top')
        ax.set_ylabel('Second Drug (To Be Applied)')
        plt.show()

    #Example, pairs that guarantee CS:
    thresh_network(init_wt, 1.0)
```

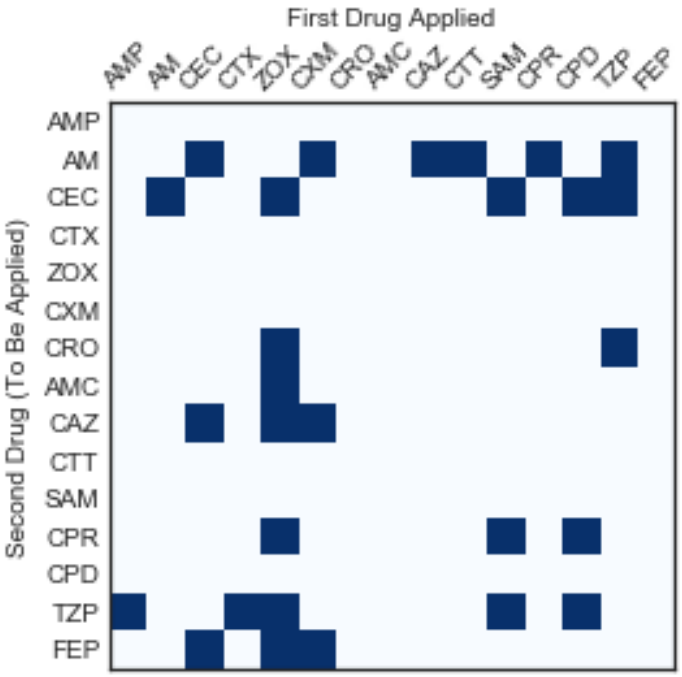

2 Evolutionary Steering

Below we derive the results presenting in Section 3.4 which outline the potential for carefully chosen drug sequences to minimise the likelihood of highly resistant disease. Here, we assume that the initial genotype for the disease population is unknown, as such we set:

```
In [420]: # Each genotype is equally likely to be the initial genotype
          init_pop = np.array([1./2**4 for i in range(16)])
```

The following code produces the results presented in Table 3.3, indicating the best sequence of 1, 2 or 3 drugs to steer evolution to **avoid highly resistant** genotypes.

```
In [39]: #============================================================================#
         # The best steering combinations are determined by an exhaustive search of all
         # sequences of 1, 2 or 3 drugs.
         #============================================================================#

         #============================================================================#
         # Determines for each drug the probability of ending at the highest fitness peak
         # genotype when starting from the initial distribution init_pop.
         #============================================================================#
         def probHighestPeak(init_pop):
             highest_pgs_probs = []
             for t in range(len(limit_matrices)):
```

```
            dist = np.array(init_pop * limit_matrices[t])[0]
            global_peak_index = convertGenotypeToInt(landscapes[t].getGlobalPeak())
            highest_pgs_probs.append(dist[global_peak_index])
            print landscapes[t].name+":", highest_pgs_probs[t]

        return highest_pgs_probs
    #============================================================================#
    # Determines for each drug the single steering drug which minimizes the
    # probability that evolution proceeds to the highest fitness peak.
    #============================================================================#
    def highPeakBestSingle(init_pop):
        best_steerers = []
        for t in range(len(limit_matrices)):
            global_peak_index = convertGenotypeToInt(landscapes[t].getGlobalPeak())
            best_prob = 1.0
            best_steerer = -1
            for s in range(len(limit_matrices)):
                dist = np.array(init_pop * limit_matrices[s] * limit_matrices[t])[0]
                if dist[global_peak_index] < best_prob:
                    best_prob = dist[global_peak_index]
                    best_steerer = s
            best_steerers.append((landscapes[s], best_prob))

            print landscapes[t].name+":", landscapes[best_steerer].name, best_prob
        return best_steerers

    #============================================================================#
    # Determines for each drug the ordered pair of steering drugs which minimizes
    # the probability that evolution proceeds to the highest fitness peak.
    #============================================================================#
    def highPeakBestDouble(init_pop):
        best_steerers = []
        for t in range(len(limit_matrices)):
            global_peak_index = convertGenotypeToInt(landscapes[t].getGlobalPeak())
            best_prob = 1.0
            best_s1 = -1
            best_s2 = -1
            for s1 in range(len(limit_matrices)):
                for s2 in range(len(limit_matrices)):
                    dist = np.array(init_pop * limit_matrices[s1] * \
                                    limit_matrices[s2] * limit_matrices[t])[0]
                    if dist[global_peak_index] < best_prob:
                        best_prob = dist[global_peak_index]
                        best_s1 = s1
                        best_s2 = s2

            print landscapes[t].name+":", landscapes[best_s1].name+\
                  " ---> "+landscapes[best_s2].name+", ", best_prob
```

```
            best_steerers.append((landscapes[best_s2], landscapes[best_s2], best_prob))
        return best_steerers

    #============================================================================#
    # Determines for each drug the ordered triple of steering drugs which minimizes
    # the probability that evolution proceeds to the highest fitness peak.
    #============================================================================#
    def highPeakBestTriple(init_pop):
        best_steerers = []
        for t in range(len(limit_matrices)):
            global_peak_index = convertGenotypeToInt(landscapes[t].getGlobalPeak())
            best_prob = 1.0
            best_s1 = -1
            best_s2 = -1
            best_s3 = -1
            for s1 in range(len(limit_matrices)):
                for s2 in range(len(limit_matrices)):
                    for s3 in range(len(limit_matrices)):
                        dist = np.array(init_pop * limit_matrices[s1] * limit_matrices[s2]
                                        limit_matrices[s3] * limit_matrices[t])[0]
                        if dist[global_peak_index] < best_prob:
                            best_prob = dist[global_peak_index]
                            best_s1 = s1
                            best_s2 = s2
                            best_s3 = s3


            print landscapes[t].name+":", landscapes[best_s1].name+" ---> "+ \
                  landscapes[best_s2].name+" ---> "+landscapes[best_s3].name\
                  +", ", best_prob
            best_steerers.append((landscapes[best_s1], landscapes[best_s2],
                                  landscapes[best_s3], best_prob))

        return best_steerers
```

In [428]:
```
%%capture
#Delete above line to print results
#============================================================================#
# Results (Table 3.3)
#============================================================================#
print "Results"
print "Final Drug: Steering combo (P(highest peak))"
print
best_singles = highPeakBestSingle(init_pop)
print
best_doubles = highPeakBestDouble(init_pop)
print
best_triple = highPeakBestTriple(init_pop)
```

The following code produces the results presented in Table 3.4, indicating the best sequence of 1, 2 or 3 drugs to steer evolution to maximise the likelihood of finding the least resistant genotype.

In [41]:
```
#============================================================================#
# The best steering combinations are determined by an exhaustive search of all
# sequences of 1, 2 or 3 drugs.
#============================================================================#

#============================================================================#
# Determines for each drug the probability of ending at the lowest fitness peak
# genotype when starting from the initial distribution init_pop.
#============================================================================#
def probLowestPeak(init_pop):
    lowest_pgs_probs = []
    for t in range(len(limit_matrices)):
        dist = np.array(init_pop * limit_matrices[t])[0]
        lpg = landscapes[t].getLowestFitnessPeak()
        lowest_pgs_probs.append(dist[convertGenotypeToInt(lpg)])
        print landscapes[t].name+":", lowest_pgs_probs[t]

    return lowest_pgs_probs

#============================================================================#
# Determines for each drug the single steering drug which maximizes the probability
# that evolution proceeds to the lowest fitness peak.
#============================================================================#
def lowestPeakBestSingle(init_pop):
    best_steerers = []
    for t in range(len(limit_matrices)):
        lowest_peak_index = convertGenotypeToInt(\
                            landscapes[t].getLowestFitnessPeak())

        best_steerer = -1
        best_prob = 0.0

        #For each possible steerer
        for s in range(len(limit_matrices)):
            dist = np.array(init_pop * limit_matrices[s] * limit_matrices[t])[0]
            if dist[lowest_peak_index] > best_prob:
                best_prob = dist[lowest_peak_index]
                best_steerer = s
        best_steerers.append((landscapes[best_steerer], best_prob))

        print landscapes[t].name+":", best_steerers[t][0].name, best_steerers[t][1]

    return best_steerers

#============================================================================#
```

```
# Determines for each drug the ordered pair of steering drugs which maximizes
# the probability that evolution proceeds to the lowest fitness peak.
#================================================================================#
def lowestPeakBestDouble(init_pop):
    best_steerers = []
    for t in range(len(limit_matrices)):
        lowest_peak_index = convertGenotypeToInt(\
                            landscapes[t].getLowestFitnessPeak())

        best_s1 = -1
        best_s2 = -1
        best_prob = 0.0
        for s1 in range(len(limit_matrices)):
            for s2 in range(len(limit_matrices)):
                dist = np.array(init_pop * limit_matrices[s1] * \
                                limit_matrices[s2] * limit_matrices[t])[0]
                if dist[lowest_peak_index] > best_prob:
                    best_prob = dist[lowest_peak_index]
                    best_s1 = s1
                    best_s2 = s2
        best_steerers.append((landscapes[best_s1], landscapes[best_s2], best_prob))
        print landscapes[t].name+":  ", landscapes[best_s1].name+" ---> "+ \
              landscapes[best_s2].name+", ", best_prob

    return best_steerers

#================================================================================#
# Determines for each drug the ordered triple of steering drugs which maximizes
# the probability that evolution proceeds to the lowest fitness peak.
#================================================================================#
def lowestPeakBestTriple(init_pop):
    best_steerers = []
    for t in range(len(limit_matrices)):
        lowest_peak_index = convertGenotypeToInt(\
                            landscapes[t].getLowestFitnessPeak())

        best_s1 = -1
        best_s2 = -1
        best_s3 = -1
        best_prob = 0.0
        for s1 in range(len(limit_matrices)):
            for s2 in range(len(limit_matrices)):
                for s3 in range(len(limit_matrices)):
                    dist = np.array(init_pop * limit_matrices[s1]
                                    * limit_matrices[s2] * limit_matrices[s3] *
                                    limit_matrices[t])[0]
                    if dist[lowest_peak_index] > best_prob:
                        best_prob = dist[lowest_peak_index]
```

```
                                best_s1 = s1
                                best_s2 = s2
                                best_s3 = s3
                    best_steerers.append((landscapes[best_s1], landscapes[best_s2],\
                                          landscapes[best_s3], best_prob))
                    print landscapes[t].name+":  ", landscapes[best_s1].name+" ---> "+ \
                          landscapes[best_s2].name+" ---> "+landscapes[best_s3].name+ \
                          ", ", best_prob

                return best_steerers
```

```
In [430]: %%capture
          #Delete above line to print results
          #============================================================================#
          # Results (Table 3.4)
          #============================================================================#
          print "Results"
          print "Final Drug: Steering combo (P(lowest peak))"
          print
          best_singles = lowestPeakBestSingle(init_pop)
          print
          best_doubles = lowestPeakBestDouble(init_pop)
          print
          best_triple = lowestPeakBestTriple(init_pop)
```

The following produces the results of Table 3.5, wherein all possible steering sequences of length 1, 2 and 3 are considered to determined the efficacy (or dangers) of **random drug ordering**.

```
In [44]: #============================================================================#
         # Exhaustive exploration of all drug sequences of length 1,2 or 3.
         #============================================================================#

         #============================================================================#
         # Determines the percentage of single steering drugs that decrease
         # and increase the probability of evolution to the highest peak of the landscape.
         #============================================================================#
         def singleSteererPercentages(init_pop, allowSAM = True):

             steering_stats = []

             overall_better = 0
             overall_worse = 0
             overall_same = 0

             for t in range(len(limit_matrices)):
                 if t!=10 or allowSAM:

                     better = 0
```

```
            worse = 0
            same = 0

            peak_genotype = landscapes[t].getGlobalPeak()
            peak_genotype_index = convertGenotypeToInt(peak_genotype)
            peak_fitness = landscapes[t].getFitness(peak_genotype)

            for s in range(len(limit_matrices)):
                steered_pop = (init_pop * limit_matrices[s] *
                               limit_matrices[t]).tolist()[0]
                straight_pop = (init_pop * limit_matrices[t]).tolist()[0]

                if np.abs(steered_pop[peak_genotype_index] - \
                          straight_pop[peak_genotype_index]) < 10**-15:

                    same+=1

                elif steered_pop[peak_genotype_index] > \
                     straight_pop[peak_genotype_index]:

                    worse+=1

                elif steered_pop[peak_genotype_index] < \
                     straight_pop[peak_genotype_index]:

                    better+=1

            steering_stats.append((better,worse,same))


            pc_better = 100*float(better) / (better+worse+same)
            pc_worse  = 100*float(worse) / (better+worse+same)
            print landscapes[t].name+", better:", str(better)+ \
            " ("+str(pc_better)+"%)", ", worse:", str(worse)+" ("+str(pc_worse)+"%)"

            overall_better+=better
            overall_worse+=worse
            overall_same+=same
            pc_ov_better = 100*float(overall_better) / \
                           (overall_better+overall_worse+overall_same)
            pc_ov_worse = 100*float(overall_worse) / \
                          (overall_better+overall_worse+overall_same)


    print "Overall, better:", str(overall_better)+\
          " ("+str(pc_ov_better)+"%)", ", worse:", str(overall_worse)+\
          " ("+str(pc_ov_worse)+"%)"
```

```
    return steering_stats

#=============================================================================#
# Determines the percentage of ordered pairs of steering drugs that decrease
# and increase the probability of evolution to the highest peak of the landscape.
#=============================================================================#
def doubleSteererPercentages(init_pop, allowSAM = True):
    steering_stats = []

    overall_better = 0
    overall_worse = 0
    overall_same = 0

    for t in range(len(limit_matrices)):
        if t!=10 or allowSAM:

            better = 0
            worse = 0
            same = 0

            peak_genotype = landscapes[t].getGlobalPeak()
            peak_genotype_index = convertGenotypeToInt(peak_genotype)
            peak_fitness = landscapes[t].getFitness(peak_genotype)

            for k1 in range(len(limit_matrices)):
                for k2 in range(len(limit_matrices)):
                    steered_pop = (init_pop * limit_matrices[k1] *
                                   limit_matrices[k2] *
                                   limit_matrices[t]).tolist()[0]
                    straight_pop = (init_pop * limit_matrices[t]).tolist()[0]


                    if np.abs(steered_pop[peak_genotype_index] -\
                              straight_pop[peak_genotype_index]) < 10**-15:

                        same+=1

                    elif steered_pop[peak_genotype_index] >\
                         straight_pop[peak_genotype_index]:

                        worse+=1

                    elif steered_pop[peak_genotype_index] <\
                         straight_pop[peak_genotype_index]:

                        better+=1

            steering_stats.append((better,worse,same))
```

```
            pc_better = 100*float(better) / (better+worse+same)
            pc_worse  = 100*float(worse) / (better+worse+same)
            print landscapes[t].name+", better:", str(better)+ \
                  " ("+str(pc_better)+"%)", ", worse:", str(worse)+\
                  " ("+str(pc_worse)+"%)"

            overall_better+=better
            overall_worse+=worse
            overall_same+=same
            pc_ov_better = 100*float(overall_better) / \
                           (overall_better+overall_worse+overall_same)
            pc_ov_worse = 100*float(overall_worse) / \
                          (overall_better+overall_worse+overall_same)



    print "Overall, better:", str(overall_better)+\
          " ("+str(pc_ov_better)+"%)", ", worse:", str(overall_worse)+\
          " ("+str(pc_ov_worse)+"%)"
    return steering_stats


#=============================================================================#
# Determines the percentage of ordered triples of steering drugs that decrease
# and increase the probability of evolution to the highest peak of the landscape.
#=============================================================================#
def tripleSteererPercentages(init_pop, allowSAM = True):
    steering_stats = []

    overall_better = 0
    overall_worse = 0
    overall_same = 0

    for t in range(len(limit_matrices)):
        if t!=10 or allowSAM:

            better = 0
            worse = 0
            same = 0

            peak_genotype = landscapes[t].getGlobalPeak()
            peak_genotype_index = convertGenotypeToInt(peak_genotype)
            peak_fitness = landscapes[t].getFitness(peak_genotype)

            for k1 in range(len(limit_matrices)):
                for k2 in range(len(limit_matrices)):
```

```
                    for k3 in range(len(limit_matrices)):

                        steered_pop = (init_pop * limit_matrices[k1] *
                                       limit_matrices[k2] *
                                       limit_matrices[k3] *
                                       limit_matrices[t]).tolist()[0]
                        straight_pop = (init_pop * limit_matrices[t]).tolist()[0]

                        if np.abs(steered_pop[peak_genotype_index] -\
                                  straight_pop[peak_genotype_index])<10**-15:

                            same+=1

                        elif steered_pop[peak_genotype_index] >\
                             straight_pop[peak_genotype_index]:

                            worse+=1

                        elif steered_pop[peak_genotype_index] <\
                             straight_pop[peak_genotype_index]:

                            better+=1

                steering_stats.append((better,worse,same))


                pc_better = 100*float(better) / (better+worse+same)
                pc_worse  = 100*float(worse) / (better+worse+same)
                print landscapes[t].name+", better:", str(better)+\
                      " ("+str(pc_better)+"%)", ", worse:", str(worse)+\
                      " ("+str(pc_worse)+"%)"

                overall_better+=better
                overall_worse+=worse
                overall_same+=same
                pc_ov_better = 100*float(overall_better) / \
                (overall_better+overall_worse+overall_same)
                pc_ov_worse = 100*float(overall_worse) / \
                (overall_better+overall_worse+overall_same)



    print "Overall, better:", str(overall_better)+\
          " ("+str(pc_ov_better)+"%)", ", worse:", str(overall_worse)+\
          " ("+str(pc_ov_worse)+"%)"
    return steering_stats
```

```
In [432]: %%capture
          #Delete above line to print results
          #============================================================================#
          # Results (Table 3.5)
          #============================================================================#
          print "Results"
          print
          print "Singles"
          singleSteererPercentages(init_pop)
          print
          print "Doubles"
          doubleSteererPercentages(init_pop)
          print
          print "Triples"
          tripleSteererPercentages(init_pop);
```

2.1 Accounting for Variance in Fitness Measurements

Typically, measurements of growth rate (which is used as a proxy for fitness here) are subject to substantial variance owing both to intrinsic biological noise and extrinsic experimental noise introduced between replicates. Unfortunately, despite reporting fitness values as the average of 12 replicates, neither the raw values nor measurements of variance are reported by Mira et al.

Here, we test the sensitivity of our results to variance in the values that comprise the fitness values. We begin by implementing a function that resamples all of the fitness values uniformly from an interval around the reported mean.

```
In [433]: def resample(landscape, variance=0.3):
              vals = landscape.landscape
              fl = FitnessLandscape(map(lambda x : np.random.uniform((1-variance)*x,
                                                                     (1+variance)*x),vals),
                                    name=landscape.name)
              return fl
```

Central to our results on evolutionary steering and the non-repeatability is the existence of multi-peaked landscapes, we begin by exploring how the number of peaks varies in each landscape under resampling of the fitness values.

```
In [436]: S=10 #1000 #in the paper/thesis
          fig = plt.figure(figsize=(18,12))
          for k, ls in enumerate(landscapes):
              plt.subplot(3,5,k+1)
              samples=S
              peaks = []
              landscape=landscapes[k]
              mean_peaks = len(landscape.getPeaks())
              for _ in range(samples):
                  fl = resample(landscape)
                  peaks.append(len(fl.getPeaks()))
```

```
        plt.hist(peaks, bins=np.arange(0,10.5,1.0), width=0.8, normed=True)
        plt.yticks(size=12)
        plt.xticks(np.arange(0.4,10.3,1.0), np.arange(0,10,1), size=12)
        plt.ylabel('Frequency', size=22)
        plt.xlabel('Peaks', size=22)
        plt.title(labs[k], size=24)
        plt.xlim(0.7,10.1)
        plt.ylim(0.0,0.6)
        plt.annotate("Peaks in \noriginal data: " + str(mean_peaks), xy=(8, 0.4),
            xytext=(0.5, .82), textcoords='axes fraction', zorder=100, size=14)
plt.suptitle('Landscape Peaks after Resampling', size=28)
plt.tight_layout(rect=[0, 0.03, 1, 0.95])
```

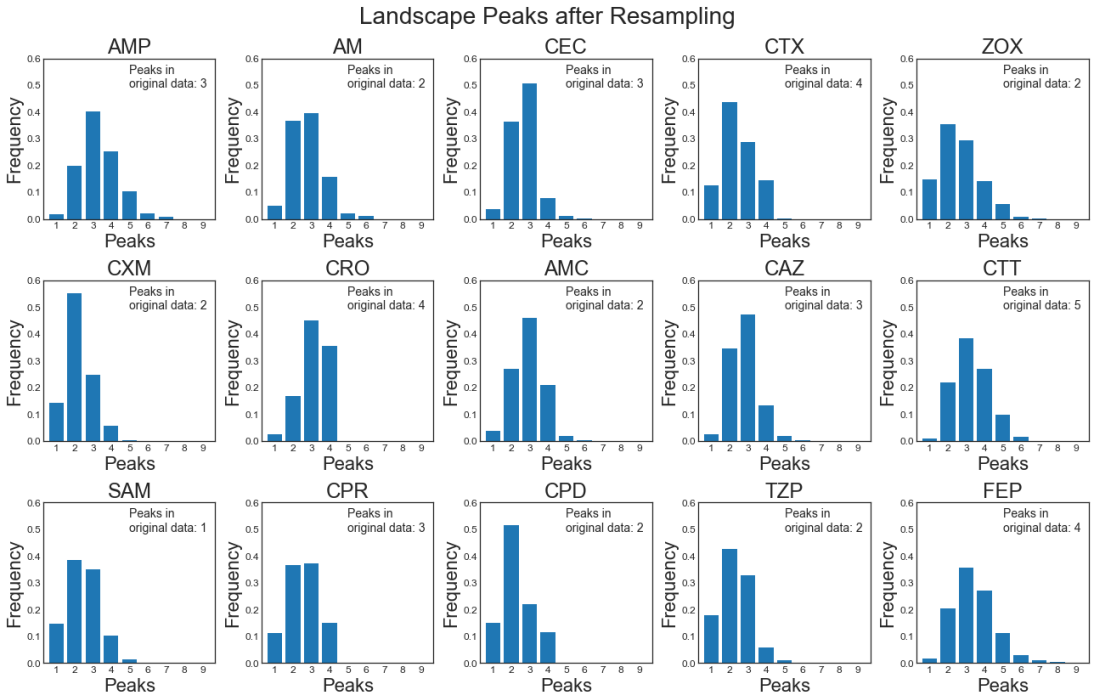

2.1.1 Collateral Sensitivity and Variance

Central to our results regarding collateral sensitivity is the phenomenon whereby a drug pair reported as collaterally sensitive from a single replicate of experimental evolution can indeed exhibit cross resistance. Here, we validate the robustness of this result to variance in the derived fitness landscape.

We assume that the variance in fitness values is biologically intrinsic, such that two separate biological replicates will occur with slightly different fitness landscapes (determined by resampling the landscape each time). We begin by computing the likelihood of cross-resistance under these assumptions

In [121]:
```
#============================================================================#
# Computes the likelihood of CR from a random drug pair where the landscape is
# resampled each time. This calculation works as follows:
# Repeat S times:
#    1. Sample a drug pair at random.
#    2. Resample the landscapes from these two drugs
#    3. Compute the likelihood of cross-resistance
# Then:
#    Average over the CR likelihoods.
#============================================================================#
def resampled_CR_prob(S):
    tot = 0.
    for _ in range(S):
        d1, d2 = np.random.choice(range(15), 2, replace=False)
        l1 = resample(landscapes[d1])
        l2 = resample(landscapes[d2])

        limit1 = limitMatrix(buildTransitionMatrix(l1))

        wt_f = l1.getFitness([0,0,0,0])
        ev_pop = init_wt * limit1
        ev_pop = np.array(ev_pop)[0]
        temp=0.
        for k in range(len(ev_pop)):
            ev_f = l2.getFitness(convertIntToGenotype(k,4))
            if ev_f > wt_f:
                temp+=ev_pop[k]
        tot+=temp

    return tot/S

#============================================================================#
# Computes the likelihood of CR for a pair of drugs reported as CS in a previous
# evolutionary replicate. This value is calculated as follows:
# Repeat S times:
#   1. Resample all landscapes, build a table of collateral response
#   2. Choose uniformly at random a drug pair from the table which exhibits CS
#   3. Resample the landscape for both drugs
#   4. Compute the likelihood of CR in thew new landscapes
# Then:
#   Average over the CR likelihoods
#============================================================================#

#============================================================================#
# Simulates a single instance of determining collateral sensitivity
# with resampled landscapes.
#============================================================================#
def resampled_col_sensitivity(d1, d2, init_pop, relative=True):
```

```
    limit_d1 = limitMatrix(buildTransitionMatrix(resample(landscapes[d1])))
    pop_dist = init_pop * limit_d1
    pop_dist = np.array(pop_dist)[0]
    peak = np.random.choice(np.array([i for i in range(16)]), p=pop_dist)

    landscape2 = resample(landscapes[d2])

    col_sens = np.log2(landscape2.getFitness(convertIntToGenotype(peak,4)) \
                       / landscape2.getFitness(convertIntToGenotype(0,4)))
    return col_sens

def resampled_generateCSN(init_pop):
    network = [[resampled_col_sensitivity(d1,d2,init_pop)\
                for d1 in range(len(landscapes))]
               for d2 in range(len(landscapes))]
    return network

#============================================================================#
# Calculates the likelihood of CR following a positive CS hit.
#============================================================================#
def resampled_CSCR_likelihoods(S):
    pr = 0.0
    for _ in range(S):
        #generate a random csn
        csn = resampled_generateCSN(init_wt)
        #Pick a pair with CS are random
        flag = True
        while flag:
            d1 = random.randint(0,14)
            d2 = random.randint(0,14)
            if csn[d1][d2] < 0.0:
                flag = False

        #find the probability of clr.
        prob_res = 0.0
        l1 = resample(landscapes[d1])
        l2 = resample(landscapes[d2])
        limit1 = limitMatrix(buildTransitionMatrix(l1))
        wt_f = l1.getFitness([0,0,0,0])
        ev_pop = init_wt * limit_matrices[d1]
        ev_pop = np.array(ev_pop)[0]
        for k in range(len(ev_pop)):
            ev_f = l2.getFitness(convertIntToGenotype(k,4))
            if ev_f > wt_f:
                prob_res += ev_pop[k]

        pr += prob_res
```

```
        return pr/S

    S=10 #10000 in the thesis
    print "Probability of CR", resampled_CR_prob(S) #approx 0.54
    print "Probability of CR following CS", resampled_CSCR_likelihoods(S) #Approx 0.53
```

```
Probability of CR 0.538365917567
Probability of CR following CS 0.531055555556
```

Next we sample an extended collateral sensitivity likelihood table as follows. For each drug pair, we resample the landscapes 100 times and determine the likelihood of CS for each resampled pair. The average of these probabilities is the *extended collateral sensitivity likelihood*.

```
In [454]: #============================================================================#
          # By first resampling the drug landscapes, returns the probability that evolution
          # under a first results in a genotype with fitness under a
          # second that is lower than the WT fitness.
          #============================================================================#
          def resampled_prob_better(d1,d2,init_pop):
              l1 = resample(landscapes[d1])
              if d2!=d1: #Do not resample twice if they are the same drug
                  l2 = resample(landscapes[d2])
              else:
                  l2 = l1

              limit_d1 = limitMatrix(buildTransitionMatrix(l1))
              pop_dist = init_pop * limit_d1
              pop_dist = np.array(pop_dist)[0]

              prob = 0.
              wt_f = l2.getFitness([0,0,0,0])
              for i in range(len(pop_dist)):
                  if l2.getFitness(convertIntToGenotype(i,4)) < wt_f:
                      prob+=pop_dist[i]
              return prob

          #============================================================================#
          # Extended Collateral Sensitivity Likelihood matrix
          #============================================================================#
          def resampled_prob_matrix(init_pop, S=10):
              mat = [[np.mean([resampled_prob_better(d1,d2,init_pop) for _ in range(S)]) \
                      for d1 in range(len(landscapes))] for d2 in range(len(landscapes))]
              return mat

          S=10 #1000 #in the thesis.
          e_csl  = resampled_prob_matrix(init_wt, S)
```

```
In [455]: #Displaying the resampled CSL and a thresholded version - Uncomment for figures
          # show_CSL(e_csl)
```

```
# show_CSL(cut_off_probs_CSL(e_csl, 0.9))
# show_CSL(cut_off_probs_CSL(e_csl, 0.8))
show_CSL(cut_off_probs_CSL(e_csl, 0.7))
```

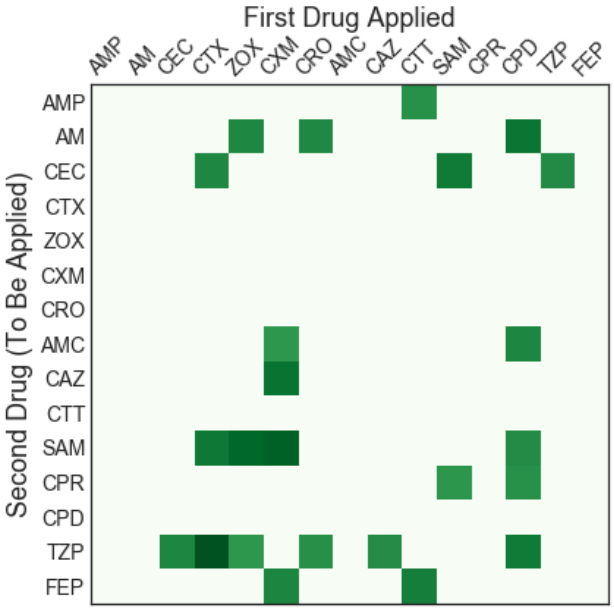

Finally, we can explore the change in CSL value that arises from resampling by considering the difference between the two CSL matrices.

```
In [456]: residuals = np.matrix(csl) - np.matrix(e_csl)
          fig = plt.figure(figsize  = (8,6.5))
          ax = fig.add_subplot(111)
          cax=ax.matshow(residuals, cmap = 'RdBu_r', vmin = -1.0, vmax = 1.0)

          cbar = fig.colorbar(cax)
          cbar.set_label('Residual (CSL - rCSL)', size=20)

          ax.set_xticks(range(len(labs)))
          ax.set_yticks(range(len(labs)))
          ax.set_xticklabels(labs, rotation=45, size='14')
          ax.set_yticklabels(labs, size='14')
          ax.axis('image')

          ax.set_xlabel('First Drug Applied', size='18')
          ax.xaxis.set_label_position('top')
          ax.set_ylabel('Second Drug (To Be Applied)', size='18');
```

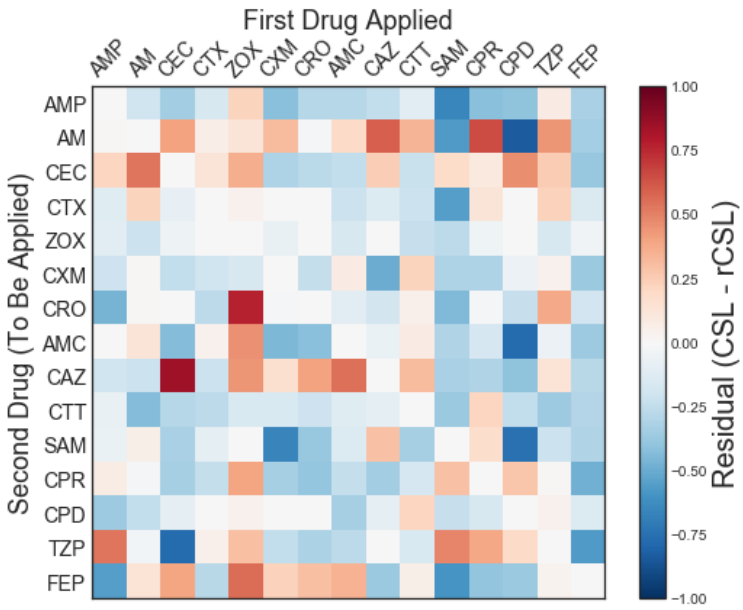

2.1.2 Steering and Variance

To conclude this analysis of the effects of fitness variance on the predictions of our steering model, we now test the efficacy of the previously derived optimal steering strategies on resampled landscapes.

```
In [457]: %%capture
          bests_hp = highPeakBestTriple(init_pop)
          bests_lp = lowestPeakBestTriple(init_pop)

In [458]: def resample_steered_hp(f_ix, init_pop):
              f_ls = resample(landscapes[f_ix])
              a1, a2, a3, p = bests_hp[f_ix]

              if a1.name != f_ls.name:
                  a1 = resample(a1)
              else:
                  a1 = f_ls

              if a2.name == a1.name:
                  a2 = a1
              elif a2.name == f_ls.name:
                  a2 = f_ls
              else:
                  a2 = resample(a2)

              if a3.name == a2.name:
```

```
            a3 = a2
        elif a3.name == a1.name:
            a3 = a1
        elif a3.name == f_ls.name:
            a3 = f_ls
        else:
            a3 = resample(a3)

        p1 = limitMatrix(buildTransitionMatrix(a1))
        p2 = limitMatrix(buildTransitionMatrix(a2))
        p3 = limitMatrix(buildTransitionMatrix(a3))
        pf = limitMatrix(buildTransitionMatrix(f_ls)) #Final drug

        dist = np.array(init_pop * p1 * p2 * p3 * pf)[0]
        return dist[convertGenotypeToInt(f_ls.getGlobalPeak())]

    # S=10000
    # for ix in range(15):
    #     dist = [resample_steered_hp(ix, init_pop) for _ in range(S)]
    #     name = landscapes[ix].name
    #     np.save('./resampled_steering/hp_'+name+'.npy', dist)
```

```
In [459]: def resample_steered_lp(f_ix, init_pop):
        f_ls = resample(landscapes[f_ix])
        a1, a2, a3, p = bests_lp[f_ix]

        if a1.name != f_ls.name:
            a1 = resample(a1)
        else:
            a1 = f_ls

        if a2.name == a1.name:
            a2 = a1
        elif a2.name == f_ls.name:
            a2 = f_ls
        else:
            a2 = resample(a2)

        if a3.name == a2.name:
            a3 = a2
        elif a3.name == a1.name:
            a3 = a1
        elif a3.name == f_ls.name:
            a3 = f_ls
        else:
            a3 = resample(a3)

        p1 = limitMatrix(buildTransitionMatrix(a1))
```

```
        p2 = limitMatrix(buildTransitionMatrix(a2))
        p3 = limitMatrix(buildTransitionMatrix(a3))
        pf = limitMatrix(buildTransitionMatrix(f_ls)) #Final drug

        dist = np.array(init_pop * p1 * p2 * p3 * pf)[0]
        return dist[convertGenotypeToInt(f_ls.getLowestFitnessPeak())]

    # S=10000
    # for ix in range(15):
    #     dist = [resample_steered_lp(ix, init_pop) for _ in range(S)]
    #     name = landscapes[ix].name
    #     np.save('./resampled_steering/lp_'+name+'.npy', dist)
```

In [462]:
```
dists = [np.load(f) for f in glob.glob('./resampled_steering/hp_*')]
dists_l = [np.load(f) for f in glob.glob('./resampled_steering/lp_*')]
ps_hp = map(lambda x : x[-1], bests_hp)
ps_lp = map(lambda x : x[-1], bests_lp)
```

In [463]:
```
plt.figure(figsize=(25,9))
plt.subplot(121)
sns.boxplot(data=dists, notch=True)

for k, p in enumerate(ps_hp):
    plt.plot([k],[p], marker='H', c='k', markersize=12)

plt.xticks(range(15), labs, rotation=45, size=16)
plt.yticks(size=16)
plt.xlabel('Final Drug', size=20)
plt.ylabel('''Probability of Evolution to Highest Peak\n
(following optimal steering strategy)''', size=20)
plt.title('Avoiding Highest Peak', size=24)

plt.subplot(122)
sns.boxplot(data=dists_l, notch=True)

for k, p in enumerate(ps_lp):
    plt.plot([k],[p], marker='H', c='k', markersize=12)

plt.xticks(range(15), labs, rotation=45, size=16)
plt.yticks(size=16)
plt.xlabel('Final Drug', size=20)
plt.ylabel('''Probability of Evolution to Lowest Peak
           \n(following optimal steering strategy)''', size=20)
plt.title('Steering to Lowest Peak', size=24)

plt.suptitle('Robustness of Steering to Noise', size=40)
plt.tight_layout(rect=[0, 0.03, 1, 0.9])
plt.subplots_adjust(wspace=0.2)
```

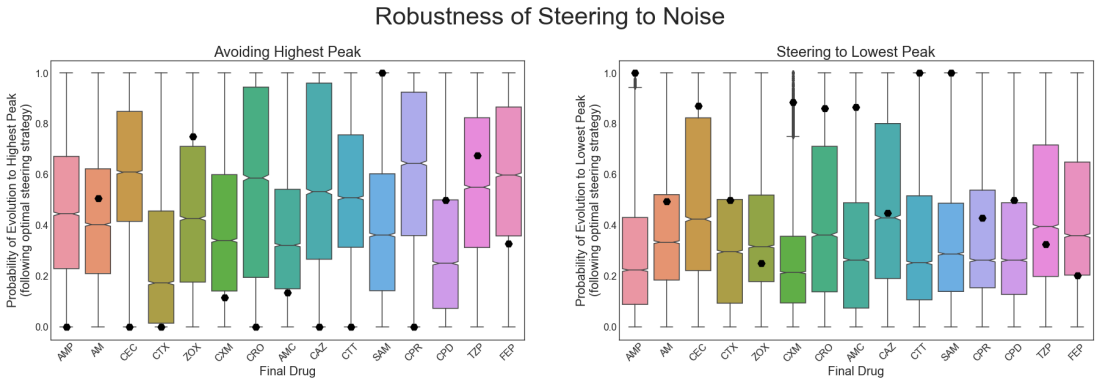
Robustness of Steering to Noise
Avoiding Highest Peak
Probability of Evolution to Highest Peak (following optimal steering strategy)
1.0
0.8
0.6
0.4
0.2
0.0
AMP
AM
CEC
CTX
ZOX
CXM
CRO
AMC
CAZ
CTT
SAM
CPR
CPD
TZP
FEP
Final Drug
Steering to Lowest Peak
Probability of Evolution to Lowest Peak (following optimal steering strategy)
1.0
0.8
0.6
0.4
0.2
0.0
AMP
AM
CEC
CTX
ZOX
CXM
CRO
AMC
CAZ
CTT
SAM
CPR
CPD
TZP
FEP
Final Drug

Chapter4

January 19, 2018

1 Chapter 4 - Modelling Non–Genetic Heterogeneity - Molecular Switches

Presented here is the python code used to generate the results in the Chapter 4 of the thesis. For brevity much of the code is contained within external Python files. This notebook serves as a guide to interact with these files.

The code used here is partially adapted from that used to generate the results in the following publication:

Nichol, Daniel, et al. "Stochasticity in the genotype-phenotype map: Implications for the robustness and persistence of bet-hedging." Genetics 204.4 (2016): 1523-1539.

1.1 Prerequisites

In [1]:
```
from copy import copy
from copy import deepcopy
import glob
import math
from random import random
import sys

import matplotlib
import matplotlib.pyplot as plt
from matplotlib import rc
from mpl_toolkits.axes_grid1 import Grid
import numpy as np
import seaborn as sns

import seaborn as sbn
sbn.set(rc={'image.cmap': 'cubehelix'})
sbn.set_style("white")

print "Python version: ", sys.version
print "numpy version: ", np.__version__
print "matplotlib version: ", matplotlib.__version__
```

```
        #Plotting style variables
        linestyles = ['-', '--', '-.',':']
        linecols = [sns.xkcd_rgb["pale red"], sns.xkcd_rgb["denim blue"],
                    sns.xkcd_rgb["medium green"],sns.xkcd_rgb["black"]]
        switchlabs=["DC Switch", "DC Switch - X Biased",
                    "DC Switch - Y Biased", "AM Switch"]

        %matplotlib inline
```

```
Python version:  2.7.13 |Continuum Analytics, Inc.| (default, Dec 20 2016, 23:05:08)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.57)]
numpy version:  1.11.3
matplotlib version:  2.0.2
```

1.2 Reaction Networks (General Case)

We begin by defining a general purpose reaction networks class that can be used to implement molecular-switch-driven bet-hedging. Note that because this implementation is purposely general it will not be optimised for a specific switch. The necessary code is available in ReactionNetworks.py where a DirectCompetition and ApproximateMajority network are defined.

```
In [2]: import reactionNetwork as rn
```

```
In [3]: DC_net = rn.DirectCompetition([1.0,1.0], [30,30])
        DCx_net = rn.DirectCompetition([1.1,1.0], [30,30])
        DCy_net = rn.DirectCompetition([1.0,1.1], [30,30])
        AM_net = rn.ApproximateMajorityNetwork(2, [30,30])
        #Example
        print DC_net
```

```
x y --1.0--> x x
x y --1.0--> y y
```

```
In [4]: #============================================================================#
        # We can generate a stochastic plot for a network by as follows:
        # (Figure 4.2, second row, example)
        #
        # Stochastic simulations are determined by a Gillespie algorithm.
        #============================================================================#
        AM_net.generateStochasticPlot(0, 10)
```

1.3 Reaction Networks (Specific Cases)

In hedges.py we define optimised look-up tables for the DC, DCx, DCy and AM switches.

```
In [5]: import hedges
        g_max = 60
        DC_hedge = hedges.DC_hedge
        DC_hedge_x = hedges.get_DC_hedge(1.1,1.0)
        DC_hedge_y = hedges.get_DC_hedge(1.0,1.1)
        AM_hedge = hedges.get_AM_hedge(120)
```

```
In [6]: #=============================================================================#
        #Figure (4.2, final row, examples)
        #=============================================================================#
        sbn.set_style("white")
        fig = plt.figure(figsize=(10,3))
        ax = fig.add_subplot(121)
        seamap = sbn.cubehelix_palette(8, start=2, rot=0, dark=0, light=.95,
                                       reverse=True, as_cmap=True)
        dc_hm = [[DC_hedge(x,x+y) for x in range(1,62)] for y in range(1,62)]
        cax = ax.matshow(dc_hm, cmap='BuGn', origin='lower', vmax=1.0, vmin=0.0,
                         interpolation='spline16')
        fig.colorbar(cax)
        ax.xaxis.set_ticks_position('bottom')
```

```
cset = plt.contour(dc_hm,np.arange(0.1,1.,0.1),linewidths=.2,colors='k')
plt.clabel(cset,inline=True,fmt='%1.1f',fontsize=10)
plt.xlabel("x0", size=12)
plt.ylabel("y0", size=12, rotation=90)
ax = fig.add_subplot(122)
seamap = sbn.cubehelix_palette(8, start=2, rot=0, dark=0, light=.95,
                               reverse=True, as_cmap=True)
am_hm = [[AM_hedge(x,x+y) for x in range(1,60)] for y in range(1,60)]
cax = ax.matshow(am_hm, cmap='BuGn', origin='lower', vmax=1.0, vmin=0.0,
                 interpolation='spline16')
fig.colorbar(cax)
ax.xaxis.set_ticks_position('bottom')
cset = plt.contour(am_hm,np.arange(0.1,1.,0.1),linewidths=.2,colors='k')
plt.clabel(cset,inline=True,fmt='%1.1f',fontsize=10)
plt.xlabel("x0", size=12)
plt.ylabel("y0", size=12, rotation=90);
```

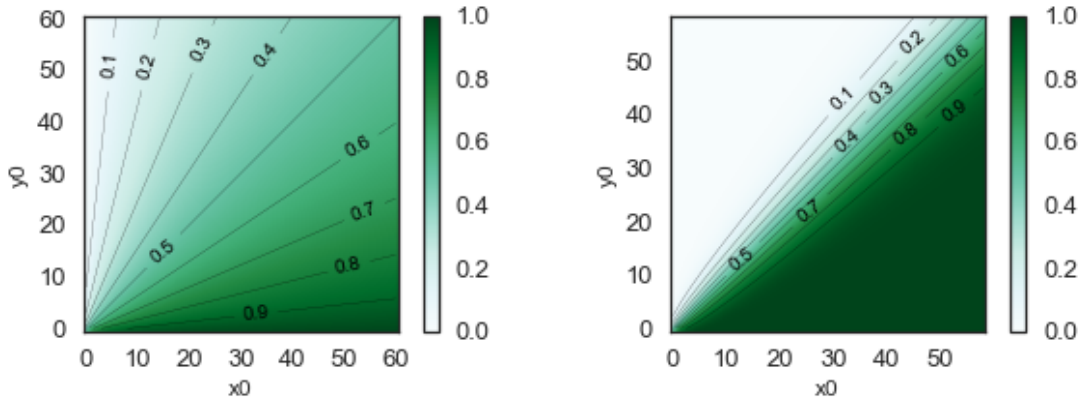

1.4 Closely Matching Abritrary p

In Figure 4.3 we demonstrate the relationsip between g_{max} and the maximum difference between $p \in [0,1]$ and $q(x_0, y_0)$ for $x_0 + y_0 < g_{max}$. The following code generates Figure 4.3.

```
In [7]: #=============================================================================#
        # Fig 4.3
        #=============================================================================#
        max_gm = 61
        diffs = [[],[],[],[]] #DC, DCx, DCy, AM
        DC = hedges.get_DC_hedge(1.0,1.0)
        DCx = hedges.get_DC_hedge(1.1,1.0)
        DCy = hedges.get_DC_hedge(1.0,1.1)
        for g_max in range(1, max_gm):
            pmdc, pmx, pmy,pmam = 0.0, 0.0 ,0.0, 0.0
```

```
    for g in range(0,g_max):
        pmdc = max(pmdc, DC(g+1,g_max)-DC(g, g_max))
        pmx  = max(pmx, DCx(g+1,g_max)-DCx(g, g_max))
        pmy  = max(pmy, DCy(g+1,g_max)-DCy(g, g_max))
        pmam = max(pmam, AM_hedge(g+1,g_max)-AM_hedge(g, g_max))

    diffs[0].append(0.5*pmdc)
    diffs[1].append(0.5*pmx)
    diffs[2].append(0.5*pmy)
    diffs[3].append(0.5*pmam)

plt.figure(figsize=(9,6))
plt.plot(range(max_gm-1),diffs[0], color=sbn.xkcd_rgb["pale red"],
         ls = linestyles[0], label="DC Switch")
plt.plot(diffs[1], color=sbn.xkcd_rgb["denim blue"],
         ls = linestyles[1], label="DC Switch - X Biased")
plt.plot(diffs[2], color=sbn.xkcd_rgb["medium green"],
         ls = linestyles[2], label = "DC Switch - Y Biased")
plt.plot(diffs[3], color=sbn.xkcd_rgb["black"],
         ls = linestyles[3], label = "AM Switch")

plt.xticks(range(max_gm-1),
           [str(0)]+[str(i) if i%5==0 else '' for i in range(2,max_gm)],
           rotation='horizontal')

plt.ylim(0,0.5)
plt.xlabel(r'$g_{max}$', fontsize=24)
plt.ylabel('Maximum Hedge Approximation Error')
plt.title('Approximating Arbitrary Hedging Probabilities')

plt.legend(bbox_to_anchor=(0., -0.2, 1., .102), loc=4,
    ncol= 4 , mode="expand", borderaxespad=0.)

plt.show()
```

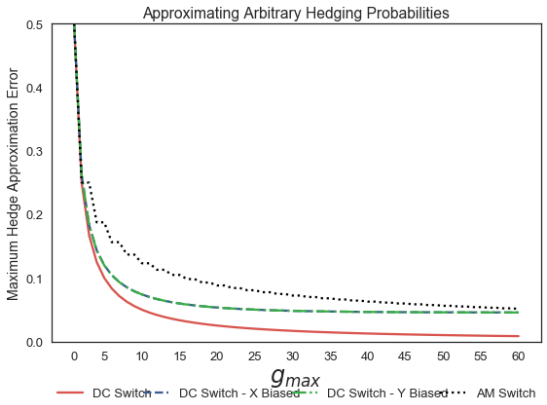

1.5 Redundancy - Section 4.3.3

In redundancy.py the extended DC networks from Figure 4.4(A) are defined. The dynamics of these reaction networks are as follows. Note that the 'switched' criterion for these switches is $x + x' = 0$ or $x + x' = 60$ and so in the following plots the final value of x can differ between stochastic realisations.

```
In [8]: import redundancy as rd
        #Note - these are not oscillatory
        rd.DCred.generateStochasticPlot('x', 10)
        rd.DCnoy.generateStochasticPlot('x',10)
```

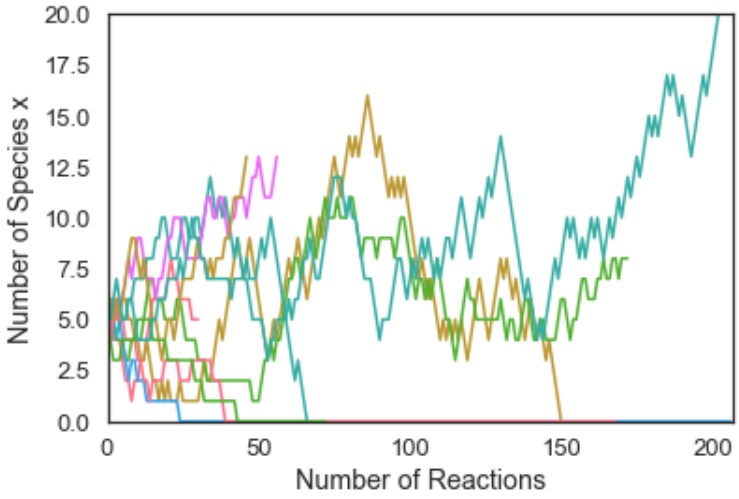

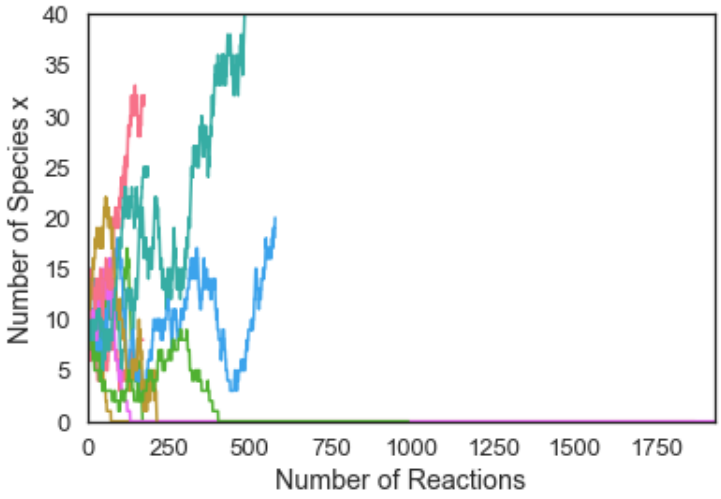

1.6 Invasion Dynamics - Section 4.5

The code in invasion.py contains the calculation of invasion probabilities from the population dynamics. This file also contains the code necessary to generate the associated figures.

In [9]:
```
#============================================================================#
# Section 4.5
#============================================================================#
import invasion as inv
#Parameters
w = [2.0,1.01] #Section 4.5.3

#Example use
M = inv.get_M(0.5,w) #population dynamics matrix for a 50/50 hedging population
print "The population dynamics matrix: ", M

#Properties of growth from the Perron-Frobenius theroem
growth_rate, stationary_dist = inv.get_lambda_v(M)
print "The exponential growth rate: ", growth_rate
print "The stationary phenotype distribution: ", stationary_dist.flatten()

# p*w[0] + (1-p)*w[1] (avg fitness, equal to the growth rate derived above)
print "Avg. fitness: ", inv.get_avg_fitness(0.5, w)

#The invasion probability for an invading bet-hedging population
p1, p2 = 0.5, 0.6
inv_p = inv.find_invasion_prob(p1, p2, w)
print "p2 invades p1 with likelihood: ", inv_p
inv_p = inv.find_invasion_prob(p2, p1, w)
print "p1 invades p2 with likelihood: ", inv_p
```

```
The population dynamics matrix:  [[ 1.     1.   ]
 [ 0.505  0.505]]
The exponential growth rate:  1.505
The stationary phenotype distribution:  [[ 0.66445183  0.33554817]]
Avg. fitness:  1.505
p2 invades p1 with likelihood:  0.111012149935
p1 invades p2 with likelihood:  0.0
```

Figures 4.5 and 4.6 are compiled from the above functions. The implementation details for the following methods are presented in invasion.py:

In [10]:
```
#============================================================================#
#Figure 4.5
#============================================================================#
```

```
sns.set_style('dark')
inv.pairwise_invasion_plot(w)
```

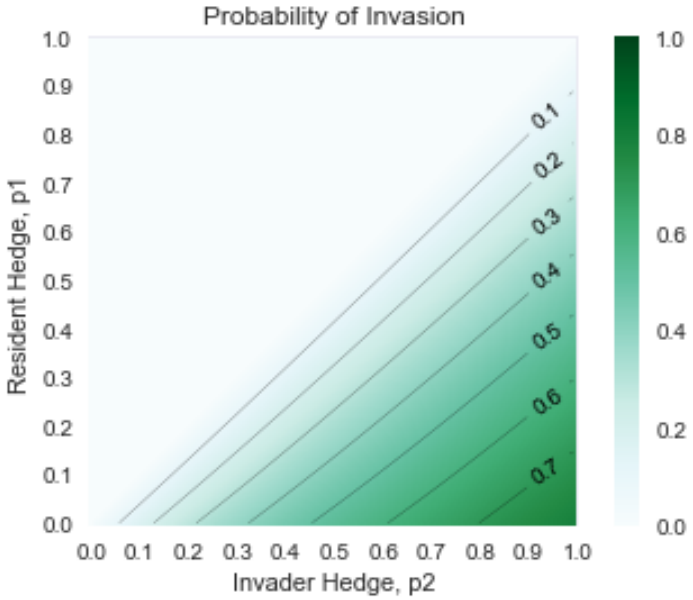

In [11]:
```
#=============================================================================#
#Fig 4.6(A)
#=============================================================================#
plt.figure(figsize=(10,10))
inv.g_vs_avgf_curve(60, w, [DC_hedge, DC_hedge_x, DC_hedge_y, AM_hedge],
                    switchlabs, linecols, linestyles) #Fig 4.6(A)
```

```
<matplotlib.figure.Figure at 0x12fc9da90>
```

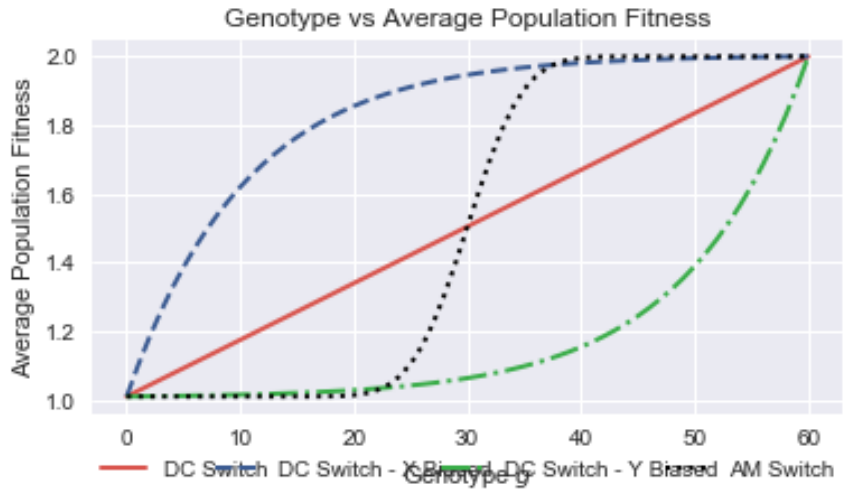

In [21]:
```
#=========================================================================#
#Fig 4.6(B)
#=========================================================================#
inv.diminishing_returns_plot(w,60,[DC_hedge, DC_hedge_x, DC_hedge_y, AM_hedge],
                             switchlabs,
                             linecols, linestyles)
```

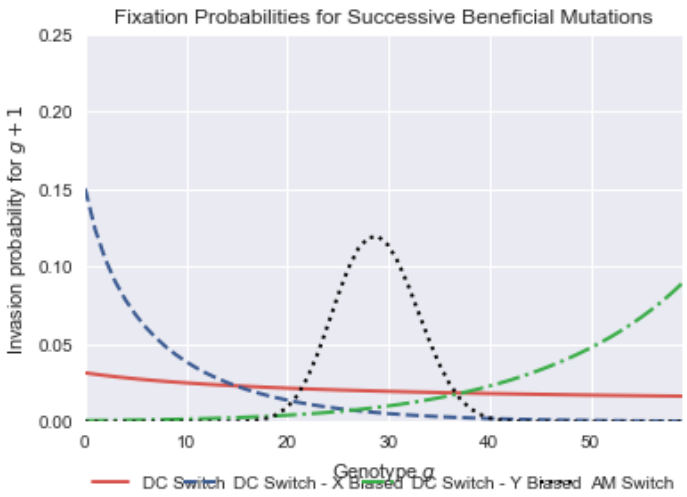

In [13]:
```
#=========================================================================#
#Fig 4.6(C) - AM example
#=========================================================================#
inv.genotype_invasion_plot(AM_hedge, w, 60, "AM Switch")
```

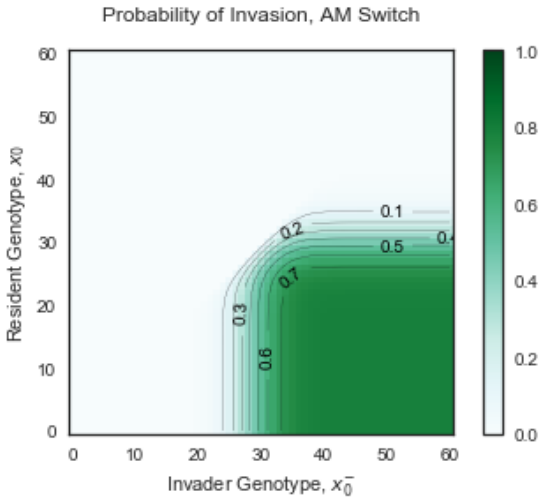

1.7 Evolutionary Convergence - Section 4.5.3

convergence.py contains the methods used to simulate evolutionary convergence.

```
In [14]: import convergence as cnv
         #============================================================================#
         #Example use - the expected convergence times in mutational events
         #
         # These values are determined by GetExpectedGenotype in Algorithm 2.
         #============================================================================#
         print cnv.expected_convergence_time(30, w, DC_hedge)
         print cnv.expected_convergence_time(7, w, DC_hedge_x)
         print cnv.expected_convergence_time(53, w, DC_hedge_y)
         print cnv.expected_convergence_time(30, w, AM_hedge)

3382.81556781
66800.1082291
195.818877031
2559425135.59


In [15]: #============================================================================#
         # The stochastic convergence curves are generated by the following method.
         # The results are saved in './Convergence.py'
         #============================================================================#
         cnv.convergence_curves(w,30) #Slow

In [22]: #============================================================================#
         # The curves can be plotted as follows (Fig 4.8(D)) (Uncomment for figs)
         #============================================================================#
         # cnv.convergence_plot_g('AM')
         # cnv.convergence_plot('AM')
```

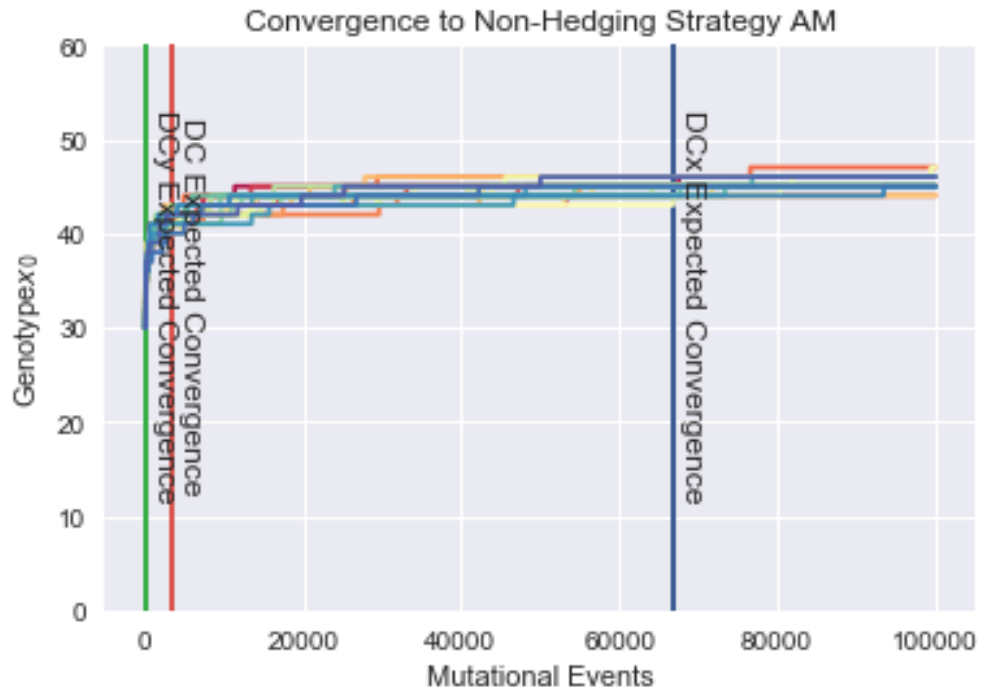
Convergence to Non-Hedging Strategy AM
Genotypex0
Mutational Events
DCy Expected Convergence
DC Expected Convergence
DCx Expected Convergence
0
10
20
30
40
50
60
0
20000
40000
60000
80000
100000

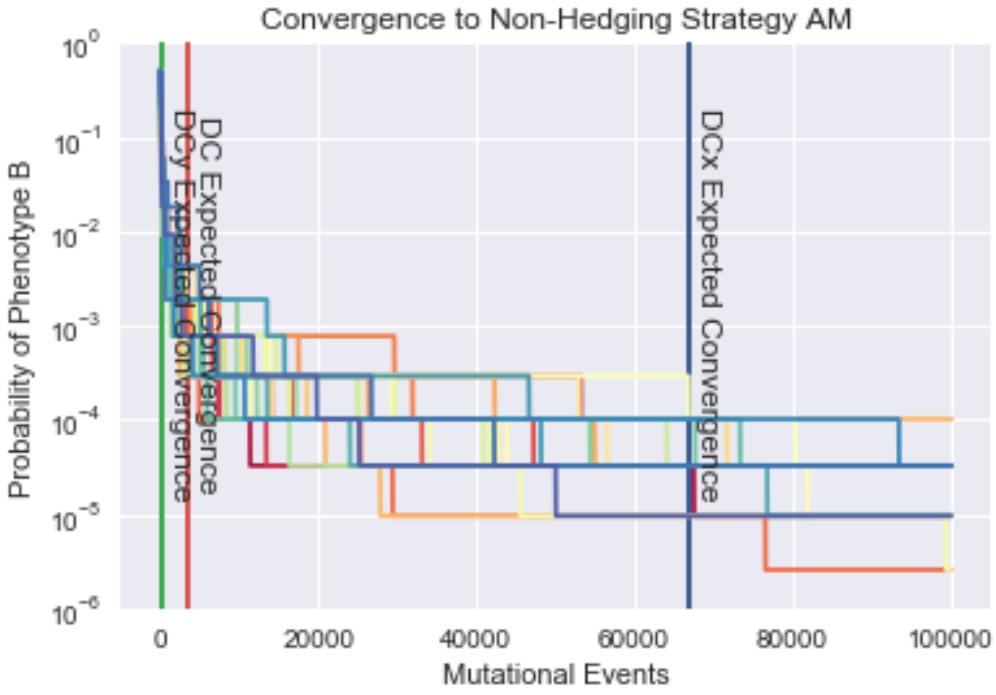
Convergence to Non-Hedging Strategy AM
Probability of Phenotype B
Mutational Events
DCy Expected Convergence
DC Expected Convergence
DCx Expected Convergence
10^{0}
10^{-1}
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-6}
0
20000
40000
60000
80000
100000

1.8 Treatment Simulation - Section 4.6

Chapter 4 concludes with an exploration of treatment holidays. The code used to generate these results is presented in treatment.py.

```
In [17]: #Parameters for the treatment sim
         # Cytotoxic
         fA,fB = 1.0,0.015
         dA,dB = 0.8, 0.005

         # Cytostatic
         # fA,fB = 0.000,0.015
         # dA,dB = 0.005, 0.005

         # Mixed
         # fA,fB = 0.50,0.015
         # dA,dB = 0.40, 0.005

In [18]: import treatment as trt

In [19]: #Insure the appropriate directories exit.
         !mkdir -p 'HistData'
         !mkdir -p 'HistData/cytotoxic'
         path = "./HistData/cytotoxic/"

         #============================================================================#
         # The following code generates the histogram data but is *very* unoptimised and
         # slow  (multiple hours running time for a large number of samples.)
         #============================================================================#
         trials = 2000
         hedges = [AM_hedge, DC_hedge, DC_hedge_x, DC_hedge_y]
         h_names = ["AM", "DC", "DCx", "DCy"]
         g0s = [30,30,7,53]
         for h_index in range(0,4):
             print "Calculating histogram for map: ", h_names[h_index]
             trt.generate_histogram_data([2.0,1.01], hedges[h_index],
                                         g0s[h_index], trials, h_names[h_index], path)

In [24]: #============================================================================#
         # Figures 4.10 - 4.12 are produced as follows
         # Some minor adjustments in shape/size may be needed.
         #============================================================================#
         trt.build_figure(path)

<matplotlib.figure.Figure at 0x118ff38d0>
```

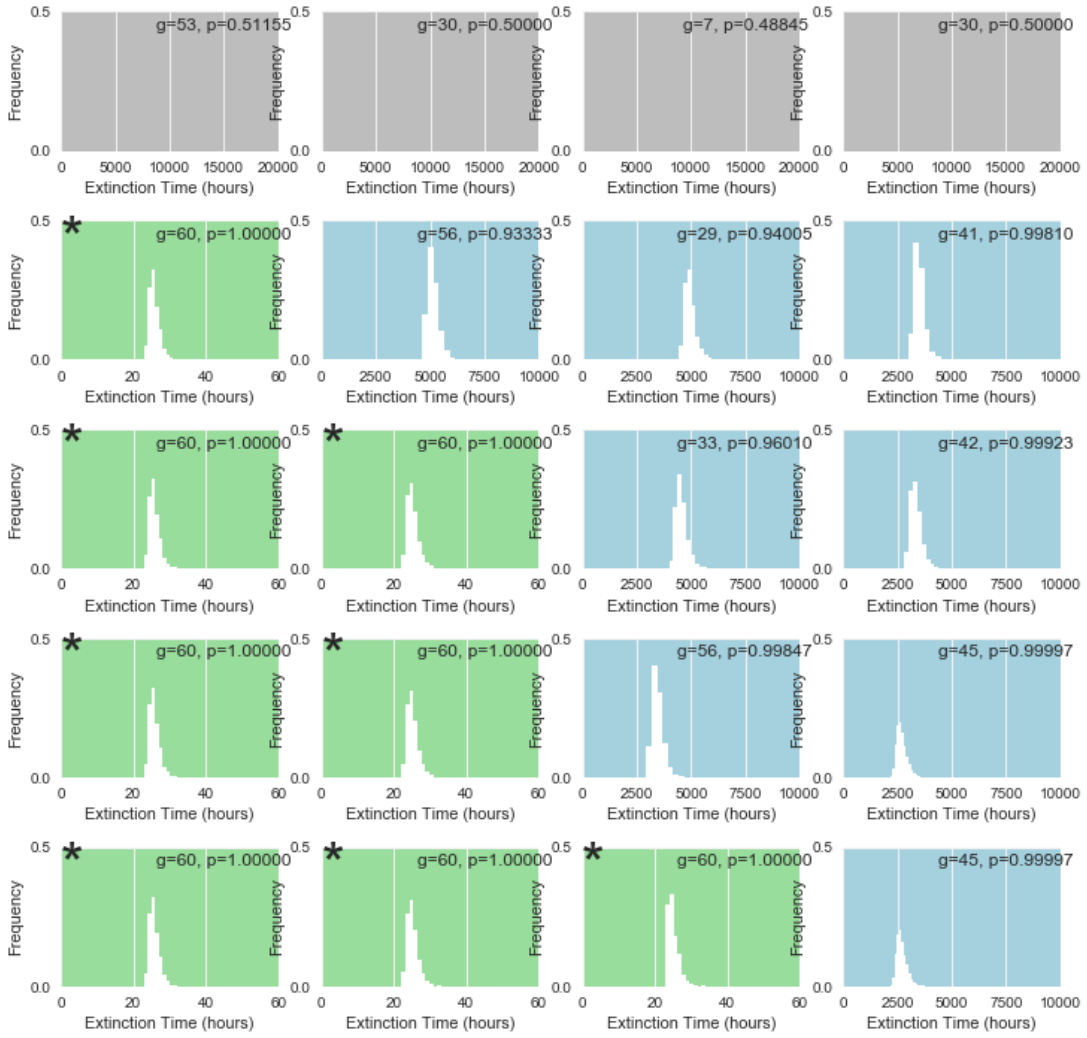
g=53, p=0.51155
g=30, p=0.50000
g=7, p=0.48845
g=30, p=0.50000
g=60, p=1.00000
g=56, p=0.93333
g=29, p=0.94005
g=41, p=0.99810
g=60, p=1.00000
g=60, p=1.00000
g=33, p=0.96010
g=42, p=0.99923
g=60, p=1.00000
g=60, p=1.00000
g=56, p=0.99847
g=45, p=0.99997
g=60, p=1.00000
g=60, p=1.00000
g=60, p=1.00000
g=45, p=0.99997
Frequency
Extinction Time (hours)

Chapter5

January 19, 2018

1 Chapter 5 - Mechanisms of Phenotypic Memory

Presented here is the code used to generate the results of Chapter 5.

1.1 Prerequisites

In [1]:
```
from copy import copy
from copy import deepcopy
import glob
import math
from random import random
import sys

import matplotlib
import matplotlib.pyplot as plt
from matplotlib import rc
from mpl_toolkits.axes_grid1 import Grid
import numpy as np
import seaborn as sns

import seaborn as sbn
sbn.set(rc={'image.cmap': 'cubehelix'})
sbn.set_style("white")

import hedges

print "Python version: ", sys.version
print "numpy version: ", np.__version__
print "matplotlib version: ", matplotlib.__version__

#Plotting style variables
linestyles = ['-', '--', '-.',':']
linecols = [sns.xkcd_rgb["pale red"], sns.xkcd_rgb["denim blue"],
            sns.xkcd_rgb["medium green"],sns.xkcd_rgb["black"]]

%matplotlib inline
```

```
Python version:  2.7.13 |Continuum Analytics, Inc.| (default, Dec 20 2016, 23:05:08)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.57)]
numpy version:  1.11.3
matplotlib version:  2.0.2
```

1.2 Intracellular Memory

The exploration of the intracellular dynamics is presented in memintracellular.py. Specifically, this module implements the equations outlined in Section 5.2.1 and provides code to generate the associated figures. Example use is as follows:

```
In [2]: import memintracellular as mmi
        sigma = 6.77*10**(-7)
        t = 60*60*24 #1day
        #Sampling the numer of molcules remaining at time t, n(t), from n_0=60 with
        # decay rate sigma
        #Note this is stochastic, these two can differ.
        print "Molecules remaining (sample 1): ", mmi.sample(t, sigma, 60)
        print "Molecules remaining (sample 2): ", mmi.sample(t, sigma, 60)

        #Finding the expected number of molecules remaining
        print "Expected remaining: ", mmi.expectation(t, sigma, 60) #Equation 5.11

        #Finding the probability of n=50 molecules remaining
        print "P(50 remaining | n0 = 60) = ", mmi.p(50,t,sigma,60) #Equation 5.10
```

```
Molecules remaining (sample 1):  57
Molecules remaining (sample 2):  58
Expected remaining:  56.5911018784
P(50 remaining | n0 = 60) =  0.00141843568099
```

We also implement functions that calculate the likelihood of an individual having an offspring of phenotype A given a specified n_0, time of reproduction t, decay rate σ, and genetic burst values x_b, y_b (for each molecular switch):

```
In [3]: t_rep = 60*60*24*5 #5days
        xb, yb = 30, 30
        print "P(A|n0=60, sigma, DC , t_rep, xb, yb) = ", mmi.probDC(t, 60, sigma, xb, yb)
        print "P(A|n0=60, sigma, DCx, t_rep, xb, yb) = ", mmi.probDCx(t, 60, sigma, xb, yb)
        print "P(A|n0=60, sigma, DCy, t_rep, xb, yb) = ", mmi.probDCy(t, 60, sigma, xb, yb)
        print "P(A|n0=60, sigma, AM , t_rep, xb, yb) = ", mmi.probAM(t, 60, sigma, xb, yb)
```

```
P(A|n0=60, sigma, DC , t_rep, xb, yb) =  0.742629113301
P(A|n0=60, sigma, DCx, t_rep, xb, yb) =  0.999750734437
P(A|n0=60, sigma, DCy, t_rep, xb, yb) =  0.0572942682524
P(A|n0=60, sigma, AM , t_rep, xb, yb) =  0.99999994845
```

Finally, the code in memintracellular.py can be used to generate the figures in the these as follows

```
In [4]: # mmi.comparing_decay() #Figure 5.3
        # mmi.molsprobtimeplot() #Figure 5.2(A)
        # mmi.exp_over_time() #Figure 5.2(B)
        plt.figure(figsize=(8,3));
        mmi.exp_reproduction_plot(sigma, 30,30,2.44*60*60*24, 24.4*60*60*24) #Figure 5.4
```

```
/Users/dannichol/miniconda2/lib/python2.7/site-packages/matplotlib/__init__.py:913: UserWarning:
  warnings.warn(self.msg_depr % (key, alt_key))
```

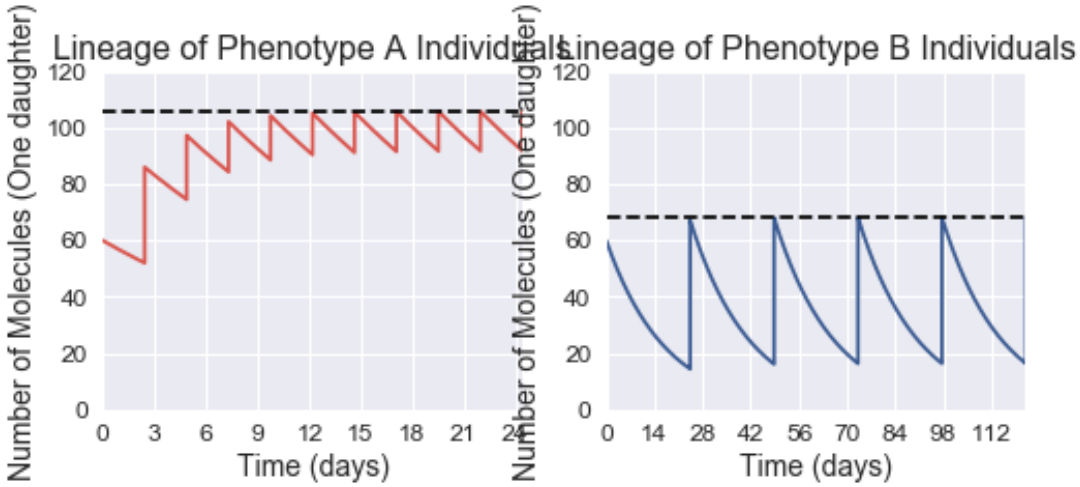

1.3 Modelling Bet-Hedging with Epigenetic Memory

The code in memmodel.py contains the implementation of the model outlined in Section 5.2.2.

```
In [5]: import memmodel as mem
```

```
In [6]: ########################################################
        # Parameters - Section 5.3
        ########################################################
        sigma = 6.77 * 10**(-7) #Equation 5.14
        # Birth Rates (per second)
        b_A = 4.74*10**(-6)
        b_B = 4.74*10**(-7)
        # Death Rates (per second)
        d_A = 1.39*10**(-6)
        d_B = 1.39*10**(-7)
        #Death Rates, drug (per second)
        d_A_drug = 5.56*10**(-6) * 2.
        d_B_drug = d_B
```

```
In [7]: #Building the population dynamics matrix
        M = mem.get_memory_M(sigma, d_A, d_B, b_A, b_B, xb, yb,
                             hedge = hedges.DC_hedge, fix_p=-1, fix_q=-1)

        ########################################################
        # Constructs a matrix model for 1hr of population growth
        # using the above method. This is P^(60*60) as in Eqn 5.33
        #
        # The memory stratgies are denoted by integers according
        # to:
        #
        #   0 - Strong Genetic
        #   1 - Epigenetic Memory
        #   2 - Contrained Memory
        ########################################################
        M1 = mem.get_one_hour_matrices(sigma, xb, hedges.DC_hedge, 1)
```

As discussed in Section 5.6.2, a drug treatment strategy is a length 14 binary vector

```
In [8]: #Example
        strat = [1,0]*7 #iota=7
```

A growth rate for this strategy can be determined as follows

```
In [9]: M_no_drug = mem.get_memory_M(sigma, d_A, d_B, b_A, b_B, xb,
                                    60-xb, hedges.DC_hedge)**(60*60*24) #1 day no drug
        M_drug = mem.get_memory_M(sigma, d_A_drug, d_B_drug, b_A, b_B, xb,
                                  60-xb, hedges.DC_hedge)**(60*60*24) #1 day drug

        print "Growth rate: ", \
                mem.get_strat_growthrate(strat, M_drug, M_no_drug, interval_len=1), \
                'doublings per week'

Growth rate:  0.594258122699 doublings per week
```

Figures 5.7 and 5.10 can be generated as follows. Note, this will take around 10minutes (on a 2012 MacBookPro).

```
In [10]: mem.fixed_e_plot() #5.7
         # mem.generate_cyclic_plot_data(hedges.DC_hedge, 'DC') #5.10 data
         # mem.make_cycling_plot() #5.10
```

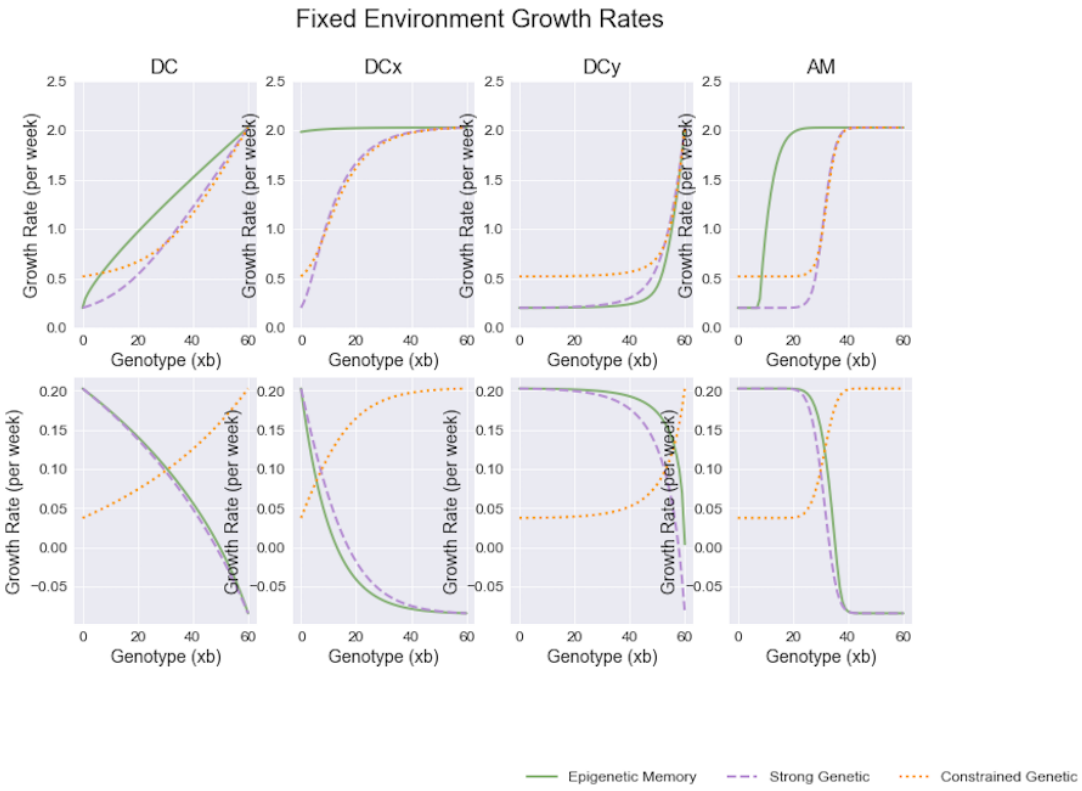

1.4 Exhaustive Exploration of All Metronomic Therapies

Section 5.6 - Implications for Therapy explores the space of all metronomic therapies comprising 14 'units' of length $\iota = 1, 7$ which can be designated as either drug or not drug. The code used to perform thiss exhaustive search is implemented in metronomic.py.

```
In [11]: import metronomic as met
```

To exhaustively evaluate all of the metronomic therapies, for all doses, all hedges, all three memory types and both genotypes (50% and 99% memory-free), call the following methods. Note that this exhaustive search is extremely slow (between 24 and 48 hours). For this reason, the methods in metronomic.py are comment out and should be uncommented to run.

```
In [12]: met.generate_all_therapy_data(interval_len = 1, num_intervals = 14)
         met.generate_all_therapy_data(interval_len = 7, num_intervals = 14)

Warning, takes a LONG time. Uncomment to run...
Warning, takes a LONG time. Uncomment to run...
```

From the exhaustive search, we can produce the figures in Section 5.6 as follows

```
In [14]: met.make_therapy_plot(interval_len=7)
```

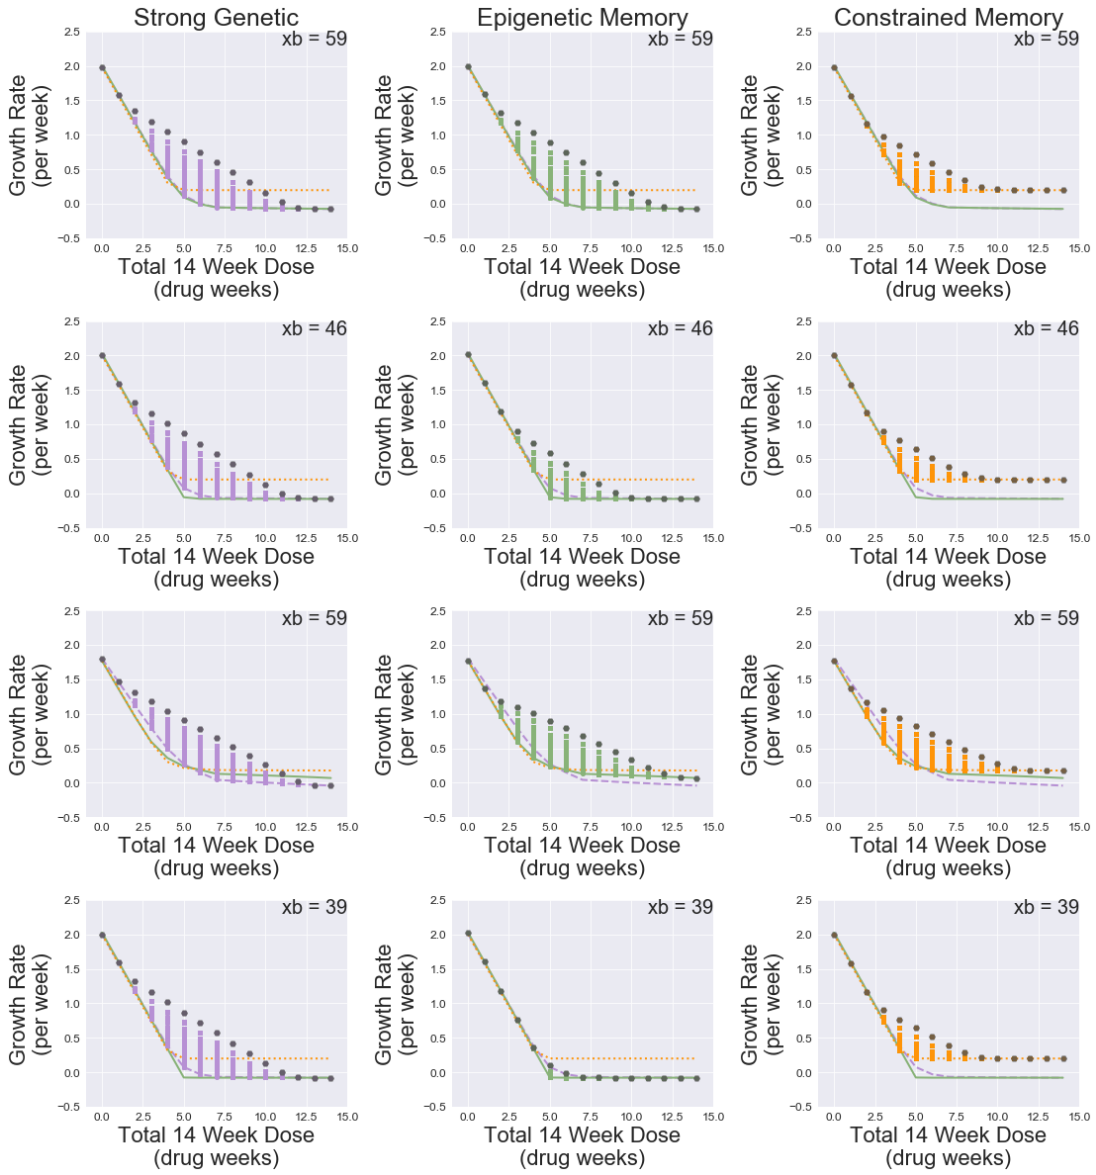

```
In [15]: met.mdd_opt_comparison(0)
```

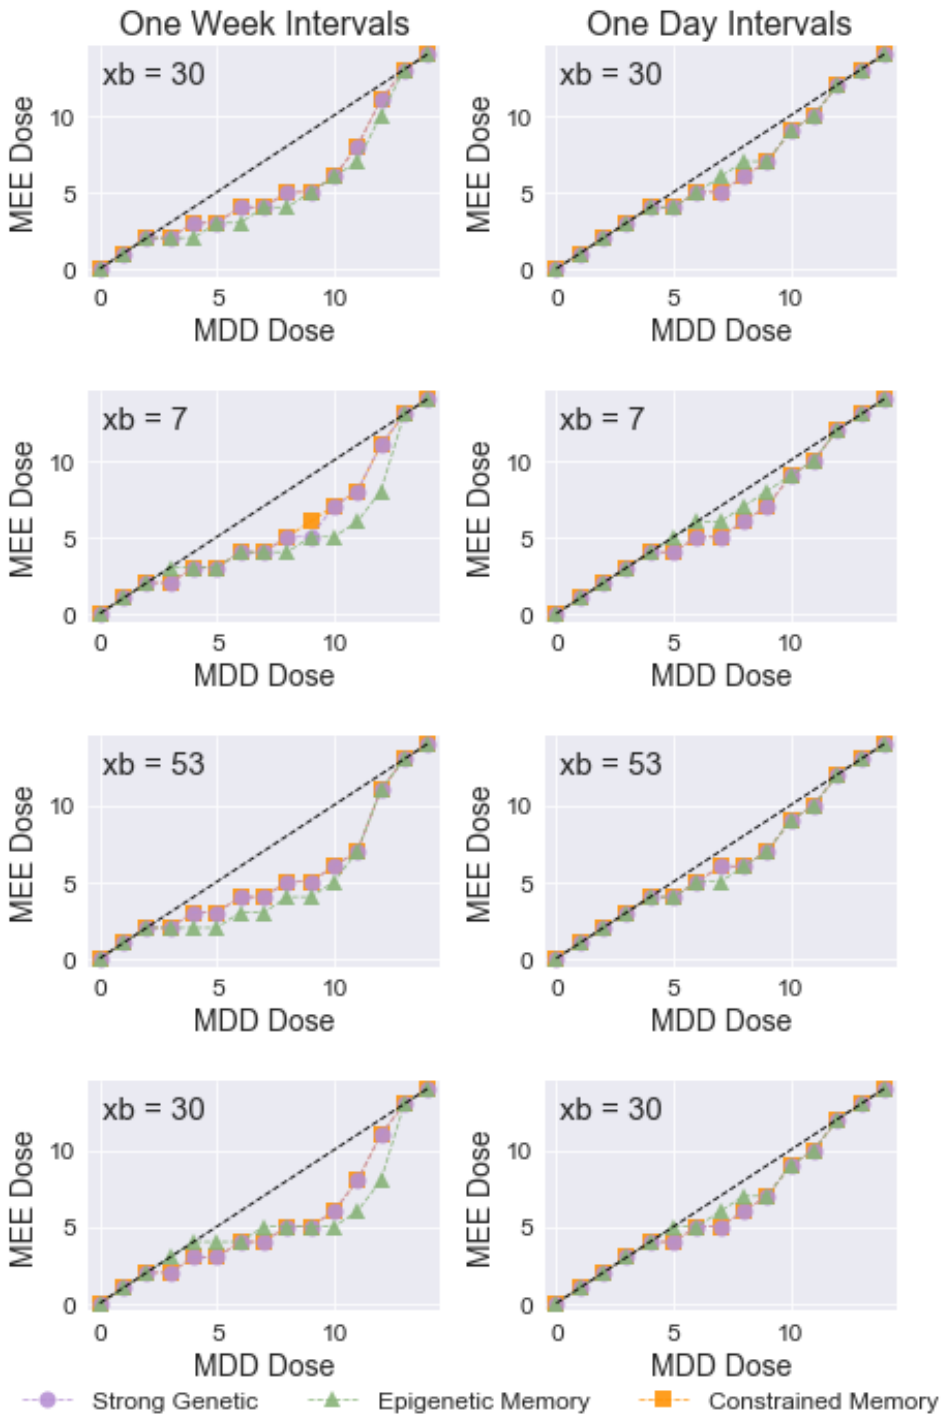

1.5 Varying σ

The results of Section 5.7 which explore how varying σ alters the efficacy of metronomic therapies are produced in varyingsigma.py.

```
In [16]: import varyingsigma as vrs
```

Figures 5.19 and 5.20 are produced as follows:
1 For each σ, the full space of metronomic therapies is explored and the best therapy is recorded
2 The σ-*rho* curves are then built and saved
3 The curves are plotted
The following code performs this task. Note that this will take a number of hours on a personal computer.

By changing dose, gs_num and intervel_len, it is possible to explore the evolution in scenarios not presented in the thesis.

```
In [17]: gs_num = 0 #Set the genotype set to 30,30,7,53 (approx. 50/50)
         dose = 7 #The dose to consider

         # The following four lines are slow.
         # vrs.generate_all_sv(interval_len=1)
         # vrs.generate_all_sv(interval_len=7)

         # vrs.sigma_rho_analysis(0,7, interval_len=7)
         # vrs.sigma_rho_analysis(0,11, interval_len=7)

         #Plot the fig (examples)
         vrs.plot_sig_curves(gs_num, dose=7)
         vrs.plot_sig_curves(gs_num, dose=11)
```

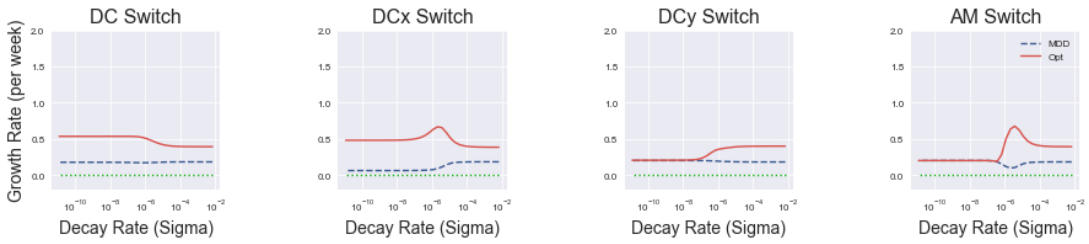

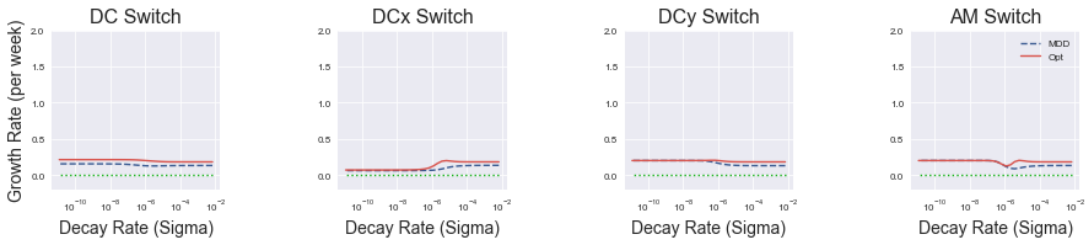

1.6 Adaptive Therapy

adaptivetherapy.py contains the code used to produce the results in Section 5.7.1 that explores the relationship between dose skipping, maximum-dose and optimal metronomic therapy.

Figures 5.21 and 5.22 can be produced as follows. plot_adaptive_curves generates the mdd and dose skipping curvea on the fly. The dose-restricted optimal curve is loaded from the data saved above (metronomic therapy). As such, the adaptive therapy figure cannot be made until after the metronomic therapy figures.

```
In [18]: import adaptivetherapy as apt
```

```
In [19]: apt.plot_adaptive_curves(hedges.DC_hedge, 'DC', 0, 1, 7, 4, hrs_to_change = 7*24)
```

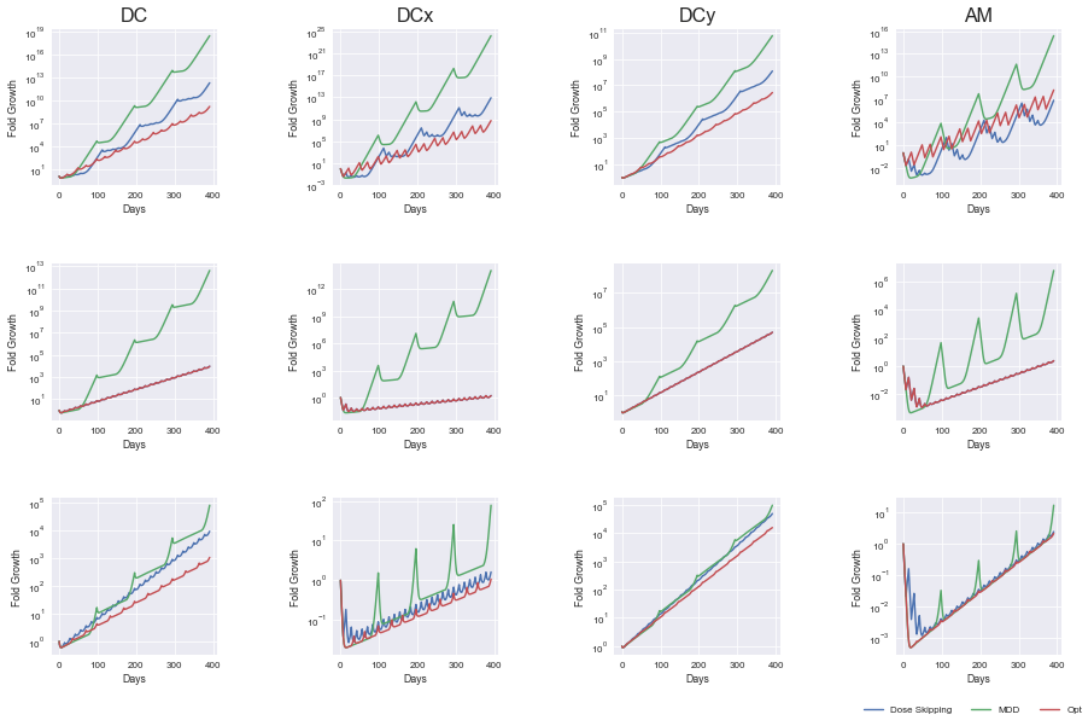

GrowthPhases

January 19, 2018

1 Chapter 6 - The GP-mapping at the population scale

In [1]:
```
import sys

        import matplotlib
        import matplotlib.pyplot as plt
        from matplotlib import rc
        from mpl_toolkits.axes_grid1 import Grid
        import numpy as np
        import seaborn as sns

        import seaborn as sbn
        sbn.set(rc={'image.cmap': 'cubehelix'})
        sbn.set_style("white")

        print "Python version: ", sys.version
        print "numpy version: ", np.__version__
        print "matplotlib version: ", matplotlib.__version__
```

```
Python version:  2.7.13 |Continuum Analytics, Inc.| (default, Dec 20 2016, 23:05:08)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.57)]
numpy version:  1.11.3
matplotlib version:  2.0.2
```

2 Growth Phases (Section 6.2)

2.1 Model and Parameters

The following code constructs the population dynamics matrices defined in Equations 6.2 and 6.3

In [2]:
```
#=========================================================================#
#Eqn 6.2 (Strong Genetic bet-hedging)
#=========================================================================#
def mat(k, p):
    a = (1-da)*(1+k*ba)*p
    b = (1-db)*(1+k*bb)*p
```

```
        c = (1-da)*(1+k*ba)*(1-p)
        d = (1-db)*(1+k*bb)*(1-p)
        m = np.matrix([[a,b],[c,d]])
        return m
    #============================================================================#
    #Eqn 6.3 (Constrained memory bet-hedging)
    #============================================================================#
    def matcm(k, p):
        a = (1-da)*(1+k*ba)*p
        b = (1-db)*(1+k*bb)*(1-p)
        c = (1-da)*(1+k*ba)*(1-p)
        d = (1-db)*(1+k*bb)*p
        m = np.matrix([[a,b],[c,d]])
        return m
```

The parameters, as specified in Section 6.2.1 are as follows

```
In [3]: da = 0.005
        db = 0.005
        ba = 1.0
        bb = 0.015
```

Finally, we defined a number of helper functions that determined population growth properties from the matrices defined above

```
In [4]: def getFreqsA(M):
            w,v = np.linalg.eig(M)
            dom_ix = np.argmax(w)
            fq = v[:,dom_ix]
            fq = fq/sum(fq)
            return fq[0,0]

        def getGrowth(M):
            w,v = np.linalg.eig(M)
            dom_ix = np.argmax(w)
            return w[dom_ix]

        def getStatSize(matfun, p):
            for k in np.arange(0.0,1.0,0.00005):
                w,v = np.linalg.eig(matfun(k,p))
                if (1.0 - np.max(w))**2 < 10**(-8):
                    return k
```

2.2 Results

The key result of Section 6.2 is that the phenotype frequency distribution is dependent on the growth phase for the memory model, but not for the strong genetic model. These results are presented in Figures 6.1 and 6.2 which are produced below.

In [5]:
```
#============================================================================#
#Produces one panel of Fig 6.1 for specified K, p and memory
#============================================================================#
def popdyn(K, p, mem=False):
    params = {
    'axes.labelsize': 14,
    'text.fontsize': 14,
    'axes.titlesize': 16,
    'xtick.labelsize': 12,
    'ytick.labelsize': 12,
    'text.usetex': False,
    'figure.bottom' : 0.16
    }

    Ksmall = 10**3
    matfun = matcm if mem else mat
    pop = np.matrix([[1.0], [1.0]])
    histA, histB = [], []
    T_start = 48
    T_end = 72 * 2
    for t in range(T_start+1):
        k = (1. - ((sum(pop) / Ksmall)))[0,0]
        P = matfun(k,p)
        pop = P * pop
        histA.append(pop[0,0])
        histB.append(pop[1,0])

    for t in range(T_end):
        k = (1. - ((sum(pop) / K)))[0,0]
        P = matfun(k,p)
        pop = P * pop
        histA.append(pop[0,0])
        histB.append(pop[1,0])

    plt.subplot(1,2,1)
    plt.plot(histA, c = sbn.xkcd_rgb["pale red"], label='Phenotype A')
    plt.plot(histB, c = sbn.xkcd_rgb["denim blue"], label='Phenotype B')
    plt.plot((T_start,T_start), (0.0,5*K), c='k', ls='-.', lw=0.5)

    plt.legend()

    plt.xlim(0,120)
    plt.ylim(1,5*K)
    plt.yscale('log')
    plt.xlabel("Time (hrs)")
    plt.ylabel("Pop. Size")
    plt.title("Growth Curves")
    plt.xticks(np.arange(0, len(histA)+1, 24),
```

```
                    [(24*i)-T_start for i in range(len(histA)+1)])

        ax = plt.subplot(1,2,2)
        freq = [histA[i] / (histA[i] + histB[i]) for i in range(len(histA))]
        plt.plot(freq, c=sbn.xkcd_rgb["pale red"])
        ax.fill_between(range(len(histA)), 0, freq,
                        facecolor=sbn.xkcd_rgb["pale red"])
        ax.fill_between(range(len(histA)), freq,
                        [1.0 for i in range(len(histA))],
                        facecolor=sbn.xkcd_rgb["denim blue"])

        #The stationary dists.
        expfq = getFreqsA(matfun(1.0,p))

        K_stat = getStatSize(matfun, p)
        statfq = getFreqsA(matfun(K_stat,p))
        plt.plot([expfq for h in histA], ls='--', c=sbn.xkcd_rgb["black"])
        plt.plot([statfq for h in histA], ls='--', c=sbn.xkcd_rgb["black"])
        plt.plot((T_start,T_start), (0.0,1.0), c='k', ls='-.', lw=0.5)

        plt.xlabel("Time (hrs)")
        plt.ylabel("Phenotype Frequency")
        plt.title("Phenotype Distribution")

        if mem:
            plt.ylim(0.0,1.0)
        else:
            plt.ylim(0.999,1.0)
        plt.xlim(0, len(histA)-1)
        plt.xticks(np.arange(0, len(histA)+1, 24),
                   [(24*i)-T_start for i in range(len(histA)+1)])
        plt.show()
```

In [6]:
```
#================================================================================#
#Example use.
#================================================================================#
plt.figure(figsize=(12,5))
popdyn(10**8, 0.99, True)
```

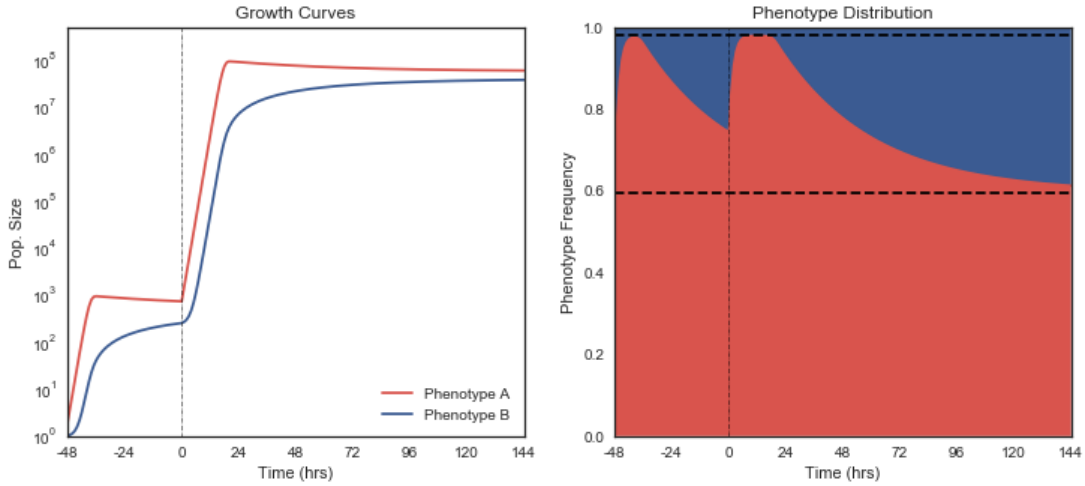

The following code produces Fig 6.2

```
In [7]: #============================================================================#
        # Fig 6.2
        #============================================================================#
        def statfig():
            from mpl_toolkits.axes_grid1.inset_locator import zoomed_inset_axes
            from mpl_toolkits.axes_grid1.inset_locator import mark_inset
            fig, ax = plt.subplots(figsize=[8, 8])

            p = 0.99
            step = 0.001
            freqs, ls = [], []
            domain = np.arange(0.0,1.0+step, step)
            for c in domain:
                M = matcm(c, p)
                a = getFreqsA(M)
                freqs.append(a)
                ls.append(getGrowth(M))

            ax = plt.subplot(2,2,1)
            ax.plot(domain, ls)

            axins = zoomed_inset_axes(ax, 30, loc=4, borderpad=2.5)  # zoom = 6
            axins.plot(domain, ls)
            axins.plot([0.008, 0.008], [0.0,1.0], c='k', ls='--')
            axins.plot([0.008, 0.0], [1.0,1.0], c='k', ls='--')
            # sub region of the original image
            x1, x2, y1, y2 = 0.0, 0.01, 0.995, 1.005
            axins.set_xlim(x1, x2)
            axins.set_ylim(y1, y2)
```

```
# fix the number of ticks on the inset axes
axins.yaxis.get_major_locator().set_params(nbins=3)
axins.xaxis.get_major_locator().set_params(nbins=3)
ax.set_xlabel('c')
ax.set_ylabel(r'growth rate $\lambda(c)$')
ax.set_title(r'p=$0.99$')
ax.set_xlim(0.0,1.0)


ax = plt.subplot(2,2,3)
ax.plot([0.0, 1.0], [max(freqs), max(freqs)], ls='--', c='k')
statfreq = freqs[np.nonzero(domain == 0.008)[0][0]]
ax.plot([0.0, 1.0], [statfreq, statfreq], ls='--', c='k')

ax.annotate('confluence frequency', xy=(0.8, statfreq),
            xytext=(0.3, statfreq+0.1),
            arrowprops=dict(facecolor='black', shrink=0.05))

ax.annotate('exp-growth frequency', xy=(0.95, max(freqs)),
            xytext=(0.3, max(freqs)-0.1),
            arrowprops=dict(facecolor='black', shrink=0.05))

plt.plot(domain, freqs)

ax.set_xlim(0.0,1.0)
ax.set_ylim(0.5,1.01)
ax.set_xlabel('c')
ax.set_title(r'p=$0.99$')
ax.set_ylabel(r'Phenotype $A$ frequency')

p = 0.99995
step = 0.001
freqs, ls = [], []
domain = np.arange(0.0,1.0+step, step)
for c in domain:
    M = matcm(c, p)
    a = getFreqsA(M)
    freqs.append(a)
    ls.append(getGrowth(M))

ax = plt.subplot(2,2,2)
ax.plot(domain, ls)
axins = zoomed_inset_axes(ax, 30, loc=4, borderpad=2.5)  # zoom = 6
axins.plot(domain, ls)
axins.plot([0.005, 0.005], [0.0,1.0], c='k', ls='--')
axins.plot([0.005, 0.0], [1.0,1.0], c='k', ls='--')
x1, x2, y1, y2 = 0.0, 0.01, 0.995, 1.005
axins.set_xlim(x1, x2)
```

```
            axins.set_ylim(y1, y2)
            axins.yaxis.get_major_locator().set_params(nbins=3)
            axins.xaxis.get_major_locator().set_params(nbins=3)

            ax.set_xlabel('c')
            ax.set_ylabel(r'growth rate $\lambda(c)$')
            ax.set_title(r'p=$0.99995$')
            ax.set_xlim(0.0,1.0)

            ax = plt.subplot(2,2,4)
            plt.plot(domain, freqs)

            ax.plot([0.0, 1.0], [max(freqs), max(freqs)], ls='--', c='k')
            statfreq = freqs[np.nonzero(domain == 0.005)[0][0]]
            ax.plot([0.0, 1.0], [statfreq, statfreq], ls='--', c='k')

            ax.annotate('confluence frequency', xy=(0.8, statfreq),
                        xytext=(0.3, statfreq-0.1),
                        arrowprops=dict(facecolor='black', shrink=0.05))

            ax.annotate('exp-growth frequency', xy=(0.2, 1.005),
                        xytext=(0.2, max(freqs)-0.15),
                        arrowprops=dict(facecolor='black', shrink=0.05))

            ax.set_xlim(0.0,1.0)
            ax.set_ylim(0.8,1.01)
            ax.set_xlabel('c')
            ax.set_title(r'p=$0.99995$')
            ax.set_ylabel(r'Phenotype $A$ frequency')

            plt.suptitle('Constrained--memory bet--hedging dynamics ')
            plt.subplots_adjust(hspace=0.3)
            plt.draw()
            plt.show()
```

In [8]:
```
#=============================================================================#
#Example use.
#=============================================================================#
plt.figure(figsize=(15,15));
statfig();
```

```
<matplotlib.figure.Figure at 0x10feed290>
```

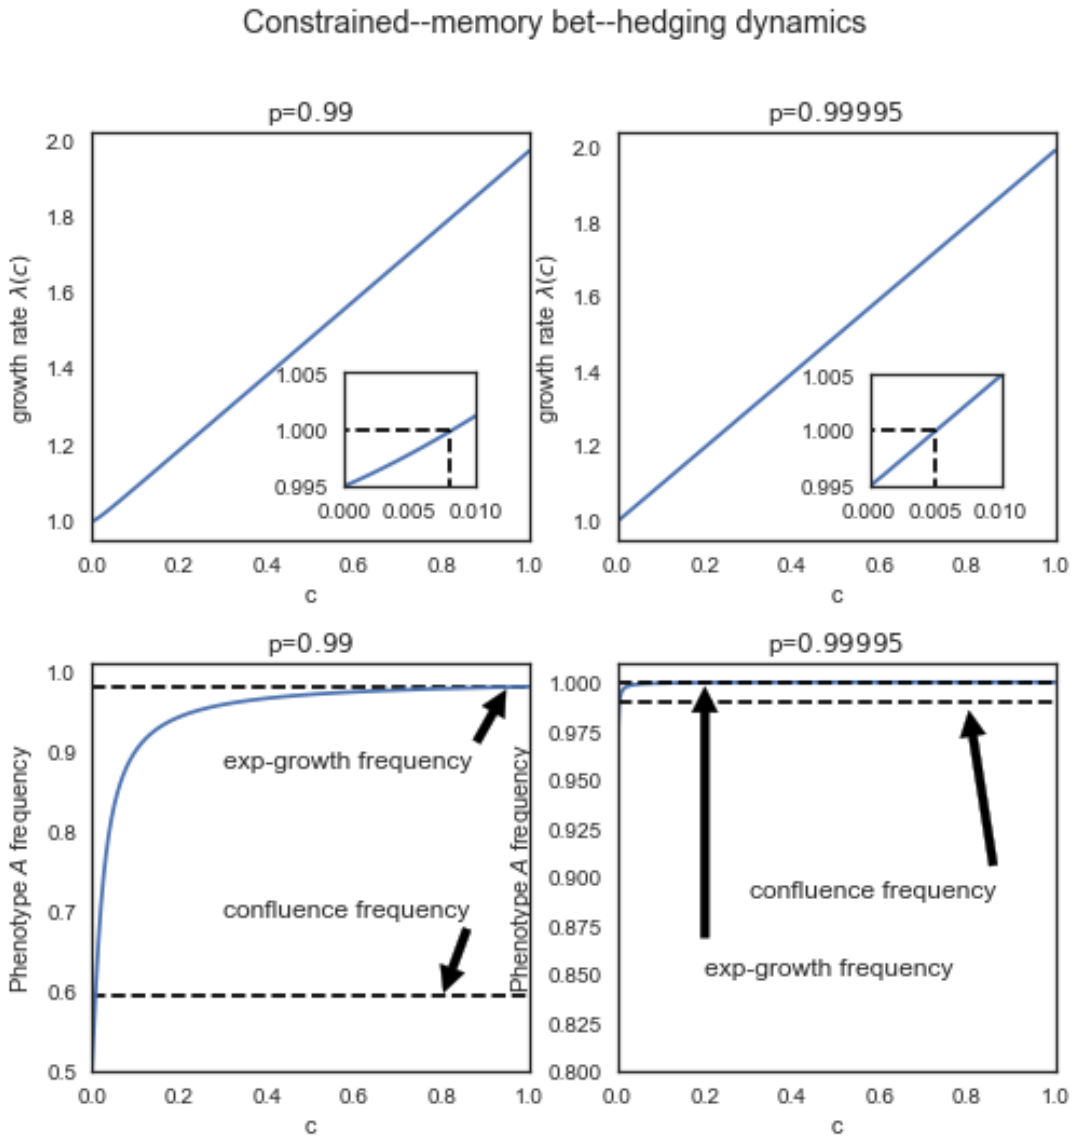

3 Mutational Spectra (Section 6.3)

The results of Section 6.3 are derived from repeated simulation of growth, mutation and genomic sequencing in exponentially growing populations exhibiting either genetically-driven resistance or bet-hedging-driven resistance. The simulations are defined in the Python files

mutsimselection.py - Simulation of growth with genetically-driven resistance

mutimBH - Simulation of growth with bet-hedging-driven resistance

These files must be run in the same folder as primes.txt which contains a list of primes used in the optimisation (as outlined in Section 6.3.1 - Model Optimisation).

To generate a sufficient dataset to create the results figures, these scripts should be run a large number of times (200 runs were performed for the results presented in the thesis)

```
In [9]: import mutsimBH as msb
        import mutsimselection as mss
```

3.1 Example model use

The simulation of population growth can be performed as follows

```
In [10]: #=============================================================================#
         #Example simulation Parameters
         #=============================================================================#
         rate = np.log(150)*1./(7*24.)
         br = 1./24
         dr = br - rate
         drugEffect = 10
         mu = 1./10.
         drugSize = 2.5 * 10**4
         finSize = 1.02 * drugSize
         Tdrug = 24. * 7. * 2
```

```
In [11]: #=============================================================================#
         # Example simulation
         #=============================================================================#
         #Run the simulation
         # ts, hist, sizes = msb.sim(finSize, drugSize, Tdrug, br, 0.1*br, dr, 0.01*dr, drugEffe
         #Extract the counts
         # cs1, cs2 = msb.extractcounts(hist) #Warning V. slow.
```

3.2 Generating the full data set

The follow code will run the models to produce the full set of mutational burden simulations. Note, this simulation is extremely slow. The results presented in the thesis were derived by running the models on a scientific computing cluster.

```
In [12]: #generate the bet-hedging data
         !mkdir -p BHsimData99sg
         # run_num = 10
         # ts, hist, sizes = mbs.sim(finSize, drugSize, Tdrug, br, 0.1*br, dr, 0.01*dr, drugEffe
         # np.save('./BHSimData99sg/curve_'+str(run_num), [ts, sizes])
         # cs1, cs2 = extractcounts(hist)
         # np.save('./BHSimData99sg/predrug_'+str(run_num)+'.npy', cs1)
         # np.save('./BHSimData99sg/progression_'+str(run_num)+'.npy', cs2)
```

```
In [13]: #generating the clonal growth data
         !mkdir -p GeneticSimData
         run_num = 10
         # ts, hist, sizes = mss.sim(finSize, drugSize, br, br, dr, dr, drugEffect, mu)
         # np.save('./GeneticSimData/curve_'+str(run_num), [ts, sizes])
         # cs1, cs2 = extractcounts(hist)
         # np.save('./GeneticSimData/predrug_'+str(run_num)+'.npy', cs1)
         # np.save('./GeneticSimData/progression_'+str(run_num)+'.npy', cs2)
```

3.3 Simulating Sequencing and Visualisation

The file specplotter.py contains the necessary code transforms the list of CCFs produced by the above simulations into in silico generate variant allele frequencies. This file also contains the code to produce Figures 6.5 and 6.6

```
In [14]: import specplotter as spp
```

Simulating the sequencing and extracting the clonal mutations is performed as follows. This method saves lists of pre- and post-therapy clonal mutation number

```
In [15]: # spp.compareGPs(depth=120)
```

The figures are built as follows:

```
In [16]: #==========================================================================#
         #Plotting example growth curves (first row)
         #==========================================================================#
         spp.example_curves()
```

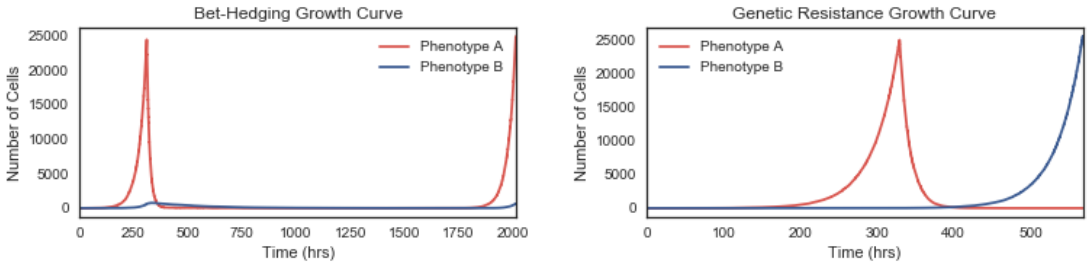


```
In [17]: #==========================================================================#
         #Plotting the matches pre and post build by compareGPs
         #==========================================================================#
         spp.plot_pre_post()
```

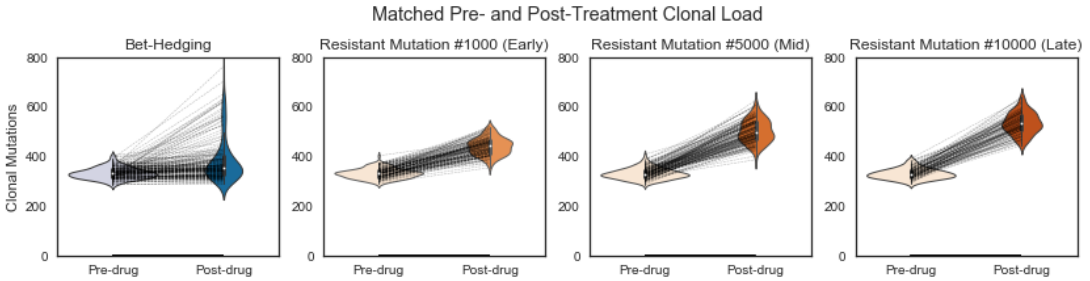

In [18]:
```
#==============================================================================#
#Finally, plotting the delta clonal mutations
#==============================================================================#
spp.plot_delta()
```

```
51.7647058824 109.5 171.993150685 204.462264151
14.0 110.0 167.0 199.5
```