

Adaptive Therapy for Cancer: Integrated modelling of tumoral competition for controlling drug resistance



Maximilian Andreas Roland Strobl
Jesus College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2021

Statement of Originality

This thesis is submitted to the University of Oxford in fulfilment of the requirements of the degree of Doctor of Philosophy. This thesis contains no material which has been previously submitted for a degree or diploma at this University or any other. Except for parts of Chapters 2 & 4, all work was carried out solely by myself. We acknowledge Professor Yannick Viossat for making us aware of the work by Hansen et al [1], and for some technical assistance in Chapter 2. Furthermore, the experimental data discussed in Chapter 4 were collected by Professor Mehdi Damaghi and Miss Samantha Byrne, and are presented here with their permission. For the analysis of these data I used Python code written by Mr Matthew Wicker, and again I present the results with his permission. All contributions which are not my own are clearly marked within the text.

Abstract

Treatment in advanced cancers often provides only temporary improvements due to the emergence of drug-resistant tumour subpopulations. Standard-of-care regimens seek to maximise cell kill in order to achieve a cure, but thereby inadvertently also free resistant cells from the intra-tumoral competition which usually constrains their growth. An emerging alternative strategy, called adaptive therapy (AT), aims to tackle this problem by prioritising tumour control over cure. The idea is to leverage intra-tumoral competition to slow, or even avoid, the expansion of resistance.

In this thesis we present an integration of mathematical modelling, clinical data analysis, and biological experiments in order to investigate: i) the factors which determine whether a patient will benefit from AT, and ii) how to best adjust treatment to control resistance. We begin by studying a simple two-population model, which we formulate first as an ordinary differential equation (ODE), and subsequently as an agent-based, model. We dissect the impact of various model parameters on the benefit of AT, and identify cell turnover as an important, but previously overlooked, factor. Furthermore, we find that the tumour's spatial organisation reflects, as well as modulates, intra-tumoral competition, and we discuss implications for intermittent androgen deprivation treatment in prostate cancer.

In the second part of this thesis, we investigate whether adaptive scheduling of poly-adenosine di-phosphate ribose polymerase inhibitors (PARPis) can improve the treatment of ovarian cancer. We develop a framework which uses an ensemble of six ODE models to dynamically guide AT decision-making. In collaboration with an experimentalist, we test our framework in an *in vitro* spheroid model, and find that it compares favourably to other adaptive dosing algorithms. However, this work also shows that high-dose, continuous treatment achieves the best overall outcome. To conclude we therefore develop an ODE model of PARPi treatment, and discuss implications for PARPi scheduling. Overall, this thesis provides novel insights into the eco-evolutionary underpinnings of AT, and highlights some of the challenges involved in translating AT to a new cancer type.

Acknowledgements

A PhD is a journey, and as any journey it comes with its ups and downs. I have been very fortunate that along my path I have had the tremendous support of my mentors, friends and family. Firstly, I would like to thank my supervisors, Professor Philip Maini and Professor Alexander Anderson. You guided and stimulated me to discover my abilities, and stood by unwaveringly whenever I faced difficulties. Thanks to you I have grown not only as a scientist, but also as a person. Furthermore, I am greatly indebted to Doctor Jeffrey West and Professor Mehdi Damaghi for their continued encouragement and support. Thank you Jeffrey, for teaching me how to be more patient and persistent in my work. Thank you Mehdi, for introducing me to the wet lab, and for your inspiring passion for cancer research.

In addition, I'm grateful to Professor Mark Coles and Professor Joe Pitt-Francis for their thoughtful feedback after my transfer of status examination. I would also like to thank Professor Jose Conejo-Garcia and Doctor Carmen Anadon for letting me experience immunology research, and Dr Robert Wenham for making the PARPi work possible. Furthermore, my thanks goes to Ms Samantha Byrne, Mr Matthew Wicker, and Mr Punit Borad for their assistance in the *in vitro* experiments. Next, I would like to thank Tom Bury for his comment on the Mathonco blog which started me on the path which made me realise the role of turnover in AT (<http://disq.us/p/23j1825>). Finally, I'd like to say a great thank you to Professor Yannick Viossat for helping me to develop a deeper understanding of the mathematics behind adaptive therapy.

This thesis was written across three countries and two continents, each of which has become a home to me in its own way. I would like to thank the many friends, new and old, who have accompanied me along this way. In particular, thank you to the members of the Wolfson Centre for Mathematical Biology, the amazing people in S04.16, and Doctor Andrew Krause, for providing my home base in Oxford. Furthermore I'm grateful to the members of the Department for Integrated Mathematical Oncology at the Moffitt Cancer Center for being both inspiring colleagues as well as great friends. Ryan, Becky, Rafael, Jill, Etienne, Gwen, Chandler, Greg, Mark, Derek, Danae, the Therapy and DIYC crews, and everyone else, I will never forget our many adventures in and around Tampa. I would also like to thank the H. Lee Moffitt Cancer Center, the SABS DTC, and Jesus College for funding my travels between Tampa and Oxford.

Finally, I would like to thank my family, who lovingly stood by my side every step of the way (quite literally for the past 9 months). You gave me the strength to persist and follow my passion, when doubts and stress were obscuring the path ahead.

To conclude, an excerpt from the thoughts of my furry, little helper, which the margins of this thesis were otherwise too narrow to contain: trrrRrghhhhh.

Glossary

- **ABM:** Agent-based model
- **AIC:** Akaike information criterion
- **AT:** Adaptive therapy
- **BIC:** Bayesian information criterion
- **CI:** Confidence interval
- **CT:** Continuous therapy
- **CTC:** Circulating tumour cell
- **ctDNA:** Circulating tumour DNA
- **DMSO:** Dimethyl sulfoxide
- **DSB:** Double strand break
- **GFP:** Green fluorescent protein
- **HR:** Homologous repair
- **HRD:** Homologous repair deficient
- **MPC:** Model predictive control
- **ODE:** Ordinary differential equation
- **PAR:** Poly-adenosine diphosphate ribose
- **PARP:** PAR polymerase
- **PARPi:** PARP inhibitor
- **PD:** Pharmacodynamics
- **PDE:** Partial differential equation
- **PFS:** Progression free survival
- **PK:** Pharmacokinetics
- **PSA:** Prostate specific antigen
- **RECIST:** Response evaluation criteria in solid tumors
- **RFP:** Red fluorescent protein
- **RMSE:** Root mean squared error
- **RPMI:** Roswell Park Memorial Institute Medium
- **SLD:** Sum of longest diameter
- **SS:** Steady state
- **ssDNA:** single stranded DNA
- **TDM:** Therapeutic drug monitoring
- **TGI:** Tumour growth inhibition
- **TTP:** Time to progression

Contents

1	Introduction	1
1.1	Overview of modern cancer therapy	2
1.1.1	What is cancer?	2
1.1.2	The oncologist’s arsenal	3
1.1.3	First-line treatment	3
1.1.4	Evaluation of treatment response	4
1.1.5	Outcomes: the challenge of recurrence	5
1.1.6	Toxicity	5
1.2	Therapy scheduling: evolving paradigms	6
1.2.1	The mathematics of treatment scheduling	6
1.2.2	The log-kill hypothesis: Total therapy cures leukaemia	8
1.2.3	The Norton-Simon Hypothesis: Kinetic resistance	9
1.2.4	The Goldie-Coldman model: Molecular resistance	10
1.2.5	Metronomic therapy: Targeting the environment	12
1.2.6	Pharmacokinetics: Challenges of drug delivery	13
1.2.7	Safety	14
1.3	Adaptive therapy	15
1.3.1	The rationale behind adaptive therapy	16
1.3.2	The current state of the field	18
1.4	Outline of the thesis	20
2	On the Role of Resistance Costs in Adaptive Cancer Therapy	22
2.1	Introduction	22
2.1.1	What is the cost of resistance?	23
2.1.2	A brief review of resistance costs in cancer	24
2.1.3	Are resistance costs necessary for AT?	25
2.1.4	Outline	25
2.1.5	Note	26

2.2	The Mathematical Model	26
2.2.1	Non-dimensionalisation	29
2.2.2	Model parameters	30
2.2.3	Numerical methods	31
2.3	Impact of a cost in the proliferation rate, r_R	31
2.3.1	Dynamics in the absence of turnover	32
2.3.2	Dynamics in the presence of turnover	40
2.4	Extension to other types of cost	50
2.4.1	Impact on the effective carrying capacity	50
2.4.2	Comparison of the AT treatment dynamics	53
2.5	Application to intermittent androgen deprivation therapy in prostate cancer	54
2.5.1	The data	55
2.5.2	Data processing and data fitting	56
2.5.3	The model is able to recapitulate treatment dynamics	57
2.5.4	Cost and turnover are required for describing patient dynamics	59
2.5.5	Neither cost nor turnover in isolation correlate with outcome	60
2.5.6	Context matters: Cost and turnover in combination predict outcome	61
2.6	Improving AT by amplifying competition	63
2.7	Discussion	64
3	The impact of space and stochasticity on intra-tumoral competition and adaptive therapy	68
3.1	Introduction	68
3.1.1	Spatial modelling of AT	68
3.2	The computational model	71
3.2.1	Model assumptions and parameters	71
3.2.2	Stochastic simulation algorithm	74
3.3	Comparison of AT and CT in the ABM	75
3.3.1	AT can be superior to CT despite random initial seeding of cells	75
3.3.2	Quantification of competitive inhibition of resistant cells during treatment	77
3.3.3	High tumour cell density and low initial resistance fraction maximise competition from sensitive cells	78
3.3.4	Turnover also increases the benefit of AT in the ABM	81

3.3.5	Turnover amplifies the impact of competitive inhibition on resistant cell growth	82
3.3.6	AT can be inferior to CT	84
3.4	Comparison between the ODE model and the ABM	90
3.4.1	The impact of the initial tumour composition	90
3.4.2	The impact of resistance costs and turnover	93
3.5	Impact of the initial resistant cell distribution	97
3.5.1	Mono-clonal resistance is more effectively controllable because it allows AT to leverage both inter- and intra-specific competition	99
3.5.2	The spatial distribution of resistance modulates the impact of cost and turnover	100
3.6	Further investigation of patient responses to intermittent androgen-deprivation treatment	101
3.6.1	Model fitting	102
3.6.2	The ABM recapitulates fast and slow cycling patients	103
3.6.3	Inter-patient variability is driven by differences in the resistance cost and turnover	105
3.6.4	Cycling speeds may reflect intra-tumoral organisation	105
3.6.5	A combination of relatively smaller cost or turnover correlate with progression	107
3.7	Discussion	107
4	Using modelling to guide adaptive therapy decision-making: two case studies in ovarian cancer	113
4.1	Introduction	113
4.1.1	A review of AT dosing algorithms	114
4.1.2	Use of mathematical models in treatment decision-making	116
4.1.3	Summary	119
4.1.4	Outline	119
4.2	Open-loop optimisation of PARPi treatment in ovarian cancer	120
4.2.1	Background	120
4.2.2	A pilot experiment comparing continuous and adaptive Olaparib treatment <i>in vitro</i>	121
4.2.3	A virtual cohort approach to improve scheduling	124
4.2.4	Results	131
4.2.5	Discussion	137

4.3	Individualised, closed-loop optimisation of PARP inhibitor treatment in ovarian cancer	140
4.3.1	Introduction	140
4.3.2	The framework	142
4.3.3	Results	156
4.3.4	Discussion	164
4.4	Concluding remarks	166
5	Integrated experimental and mathematical evaluation of the oppor- tunities for adaptive PARPi treatment	169
5.1	Introduction	169
5.1.1	Review of PARPi biology and clinical use	170
5.2	Characterisation of <i>in vitro</i> PARPi response	173
5.2.1	Experimental methods	174
5.2.2	Response to continuous treatment	175
5.2.3	Response to treatment withdrawal	177
5.2.4	Characterisation of drug resistant cells	178
5.2.5	Discussion	178
5.3	A mathematical model of PARPi treatment	180
5.3.1	The previously used ODE model is insufficient to describe the experimental data	181
5.3.2	Model development	184
5.3.3	Application to adaptive PARPi treatment	191
5.4	Discussion	195
6	Discussion and Future Work	198
7	Conclusion	204
A	Appendix for Chapter 2	205
B	Appendix for Chapter 3	206
C	Appendix for Chapter 4	214
	References	220

Chapter 1

Introduction

“Thus we may make the strange assertion that the more faulty the spraying this year, the easier it will be to control the scale the next year.”

Axel Melander, 1914

“Can insects become resistant to sprays?” - this is the question entomologist Axel Melander raised in an article of the same title in 1914 [2]. The answer seems obvious today, when in the USA alone more than \$2 billion are spent each year on managing pesticide resistance [3]; but this was not true at the time. Melander reported that farmers in Clarkston valley, USA, were increasingly struggling to control San Jose scale, an insect pest which was damaging local orchards. Whilst normally scales could be effectively killed with sulphur-lime pesticide, more and more insects were found to survive the spraying each year. Construction of a local factory to ensure correct manufacturing, and power spraying of the trees to drench them in pesticide, brought temporary remission, but did not fix the problem. Thus, Melander reasoned that something must have changed about the insects. He applied the same batch of insecticide with the same protocol to scales from Clarkston valley and from orchards in other regions. Remarkably, he found that while he was able to consistently eradicate the scales from the surrounding regions, between 4% and 17% of the scales from Clarkston survived the experiment [2].

This may seem like a strange start to a thesis on cancer research. Why should the problems of fruit growers from 1914 matter to cancer patients today? The reason is that both are battling similar opponents: ecological systems in which treatment triggers Darwinian evolution. Motivated by this similarity, we will explore whether concepts used in the management of resistant insects may also help to improve the treatment of cancer. We will begin by reviewing the current state of cancer treatment, focussing in particular on the scheduling of systemic therapies. We will discuss the

important role which mathematical modelling has played in treatment scheduling, but also how our approaches still often fall short in the metastatic setting. Subsequently, we will show how an eco-evolutionary perspective of cancer can potentially explain these failures, and review a novel paradigm, called “Adaptive Therapy” (AT). AT was inspired by strategies from agricultural pest management, the foundations of which were laid out by Melander in 1914 (see epigraph; [2]). This will motivate the four questions which will form the focus of this thesis.

1.1 Overview of modern cancer therapy

Until the 1950s, the options for treating cancers were scarce and rarely successful. Whilst certain tumours could be removed by surgery or radiation therapy, in advanced cases these approaches often had little benefit [4]. Hippocrates, one of the fathers of medicine, even suggested that cancer was “*best left untreated, since patients live longer that way*”¹. However, thanks to the development of systemic therapies over the past 60 years, an unprecedented range of options exists today. In the following, we will review how a cancer diagnosis is currently approached in the clinic, in order to introduce key concepts and terminology which we will use within this thesis.

1.1.1 What is cancer?

Cancers are a group of diseases characterised by abnormal cell growth. A local cell mass is called a *tumour* (Greek: burden), which can be either *benign*, if growth is restricted to a small region, or *malignant*, if cells have started to invade the surrounding tissue. If cells spread to multiple regions in the body, then this is termed *metastatic* disease. In 90% of patients, the cause of death is not the primary tumour, but the disruption of normal body function caused by pervasive metastatic growth [6, 7].

Cancers are distinguished by their tissue of origin [8]. *Liquid* tumours are cancers originating from the bone marrow (e.g. leukaemias), whereas *solid* tumours are cancers which grow as masses in the patient’s tissues (e.g. breast or lung cancer). Furthermore, cancers are *graded* and *staged* to describe the extent to which the tumour cells differ from those in their tissue of origin, and the degree to which they have begun to spread around the body. Localised tumours are early-stage, and typically low-grade, whereas metastatic cancers are considered late-stage disease, and are often high-grade. In this thesis, we will focus on solid tumours, in particular advanced prostate and ovarian cancer.

¹Translation of Aphorism 38 as provided in [5].

1.1.2 The oncologist's arsenal

Modern cancer therapy is based on three main approaches: i) surgery, ii) radio-therapy, and iii) systemic therapies. Surgery and radio-therapy allow effective localised treatment of individual tumour lesions, provided these can be identified and accessed with existing technology [8]. Here, we will focus on *systemic therapies*. These are drugs which are administered by infusion or are taken orally, and are aimed at killing cells, which can not be reached by surgery or radio-therapy. There are four main classes of systemic therapies: i) chemotherapies, ii) targeted therapies, iii) anti-angiogenic therapies, and iv) immunotherapies. Chemotherapies, such as 5-Fluoro-Uracil, carboplatin, or methotrexate, interfere with the cell division machinery by causing DNA damage and cellular stress, which results in the arrest of division (*cytostatic* therapies) or cell death (*cytotoxic* therapies). Targeted therapies take a more nuanced approach and leverage specific weaknesses of a cancer. For example, many prostate cancers exploit hormonal androgen signalling (e.g. testosterone) to drive their growth. Anti-androgen therapies, such as Arbiraterone or Leuprorelin, specifically target proteins in this pathway, and can thereby kill cancer cells in a more targeted fashion [7, 8].

In contrast, anti-angiogenic treatment and immunotherapies are focussed not on the tumour cells, but on their environment [7]. For example, anti-angiogenic agents, such as bevacizumab, inhibit the formation of new blood vessels in the tumour, thereby cutting-off the tumour's nutrient supply. While these approaches have yielded promising results in recent years, chemo- and targeted therapies remain the back stay of treatment in most tumours, and will be the focus of this thesis.

1.1.3 First-line treatment

Following diagnosis, patients are typically treated with a combination of all three treatment modalities. While the details depend on the type and stage of the disease, the standard-of-care protocol for many solid tumours is organised in a similar fashion (Figure 1.1). Treatment is centred around surgery, in which the primary tumour and visible metastases are removed. In preparation for surgery, patients may receive *neo-adjuvant* therapy to reduce the tumour size, and facilitate its removal. Following surgery, patients receive *adjuvant* therapy to eradicate any remaining cells. Adjuvant and neo-advujant treatment almost always includes systemic therapies, and may also make use of radio-therapy [8, 7].

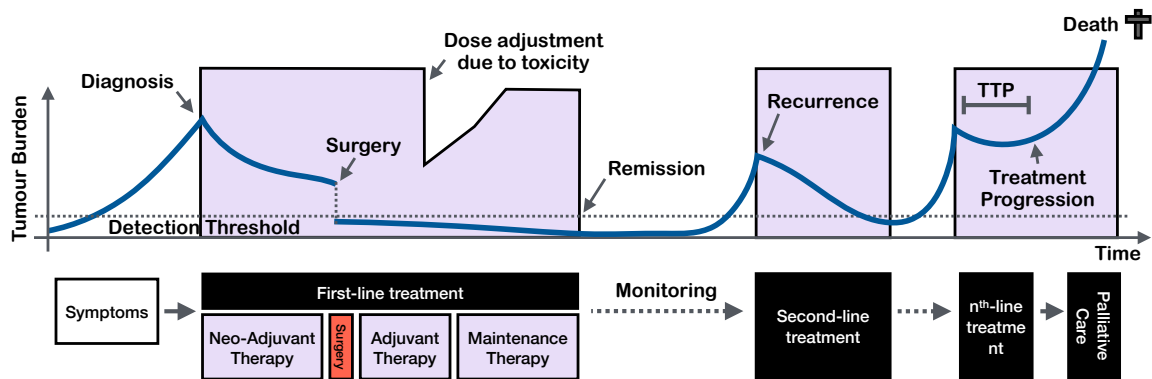


Figure 1.1: Exemplary illustration of the clinical trajectory of a cancer patient, for whom therapy is unable to cure the disease. While therapy may initially achieve remissions, these are temporary, and shorten over the course of treatment. Note: the details of each treatment, and its length will vary between cancer type, and between patients. TTP: Time to Progression.

1.1.4 Evaluation of treatment response

To plan therapy and evaluate its effect, tumour burden is monitored during treatment. Where tumours grow on or near the body’s surface, they may be palpated, or evaluated using endoscopic or laparoscopic devices (e.g. colorectal tumours). More quantitative information is provided by radiologic imaging using, for example, Computed Tomography or Magnetic Resonance Imaging. High resolution molecular detail can be obtained through tumour biopsies. However, most of these methods are applicable only once lesions reach a certain size, potentially require invasive procedures, and are relatively expensive. As a result, clinical time-course data are often sparse [8].

An alternative technology, which is becoming increasingly important, is blood-based disease monitoring (so-called *liquid biopsies*; [8, 9]). Firstly, certain cancers shed specific proteins which can be detected in the blood. For example, prostate and ovarian cancer burden can be monitored using serum levels of prostate specific antigen (PSA) and CA-125, respectively. Moreover, thanks to advances in technology it has become possible to extract and characterise circulating tumour cells (CTCs) and circulating tumour DNA (ctDNA) from the blood. These provide a measure of tumour burden, as well as details of the tumour’s molecular composition, and can, for example, give early warning signals of developing drug resistance (e.g. [10, 11, 12]). However, a key issue with blood-based measures is that they may not be representative of the actual tumour dynamics, and can be confounded by other patient-specific factors, such as age or diet [8, 9]. For example, while a PSA level greater than

4ng/mL is indicative of prostate cancer, the association with actual tumour burden is complex, and new metastases can emerge without a change in PSA level [13]. CTC- and ctDNA-based methods are potentially more robust to such issues, but are only just starting to make their way into the clinic [9, 14].

1.1.5 Outcomes: the challenge of recurrence

Based on an integration of these measures, the patient's response is evaluated. In clinical trials, this is typically done using the response evaluation criteria in solid tumors (RECIST), which identify four categories of response [15]: Complete response (CR), partial response (PR), stable disease (SD), and progressive disease (PD). Assessment is based on the sum of the longest diameter (SLD) of the target lesions. For example, CR is a disappearance of all target lesions, whereas PD indicates a greater than 20% increase in SLD relative to the nadir, or the appearance of a new lesion. The *nadir* is the smallest SLD recorded since treatment started. Outside of clinical trials, evaluation of response is more subjective and may also take into account the patient's symptoms [15, 8].

If the tumour shrinks or even disappears by the end of treatment (CR or PR), then the patient is considered to be in *remission*. At this point the patient continues to be monitored for signs of *recurrence* of the disease, and may take *maintenance therapy* to try and eradicate any remaining disease (Figure 1.1). If the patient is cancer free for several years, then they may be considered *cured* [8].

If the tumour does not respond, or even grows under treatment (*refractory disease*) or recurs, then a *second-line treatment* is started (Figure 1.1). This follows similar principles as the first-line treatment, but may use different, or additional, systemic therapies. The time interval between starting a particular treatment and the disease getting worse (again) is termed time to progression (TTP), and unfortunately often becomes shorter the longer a disease has been treated (Figure 1.1; [8]).

1.1.6 Toxicity

Due to the close similarity between cancer cells and normal body cells, most anti-cancer drugs also damage healthy tissue resulting, for example, in rashes, nausea, or anaemia. So-called Phase I studies map out the toxicity spectrum of a novel drug, and establish the *maximum tolerated dose* (MTD) to be administered. Side-effects are constantly monitored during therapy and, if necessary, the drug dose is adjusted accordingly, or treatment is delayed until it is safe to resume (Figure 1.1; [8]).

1.2 Therapy scheduling: evolving paradigms

Recurrent and refractory disease are a major challenge in modern cancer therapy. For example, in ovarian cancer, first-line treatment is extremely effective, achieving responses in up to 80% of patients [16]. However, 60% of these patients will see their disease return within 3 years, at which point it becomes increasingly harder to treat [17]. For stage III or IV disease, 5-year survival rates are currently below 30% [18].

The main strategy to address this issue has been the development of new drugs which either target the resistance mechanism, or kill the cell through a different route. While great progress has been made, for example with the development of targeted therapies, a “magic bullet” remains elusive. As an alternative, but complimentary, approach, researchers and clinicians have explored the question of how the benefit of existing treatments may be maximised through strategic treatment scheduling (Figure 1.2). In the following, we will review the key paradigms, their underlying mathematics, and the evolving understanding of resistance which they reflect. For this discussion, in particular in Sections 1.2.2 - 1.2.5, we will closely follow the excellent review by Benzekry et al [19].

1.2.1 The mathematics of treatment scheduling

A key challenge in schedule optimisation is that the number of possible candidates is usually very large. For example, there are almost 600,000 ways in which 6 doses of treatment could be given over a 6 week time frame (assuming no more than one dose a day, and excluding weekends). Mathematical modelling has played a pivotal role in helping to overcome this issue, and is a large and active area of research. An exhaustive review is beyond the scope of this thesis, but we will highlight those papers which have been important in shaping the thinking about scheduling, and the tools which we will employ in this thesis.

Mathematically, the contributions to treatment scheduling may be broadly classified according to:

1. The nature and assumptions of the tumour model: Tumour models may be *discrete*, if every cell is accounted for individually, or *continuous*, if it can be assumed that cell densities are sufficiently large. In addition, they may explicitly account for space, or not. Continuous, non-spatial models are most common, and typically describe the tumour’s temporal evolution using one or more *ordinary differential equations* (ODEs). Key spatial modelling techniques are *partial differential*

equation (PDE) models, *agent-based models* (ABMs), or hybrids of these (see also the excellent reviews by Araujo et al [20] and Altrock et al [21]). A further challenge arises in linking these models of tumour growth in individual patients to the heterogeneity of dynamics observed in clinical cohorts in practice. This is the focus of so-called *Tumour Growth Inhibition (TGI) Models*, which use statistical models such as Cox-regression or hazard functions to link tumour models (typically simple analytic or ODE models) to clinical measures of response and outcome. TGI models have found particular application in pharmacometrics and industry and have been, for example, used to inform clinical trial design (see Benzekry [22] and Bruno et al [23] for reviews). In order to represent variability between patients in a clinical cohort these approaches often make use of *mixed effect models* (also referred to as *population models*) in which virtual patient cohorts are simulated by drawing a subset of model parameters from statistical distributions [24, 25].

2. The objective function and its constraints: Most commonly, the objective is to minimise tumour burden, either at the nadir or at some end point. Alternatively, approaches may minimise cumulative tumour burden, the size of a particular subpopulation (e.g. resistant cells), or toxicity, where the latter can also be implemented as a constraint (see Engelhart et al [26] and Shi et al [27] for reviews). Multi-objective approaches are also used, and can help to regularise the optimisation problem [27, 28]. A further important distinction is whether optimisation is carried out only once for the complete treatment period (*open-loop control*) or whether the schedule is regularly updated based on the tumour's response (*closed-loop control*; [27]). This is a topic we will discuss further in Chapter 4.

3. The optimisation method: Scheduling can be framed either as a *discrete* or a *continuous optimisation problem*, depending on whether treatment is represented as a discrete set of doses given on specific days, or is allowed to vary continuously. Continuous problems are often solved using optimal control theory. The Pontryagin Maximum Principle enables the derivation of equations from which, either numerically or analytically, an optimal solution can be computed (see Sharp et al [29] for an introduction). Discrete problems are typically solved numerically, using, for example, genetic algorithms. Alternatively, optimisation may be based on less formal, heuristic or geometric arguments. We refer to Shi et al [27] for an excellent review of this topic.

1.2.2 The log-kill hypothesis: Total therapy cures leukaemia

The first theory for treatment scheduling was proposed by Skipper, Schabel and Wilcox [30] in the 1960s. They compared the effect of treating leukaemic mice with either a daily, low dose of chemotherapy, or a single high dose. They found that the single dose resulted in a greater tumour burden reduction and a higher probability of cure, even though the low dose schedule administered a greater cumulative dose. To explain their observations, they tested the effect of a single dose of treatment for different tumour sizes (numbers of leukaemic cells), and found that always the same proportion of cells was killed [30, 31].

Based on their observations, they formulated the *log-kill hypothesis*, which states that the fraction of cells killed by treatment depends solely on the treatment dose [30]. Mathematically, this may be formulated as follows: Let $N(t)$ be the number of tumour cells as a function of time, t , and assume that $N(t)$ is a continuous function. Further, let $D(t)$ be the level of treatment over time. Then the dynamics of the tumour in response to treatment may be described by the following ODE:

$$\frac{dN(t)}{dt} = rf(N, t)N - d(D)N, \quad (1.1)$$

where r is the cell division rate, $f(N, t)$ is the fraction of dividing cells, and $d(D)$ is the rate at which tumour cells are killed at treatment level D . The relationship between d and D is called the *drug-effect curve*.

Based on their results, Skipper et al [30] concluded that to maximise the chance of cure, treatment should be given at the highest dose which toxicity permits, and continue for as long as possible, even if the tumour disappears (Figure 1.2). While at this point it had been possible to achieve tumour size reductions with chemotherapy, these were only partial and short-lived [4, 5]. Based on Skipper's observations, Frei and Freireich proposed a new regimen in which four chemotherapies were given simultaneously at maximum dose and at maximum frequency to treat children with leukaemia. The side effects were excruciating, but the regime consistently produced long-lived remissions, and in some cases even cures [32]. This revolutionised the treatment of leukaemias, and established chemotherapies as a treatment option for cancer. However, despite ever more aggressive regimes, referred to by some as "total therapy" [33], some leukemia patients still relapsed, and the success in solid tumours remained limited [4, 5].

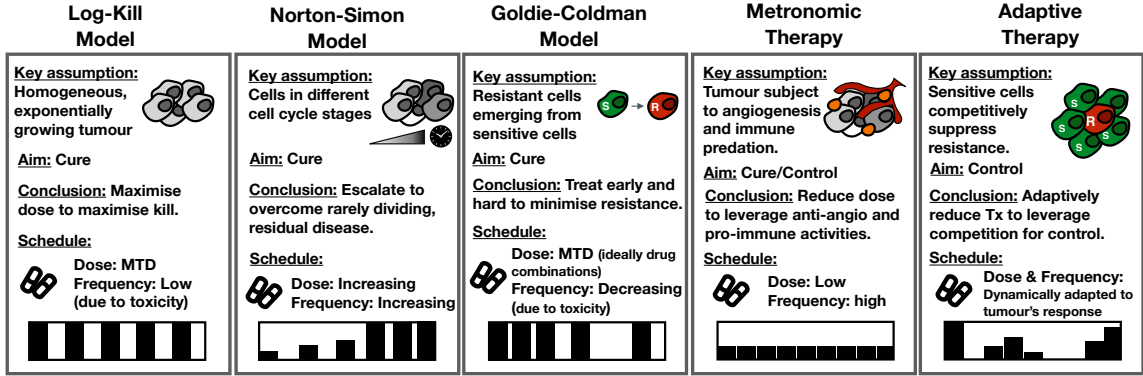


Figure 1.2: Key theories in cancer treatment scheduling. Figure re-drawn and modified with permission from Benzekry et al [19].

1.2.3 The Norton-Simon Hypothesis: Kinetic resistance

In the 1960s and 70s, detailed studies of tumour growth kinetics were carried out (e.g. [34, 35, 36]). These showed that, unlike leukaemias, solid tumours grew in a saturating fashion, where the fraction of dividing cells decreased with the size of the tumour. Possible ways to describe this mathematically were explored, such as: i) *logistic growth* ($f(N) = (1 - \frac{N}{K})$), which assumes a simple, linear density-dependence, ii) the *Gompertz model* ($f(N) = \log K/N$), which was adopted from the modelling of human life spans [37], or iii) a *power-law model* ($f(N) = N^{\nu-1}$), where ν can be seen to reflect the surface tumour's geometry (e.g. $\nu = 2/3$ for a sphere, which is a special case of the *von Bertalanffy model*²). Here, $K(> 0)$ denotes the tumour's *carrying capacity*. The Gompertz and power-law models were found to provide especially good descriptions of tumour growth (e.g. [39, 40, 41]; see also Chapters 2 & 5).

In 1977, Norton and Simon [42] pointed out that saturating growth has important implications for treatment scheduling (Figure 1.2). Many cancer drugs act on dividing cells, so that the number of cells killed will depend on the tumour's growth rate. To explore how this affects scheduling, they proposed the following ODE model:

$$\frac{dN(t)}{dt} = rf(N, t)N - d(D)rf(N, t)N = rf(N, t)(1 - d(D))N. \quad (1.2)$$

If one assumes that $f(N, t)$ is a constant, as it would be in exponentially growing leukaemias, then this recovers the model by Skipper et al (Equation (1.1)). However, if one assumes saturating growth (e.g. Gompertzian growth), then Norton and Simon showed using numerical simulation that treatment induces maximum size reductions

²The full von Bertalanffy model reads: $f(N) = N^{2/3} - b$, where b is the cell death rate [38].

at intermediate tumour sizes (at the maximum of $f \times N$). Importantly, this implies that both large as well as small tumours experience a certain protection from treatment, termed *kinetic resistance*. At large sizes this occurs because only a small fraction of cells is actively dividing, and at small sizes because it is unlikely that one treats at exactly the time when one of the few remaining cells divides [42].

The aggressive approaches motivated by Skipper's work often require dose reduction over the course of treatment due to toxicity. Using numerical simulations of their model, Norton and Simon [42] showed that this is problematic because it means that by the end of treatment, when the tumour population is small, treatment is insufficient to eradicate the disease. Moreover, they argued that the benefit of high initial doses is limited due to kinetic resistance. Instead they proposed that treatment should gradually be escalated so that small tumours are treated at MTD and maximum frequency [42]. Moreover, when several drugs are administered, these should be given sequentially to allow maximal intensification of each drug (maximum dose and frequency).

The guidelines proposed by Norton and Simon have been successfully tested in a number of clinical trials (see [43] for an overview). For example, in breast cancer it was found that densification of adjuvant treatment, from a three week to a two week interval, decreased the annual probability of disease recurrence by 26% and the annual probability of death by 31% [44]. As such, this work represents an important success, and has inspired further theoretical investigations of kinetic resistance [4, 19]. In particular, explicit modelling of the cell cycle using, for example, ODE or age-structured PDE models, has shown possible resonance and interference effects between cell and treatment cycles, with implications for tumour kill and toxicity (see Gaffney [45] for a review).

1.2.4 The Goldie-Coldman model: Molecular resistance

Around the same time, Goldie and Coldman considered how to tackle a different type of resistance: *molecular resistance*. Through a variety of mechanisms, cancer cells can reduce, or even eliminate, their sensitivity to treatment. While the details depend on the particular drug and cancer, the routes to resistance may be broadly summarised as follows (adopted from Holohan et al [46]):

1. **Drug elimination and metabolism:** Drugs can be pumped out of a cell through efflux pumps (e.g. ATP-binding cassette transporter proteins), or can be deactivated, or prevented from being activated, by metabolic enzymes.

2. **Alterations in drug targets or pathways:** Mutations in the targeted proteins or pathways can leave the drugs unable to bind, or stop their down-stream effects (e.g. deregulation of apoptosis).
3. **Oncogenic bypass and pathway redundancy:** Many cellular tasks can be carried out by multiple pathways, and this redundancy ensures that function can continue should one of them fail, for example, due to poisoning or infection. Cancer cells can exploit this redundancy, to bypass pathways targeted by treatment.
4. **Other survival mechanisms:** Cells can also make use of generic survival enhancing mechanisms, such as up-regulation of pro-survival signalling (e.g. EGFR signalling), autophagy, or transition to a more adaptable and motile, mesenchymal phenotype (epithelial to mesenchymal transition).

These various resistance mechanisms may become activated through genetic mutations, or simply through changes in gene expression (*epi-genetic changes*), and often cells employ more than one resistance modality. Also, within a tumour, different cells may make use of different mechanisms, which is termed *poly-clonal* resistance. For further details, we refer to the reviews by Holohan et al [46], and Burrell and Swanton [47].

To investigate how to best tackle molecular resistance, Goldie and Coldman [48, 49, 50] explored stochastic branching process models, in which they assumed that resistance arises randomly through mutation. Their numeric and analytic studies showed that the probability of there being at least one resistant cell within the tumour is very sensitive to tumour size and, importantly, increases non-linearly over a short period of time as the tumour expands [48]. They extended their model to consider more realistic growth models [50], and the possibility of multi-drug resistance [49].

Based on their results they concluded that treatment should be started as early and aggressively as possible to minimise the chance of resistance (Figure 1.2; [48, 49]). Moreover, if toxicity does not permit simultaneous administration of multiple drugs, they proved analytically that the treatments should be rapidly alternated to maximise the chance of cure [51]. By rapid alternation, the sizes of different sub-populations, which are resistant to one of the treatments, can be kept as small as possible, thereby reducing the chance of multi-drug resistance [51, 52].

However, clinical testing of the Goldie-Coldman hypothesis has yielded mixed results. Multiple analyses have indeed confirmed that drug-tolerant, or resistant,

cells often exist prior to any treatment (see [47] for a review), and some studies in breast cancer have reported that a delay between surgery and adjuvant therapy is associated with a shorter TTP [52]. However, Bonadonna et al [53], for example, found that sequencing of drugs, rather than alternation, achieved longer survival time in breast cancer.

As such, the management of molecular resistance remains an important, and open, question in oncology. The work by Goldie and Coldman has been expanded, for example, by Day [54], and later Murray [55], to consider possible asymmetry in the efficacy of the two treatments. Such asymmetry, resulting in an improvement in some, but a worsening in other, patients could explain the mixed clinical success [55]. More recently, Chmielecki et al [56] derived a 2-population model of drug-sensitive and resistant cells from *in vitro* experiments of targeted treatment of lung cancer. Through numerical simulations they predicted that the best control of resistant cells would be achieved by a combination of high and low doses, as the short bursts of intense treatment would help to kill resistant cells whilst also being better tolerated [56]. The proposed schedule outperformed standard of care in a mouse model [56], but did not have the desired effect in patients [57]. Due to these mixed results, there is a growing understanding that while aggressive treatment may help to avoid resistance, it may not be the most effective method to contain it once it has emerged [58, 59, 60]. This idea will be a key focus of this thesis.

1.2.5 Metronomic therapy: Targeting the environment

The discovery that genetic mutations are a root cause of cancer was a pivotal breakthrough in cancer research, and underlies many innovations in treatment and prevention. However, increasingly it has become clear that cancer is not simply a “genetic disease” [61]. Tumours contain a plethora of normal body cells, which have been recruited either to support tumour growth, for example, by the construction of vasculature, or to suppress it (immune cells; [7]). For example, in pancreatic cancer, up to 90% of the tumour is made up of such non-tumour cells [62].

Administration of drugs at their MTD, as recommended by the prior paradigms, often requires significant rest periods between administrations to allow the patient to recover. However, in the late 1990s, Folkman, Hannahan, Kerbel and others argued that such breaks also allow the tumour vasculature to recover, thus providing the ground for recurrence [63, 64, 65]. To avoid these breaks, they proposed *metronomic therapy* (MT) in which treatment is given at a lower dose, but higher frequency, “*akin to the uninterrupted ticking of a metronome*” [66] (Figure 1.2).

MT has since been investigated in a series of modelling, pre-clinical and clinical studies (see also [19] and [67] for in-depth reviews). Hahnfeldt et al [68, 69] developed two simple ODE models of the effect of treatment on angiogenesis, and showed analytically that frequent, low-dose treatment indeed outperforms an MTD approach. d’Onofrio et al [70], and Ledzewicz and Schättler [71], have carried out more formal optimal control analyses, and demonstrated, for example, that MT may benefit from personalisation to individual patients [71]. Important experimental validation was provided by two seminal studies in 2000, which showed that metronomic chemotherapy, in combination with anti-angiogenic therapy, could eradicate tumours *in vivo*, even if these tumours had previously developed resistance to chemotherapy [63, 64]. One challenge in developing metronomic schedules is identification of doses, which can be given at a high frequency, but remain effective against the tumour (*minimum effective dose*; [72]). Recently, a group in France has demonstrated how mathematical modelling can help to address this problem. They carefully developed and calibrated an ODE model of vinorelbine chemotherapy, and used simulations to derive a metronomic schedule, which has yielded promising results in a Phase I study in lung cancer and mesothelioma [72, 73].

Moreover, investigation of MT has revealed other benefits. It is cheaper and better tolerated, thus it is easier to carry out in poorer countries [67]. Furthermore, it has immuno-stimulatory effects. The lower doses cause less damage to immune cells, and induce a more immunogenic form of cell death [66, 74]. This observation has motivated modelling efforts which also study the effects of therapy on the immune system (e.g. [75, 76, 77]). In addition, modelling by Benzekry and Hahnfeldt [78], using a combination of ODEs to describe the primary tumour and PDEs to capture the metastatic burden, has suggested that MT may be better able to control metastatic disease than an MTD approach. While MTD approaches remain the mainstay in the clinic, successes in Phase I and II trials are spurring a growing interest in metronomic strategies, and are starting to change the paradigm that “more is better” [67].

1.2.6 Pharmacokinetics: Challenges of drug delivery

While this thesis will focus on the effect of treatment on the tumour, the so-called *pharmacodynamics* (PD), a further crucial aspect of scheduling is *pharmacokinetics* (PK). This describes how a drug circulates throughout the body, and how it is removed by the different drug clearance mechanisms [79]. Insufficient drug delivery into the tumour is an important obstacle in therapy development [80]. For example, in Frei and Freireich’s study, most of the recurrences originated not in the bone marrow, but in

the patients' brains, a *pharmacologic sanctuary* protected by the blood-brain barrier [32]. As such, pharmacokinetic modelling and quantitative systems pharmacology (integration of PK with PD models) have become key steps in the industrial drug development process [79, 22, 23]. For example, Morrissey et al [81] recently developed a combined PK and TGI model for the immunotherapy atezolizumab. Using model simulations they compared the approved protocol which administered treatment once every 3 weeks, with alternative protocols which administered treatment every 2, or every 4, weeks. Their simulations predicted that all three schedules would achieve comparable drug exposure and tumour response, which helped pave the way for FDA approval of the bi-weekly and four-weekly schedules [81, 23].

In addition, drug delivery into tumours is impeded by the fact that tumours are often poorly perfused, due to their disordered vasculature. PDE and ABM modelling has provided important insights into the dynamics of tumour angiogenesis and its impact on nutrient and drug distributions within the tumour (e.g. [82, 83, 84]).

1.2.7 Safety

Finally, a key factor in the development of treatment schedules is the management of side effects. As we discussed in Section 1.2.3, the proposition by Norton and Simon [42] to gradually escalate treatment over time was motivated by the observation that toxicity often prevents the completion of the regular, high-dose schedules foreseen by Skipper et al [30]. Over time, these verbal models have been formalised into mathematical models, such as the seminal paper by Friberg et al [85] which provides an ODE model of the haematological toxicity of chemotherapeutic agents. Today, models play an important role in optimising the safety of new agents and of drug combinations [22, 86]. For example, Meille et al [87] recently used a multi-compartment TGI ODE model to optimise the scheduling of the combination of docetaxel and epirubicin in breast cancer, which prior studies had shown to be effective but which had also caused lethal side effects in some patients. A subsequent clinical trial validated the model's ability to accurately predict the time-course of different markers of haematological toxicity, and demonstrated that the model-derived schedule indeed achieved a more acceptable toxicity profile [88, 87].

A key challenge in managing side effects is inter-patient variability in PK parameters, such as the rate of drug clearance. This variability arises from physical (e.g. body weight), genetic (e.g. drug transporter polymorphisms), and environmental factors (e.g. diet or concurrent medications) and means that a dose/schedule may be well tolerated in one patient but can result in toxic drug levels in others [86]. In

order to more carefully tailor treatment to the individual patient, so-called *therapeutic drug monitoring* (TDM) and *adaptive dosing* approaches have been developed [80, 86]. These monitor drug levels (or surrogates) and/or markers of toxicity (e.g. neutrophil or platelet counts), and dynamically adjust treatment to ensure both efficacy as well as safety. There is a strong tradition in this area that goes back as far as the 1970s [89], of using mathematical modelling to guide decision-making. In oncology, these approaches are most established in treatment with the chemotherapeutic agent methotrexate, where there is a high risk of lethal side effects, in particular in paediatric patients [90, 91, 92, 93, 94]. For example, Evans et al [90] found in a clinical trial that adaptive dosing based on predictions of a simple PK model, calibrated with data from the first treatment cycle, was superior to dosing based simply on body weight. With the increasing role of personalised treatment in oncology, and the use of combination therapies, these approaches are also finding interest in other clinical settings [95, 86]. For example, in the aforementioned work by Meille et al [87], the model was fitted to the data from the first treatment cycle using a Bayesian framework and used to generate personalised predictions to decide on each patient's subsequent treatment plan. But not all adaptive dosing approaches succeed, as illustrated by the clinical study by Ratain et al [96] which did not find any additional benefit in model-guided adaptive etoposide treatment. Importantly the authors were able to attribute this lack of benefit to the fact that the chosen model did not correctly capture the treatment's impact on myelosuppression, which shows that great care needs to be taken in choosing *how* to adapt treatment and *the reasons* for doing so [96].

1.3 Adaptive therapy

Comparing the development of cancer therapy with the report by Melander [2] at the start of this chapter, the similarities are striking: Treatment is initially effective, but loses its efficacy over time; maximising exposure through more effective and intensive treatment only brings limited, usually short-lived, improvements. These similarities were first pointed out in 2009 by a group headed by Robert Gatenby, who argued that the problem in both cases was the same: Darwinian evolution [60, 98]. They proposed a new treatment approach, called *adaptive therapy* (AT; Figures 1.2 & 1.3). In the following, we will review its motivation, and the current state of this relatively young field. This will motivate the questions addressed within this thesis.




Figure 1.3: The principles of AT. (a) Tumours as ecosystems in which drug-sensitive and resistant cells compete for resources. (b) Aggressive treatment causes the competitive release of resistant cells. (c) AT dynamically reduces treatment to delay competitive release whilst maintaining the tumour at a tolerable burden. Illustrated is the AT algorithm used in the first-in-human AT trial in androgen-deprivation treatment of metastatic prostate cancer, carried out by Zhang et al [97].

1.3.1 The rationale behind adaptive therapy

Cancer as an eco-evolutionary system: The paradigms reviewed in the previous section reflect an evolving understanding of cancer. As new theories and treatments have emerged, it has become increasingly clear that cancer, in particular metastatic disease, is a highly complex system. No two tumours, even of the same type, share the same molecular make-up [99]. Nevertheless, all share certain “hallmarks” [100]. They are composed of cancer cells in different stages of the cell cycle, with different sensitivities to treatment, which interact with normal tissue cells in order to engineer themselves an environment which supports their growth, and which protects them from the body’s defence mechanisms, and from clinical treatment [100, 101].

To explain these observations, an eco-evolutionary theory of cancer has begun to emerge [99, 102, 103, 101, 104]. In this theory, tissues are viewed as ecosystems in which cells compete with each other for space and resources, are predated upon by the immune system, and can enter symbiotic or parasitic relationships with each other. Genetic mutations and epi-genetic changes generate diversity between cells. The removal of cells through turnover, starvation, predation, or wounding, then induces selection for the fittest cell type, or the fittest collective of cells [99, 102, 103, 101, 104].

The competitive release hypothesis: Based on this “Darwinian” view of cancer, Gatenby and others [58, 59, 98, 104] have argued that the failure of aggressive MTD approaches in metastatic cancers is no surprise. The fact that cancers respond, at least

temporarily, shows that resistant cells are usually suppressed by their environment, for example, by competition with other cancer cells (Figure 1.3a). In fact, resistant cells may even be at a fitness deficit prior to treatment due to the fact that they may be paying a *resistance cost*, such as the energy required to run drug-efflux pumps (see also Chapter 2). However, aggressive treatment rapidly removes large portions of the tumour, leading to the *competitive release* of these resistant cells, and their rapid expansion (Figure 1.3b).

Control and conquer: Cancer is a devastating disease, and it is only natural that research and clinical practice have gravitated towards radical approaches for its eradication. However, a growing body of literature, and the success of MT, show that “less may be more”, especially if cures are unlikely. Indeed, as the quote by Melander at the beginning of this chapter illustrates, similar conclusions have been drawn also in other fields, such as in pesticide or antibiotic resistance [105, 106, 107]. For example, in the management of insecticide resistance, untreated “refugee patches” are used to maintain sensitive insect populations, which breed and compete with emerging resistant strains [108].

Drawing on this analogy, Gatenby [98] proposed that in cancers, in which cures are rare, the focus should switch from eradication to management. This means that rather than trying to eliminate resistant cells, the aim should be to slow their expansion, and ideally turn the cancer into a chronic disease. To do so, treatment should be de-escalated in order to maintain sensitive cells, which will competitively inhibit resistant cells (Figure 1.3c). Similar to the MT paradigm, this de-escalation may also help to target the tumour’s environment, and reduce toxicity and financial costs [98, 60].

Evolving treatment for an evolving disease: The concept that treatment de-escalation may forego the chance of a cure, but increase survival, has previously been suggested by others, in particular theoreticians. Most notably, Martin et al [58] proved, using optimal control theory, that in a simple ODE model of sensitive and resistant cells the best way to avoid the tumour exceeding some maximum size is to maintain the tumour at the largest tolerable size for as long as possible [58]. Similar conclusions were also drawn by researchers investigating MT, such as Hahnfeldt et al [69], Monro and Gaffney [59], and Carrere [28]. Moreover, in certain clinical settings, such as prostate cancer [109, 110, 111] or melanoma [112], clinical trials have investigated intermittent treatment schedules where treatment is withdrawn at regular points in time (see also Section 2.5). However, in all the aforementioned cases,

treatment is assumed to be administered according to fixed, pre-determined schedules. But Gatenby et al [60, 98] argued that if tumours are evolving systems, then our treatment also has to evolve. This means that timing and dosing are driven not by a calendar, but by the current state of the individual patient’s tumour, i.e. *adaptive* therapy (Figure 1.3c). As such, AT advocates dynamic personalisation of treatment similar to the TDM or adaptive dosing approaches discussed in Section 1.2.7. The novelty of AT is that dosing is driven not by considerations of drug exposure, safety, and maximal tumour size reduction, but by eco-evolutionary principles seeking to maintain tumour control via competition.

1.3.2 The current state of the field

Pre-clinical experiments: To illustrate AT, Gatenby et al [60] proposed a treatment algorithm which increases dosing when the tumour grows, but decreases, or even withdraws it, if the tumour regresses. In their seminal paper they demonstrated the ability of this algorithm to outperform MTD treatment in a mouse model of ovarian cancer chemotherapy treatment [60]. Over the past ten years, these results have been corroborated in *in vivo* models of breast cancer [113] and melanoma [114].

In addition to testing whether AT works, research has also investigated how it works. Enriquez-Navas et al [113] showed *in vivo* that AT helps to normalise the tumour vasculature, thereby improving drug perfusion. Bacevic et al [115] demonstrated that they could control drug resistant cells in a spheroid, but not in a less constrained 2-D *in vitro* model of targeted therapy, indicating the importance of spatial competition [115]. Smalley et al [114] observed that in both mouse and human samples of melanoma there was a re-sensitisation of the tumour cells after a drug holiday. Finally, a recent pre-print by Farrokhian et al [116] has demonstrated competitive release *in vitro*, although the dynamics appear to be more complex than predicted by current models.

Modelling: Mathematical modelling has played an important role in the development of AT. We will review the details of prior modelling studies in the chapters to which they pertain, and so restrict ourselves here to a broad overview. The first two models of AT, an analytic, stochastic model, and a discrete, computational model, were presented in the seminal study by Gatenby et al [60]. They provide proof-of-principle illustrations of how AT can improve control of resistant cells through competition, and by modifying the tumour environment. Silva et al [117] simplified the computational model by dropping the tumour-environment interactions, and

focussing just on the competition between drug sensitive and resistant cells. Subsequently they discussed how AT can be improved by inducing a resistance cost via drugs which starve tumour cells [117]. Zhang et al [97] introduced a 3-population ODE model of the eco-evolutionary dynamics of androgen-deprivation treatment in prostate cancer, which helped to motivate the first clinical trial of AT. This model has been analysed more formally by Cunningham et al [118, 119], and has been extended to include the effect of chemotherapy by West et al [120]. West et al [121, 122] have also studied a more game-theoretic formulation of this model. An alternative, stem-cell focussed, model of prostate cancer was recently proposed, and carefully calibrated with clinical data, by Brady-Nicholls et al [123]. Via numerical simulations the authors made quantitative predictions on the benefit of AT for prostate cancer patients, and suggestions of how it may be further improved [123].

To gain a deeper understanding of the mechanisms behind AT, Bacevic et al [115]³ studied a 2-population model, first as an ODE, and subsequently as an ABM, and highlighted the importance of spatial constraints on resistant cell growth. This conclusion was corroborated by Gallaher et al [124], who also considered the role of intra-tumoral heterogeneity, cell movement, and phenotypic drift in an off-lattice ABM. Concurrent with publication of work from this thesis [125], Viossat and Noble [126, 127] have presented a general theory of AT in 2-population ODE models, inspired by work by Hansen et al [1]. The latter have recently also extended their work to consider the role of stochasticity in AT [128], and how competitive inhibition may be maximised [107, 129].

A final line of research has focussed on how therapy should be modulated. Smalley et al [114]⁴ used an ODE model to directly guide AT decision-making in a melanoma mouse model. Newton and Ma [130] and Gluzman et al [131] have recently illustrated how optimal control theory may be used to formalise decision-making. Finally, West et al [132], Yoon et al [133], and Ma and Newton [134] have proposed frameworks for multi-drug AT.

Clinical trials: Five clinical trials of AT have been launched so far, all carried out at the Moffitt Cancer Center. The most advanced of these is a trial in metastatic castration resistant prostate cancer, which investigates an adaptive algorithm for the administration of androgen-deprivation treatment (clinicaltrials.gov identifier: NCT02415621; [97]). This algorithm, which was derived via integration of numerical

³Modelling work carried out by Dr Robert Noble.

⁴Modelling work carried out by Dr Eunjung Kim.

simulation of the game-theoretic model by Zhang et al [97] and clinical considerations, administers treatment until the blood PSA level has dropped by 50% relative to the pre-treatment value. Subsequently, treatment is withdrawn and only restarted if the PSA level again exceeds pre-treatment levels (see Figure 1.3c for an illustration). While data collection has not yet finished, an interim analysis in 2017 [97] reported a 10 month increase in median TTP compared to standard-of-care continuous treatment, and a 53% reduction in cumulative drug usage. A trial using the same AT algorithm in the chemotherapeutic treatment of rhabdomyosarcoma is ongoing (NCT04388839), and one in osteosarcoma, is in preparation [135]. A trial using this algorithm in targeted treatment of thyroid cancer had to be stopped in September 2020 due to failures (NCT03630120). Motivated by insights from theoretical studies [97, 120, 126], a modified form of the algorithm, which allows PSA to exceed twice the baseline value before treatment is applied and also adapts which of two lines of treatment is used based on testosterone levels, is currently being tested in castration sensitive prostate cancer (NCT03511196). Finally, the experimental and modelling work by Smalley et al [114] has provided the base for a trial of AT (albeit of a pre-determined 3w-on/2w-off/3w-on protocol) in targeted treatment of melanoma (NCT03543969).

1.4 Outline of the thesis

Motivated by the aforementioned promising pre-clinical and clinical results, there is growing interest in extending AT to other cancers. However, in order to do so it is crucial to understand the factors which determine whether a patient will benefit from AT, or not, and how to best adjust treatment. In this thesis we present four studies which integrate mathematical modelling, clinical data analysis, and biological experiments to investigate these questions. The first two of these are more theoretical, and leverage historical data from prostate cancer, in order to better understand the factors which are important to AT. In the second part of this thesis, we then seek to translate AT to PARP inhibitor (PARPi) maintenance therapy for ovarian cancer. This work was carried out in collaboration with Professor Mehdi Damaghi at the Moffitt Cancer Center, and includes wet-lab experiments done by myself:

- **Chapter 2:** We analyse a simple 2-population ODE model in order to explore the role of resistance costs in AT. Our analysis shows that the impact of resistance costs depends strongly on the characteristics of the tumour, as described

by the model parameters, and we identify cell turnover as an important, but previously overlooked, factor in AT.

- **Chapter 3:** We extend our prior work to consider the impact of space on AT. We develop a simple ABM and show that, while the ABM qualitatively replicates our ODE results, there are significant quantitative differences. Moreover, we carefully dissect the spatial competition between cells in the ABM. This shows that the tumour's spatial organisation shapes intra-tumoral competition and is a further important factor in AT.
- **Chapter 4:** We then ask whether we can use mathematical modelling to directly guide treatment decision-making in AT. We develop a model-guided framework which uses an ensemble of six ODE models to inform decision-making. In collaboration with Professor Damaghi, we test this framework *in vitro* in the context of PARPi maintenance therapy for ovarian cancer. While our approach shows promise, the experiments indicate that continuous MTD treatment does best.
- **Chapter 5:** To conclude, we investigate why the experimental results in Chapter 4 suggest that MTD treatment does best. We carry out *in vitro* experiments in order to build and calibrate a more detailed model of PARPi action. Our analyses indicate that treatment breaks are problematic in this case, and we discuss preliminary results on a possible alternative scheduling strategy.

To conclude, in Chapters 6 & 7 we will critically review our results, and discuss how they may be extended.

Chapter 2

On the Role of Resistance Costs in Adaptive Cancer Therapy

2.1 Introduction

An important topic in the development of novel therapies is the identification of patient selection criteria. Success in prostate cancer has spurred interest in extending AT to other cancers, and, for example, trials in melanoma or rhabdomyosarcoma are already ongoing (NCT03543969 and NCT04388839, respectively). However, how do we identify which cancers, patients and drug combinations will benefit from AT and which will not? In a recent review paper, Gatenby and Brown, two key figures behind AT, stated that: “*In evolutionary cancer treatment, a major component of Darwinian dynamics is the cost of resistance.*” [104]. Motivated by this statement, the aim of this chapter is to explore whether the presence of a fitness cost of resistance would provide a useful selection criterion for AT. To do so, we study a 2-population ODE model to investigate whether the presence of a cost is a necessary condition for AT to extend the TTP beyond that of a standard-of-care CT. We find that for tumours close to their environmental carrying capacity, such a cost of resistance is not required. However, for tumours growing far from carrying capacity, a cost may be required to see meaningful gains. Notably, we show that in such cases it is important to consider the cell turnover in the tumour and we discuss its role in modulating the impact of a cost of resistance. Overall, this chapter helps to clarify under which circumstances AT may be beneficial, and suggests that turnover may play an unexpectedly important role in the decision-making process ¹.

¹The work in this chapter is based on the publication Strobl et al, 2021. Turnover modulates the need for a cost of resistance in adaptive therapy. *Cancer Research*. 81(4), 1135–1147. Figures and parts of the text have been re-used.




Figure 2.1: The three types of experiments typically conducted to test for a cost of resistance. All experiments are carried out in the absence of drug, either *in vitro* or *in vivo*. (a) Mono-culture experiments study the resistant strain in isolation, examining changes in growth rate, size, migration rate etc. (b) Competition experiments compare the abundance of sensitive and resistant cells over time. The higher the cost, the faster the resistant population is expected to disappear. (c) Reversal experiments examine the rate at which drug resistance is lost, if the drug is withdrawn. To do so, the resistant population is cultured in a drug-free environment and its drug response is tested at regular time intervals.

2.1.1 What is the cost of resistance?

The *cost of resistance* refers to the notion that a resistance adaptation not only confers a fitness advantage under drug exposure, but that it might also come at the *cost* of a fitness disadvantage in a treatment-free environment [98, 136]. This is because it is assumed that there is a trade-off between the traits required for life under drug, and life under drug-free, conditions, and that a cell cannot be perfectly adapted to both. For example, the running of drug-efflux pumps requires energy which leaves the cell with fewer resources to support proliferation [137].

Resistance costs are experimentally measured in one of three ways (Figure 2.1): i) Mono-culture experiments, which measure the growth rate (or other characteristics) of the resistant population in the absence of drug in isolation, ii) Competition experiments, which study the relative frequency of sensitive and resistant cells over time when cultured together, and iii) Reversal experiments, which explore how long it takes for resistant cells to lose their resistance mechanism if grown in absence of drug for long times. Interestingly, while it has only recently become a focus of study in cancer, it is a concept that dates back to Dobzhansky [138] in the 1950s, and has received significant attention in the agricultural pest and antibiotic resistance communities [139, 140].

2.1.2 A brief review of resistance costs in cancer

While a comprehensive review is beyond our scope, we will briefly provide examples of experimental evidence for resistance costs in the cancer literature. This will serve to illustrate that while costs have been reported by a number of authors, the magnitude of these costs appears to depend strongly on the drug, the experimental model system, and even the cell line. This explains the controversy around the topic and shall motivate the work in this chapter.

To the best of our knowledge, some of the largest costs have been reported by Gallaher et al [124] who found that the doubling time of chemotherapy (doxorubicin) resistant MCF7DOX breast cancer cells is increased by 50% in *in vitro* 2-D cell culture when compared to their sensitive counterparts. Consistent with this significant fitness cost, the authors further reported that in competition experiments the sensitive cells rapidly out-competed the resistant cells in absence of the drug [124]. That being said, others have reported more moderate differences. For example, Jensen et al [141] and Kanzawa et al [142] have also reported increases in doubling times in chemotherapy-resistant colorectal and lung cancer cell lines, respectively, but their cost estimates are more on the order of 20%. Even lower values have been reported in experiments on targeted therapies, with studies in lung cancer [143, 28, 144], colorectal cancer [115], or melanoma [114] reporting differences of only 5-10%. This potentially suggests that chemotherapy resistance may be more costly than resistance to targeted agents. This intriguing hypothesis was outside our scope, but warrants further investigation.

However, importantly, not all resistance mechanisms appear to come at a cost. For example, Behrens et al [145] reported that A2780 cis-platin resistant ovarian cancer cells grew even faster than the sensitive parental cell line (doubling times of 22.1h compared to 25.3h). Furthermore, Kaznatcheev et al [144] observed that even though their resistant cells grew more slowly in mono-culture, these cells were in fact supported by co-culture with sensitive cells [144]. Moreover, even within the same study, there can be significant heterogeneity between cell lines, indicating that costs also depend on the specific genetic and histologic origin of the cell [141, 114].

In addition, even if a resistance cost is present in the mono-culture growth rates, this does not necessarily guarantee a swift return to sensitivity. For example, Bacevic et al [115] found that even though their resistant cells paid a resistance cost of 10%, there was no clear dominance of sensitive over resistant cells in *in vitro* 2-D cell culture competition experiments (although there was in 3-D). Furthermore, several authors have reported that there is no return to sensitivity, even if resistant cells are cultured in drug-free medium for many months [143, 115, 114].

2.1.3 Are resistance costs necessary for AT?

In his original essay on AT, Gatenby [98] argued that resistance costs are one of the main reasons why it might be possible to control resistant cells through intra-tumoral competition. Indeed, it is intuitive that a cost would facilitate keeping resistance at bay. However, given how strongly resistance costs appear to depend on the specific environmental and genetic context of the cancer cells, two important questions arise: how much of a cost is required, and can AT succeed even if it is absent?

The first study to consider the role of resistance costs in AT was carried out by Silva et al [117]. They modelled the growth of a chemotherapy-sensitive and resistant breast cancer population using a set of two difference equations, and showed through numerical simulations that the larger the cost and the smaller the initial frequency of resistant cells, the greater the benefit of AT over aggressive treatment. This work was subsequently extended by Bacevic et al [115], who recast the model as two ODEs, and assumed that instead of growing exponentially, the cells were growing in a Gompertzian fashion. Furthermore, they explored different hypotheses on how the fitness of resistant cells (their growth rate) may change with the frequency of resistance in order to reflect different assumptions about intra-tumoral competition. The authors concluded that a relatively large cost of around a factor of 2 would be required for AT to show meaningful benefits [115]. That being said, Bacevic et al [115] also subsequently considered an ABM model in which they find that removing the cost has relatively little effect on their ability to control resistant cells with competition. Overall, this indicates that while resistance costs generally appear to confer a benefit to AT, the extent to which they are necessary to see a meaningful benefit depends on other characteristics of the tumour, such as the initial resistance fraction. Further characterisation of these factors and their interactions with resistance costs in AT remains an important, but open, question.

2.1.4 Outline

The aim of this chapter is to clarify the role of resistance costs in AT. We will use a simple mathematical model in which we divide the tumour into drug-sensitive and drug-resistant cells and model their growth with a pair of ODEs. We will compare the TTP under two treatment strategies: i) continuous treatment (CT) at MTD, representing an aggressive approach aimed at curing the cancer and often used in the dosing of targeted therapies, and ii) the AT algorithm used in the clinical trial

by Zhang et al [97]. We will compare the two treatment strategies first in the absence, and subsequently in the presence, of a cost of resistance. This will allow us to develop a mathematically rigorous understanding of AT, and of the impact of a cost of resistance. Moreover, through our analysis we show that cell turnover is a key factor to consider to understand the impact of resistance costs and to assess the ability to control a tumour with AT more generally. To conclude, we present evidence for this hypothesis from an analysis of longitudinal data from 67 patients undergoing intermittent androgen deprivation therapy for prostate cancer.

2.1.5 Note

Working in parallel to us, Professor Yannick Viossat and Professor Robert Noble [126] also studied simple 2-population models with the aim of characterising the “logic of AT” (see [127] for the final version of this work, published after submission of this thesis). We had a brief discussion in the summer of 2019, in which Professor Viossat directed me to the work by Hansen et al [1]. However, we did not realise the close overlap of our interests until I sent him a copy of a first draft of our manuscript in January 2020. As a result, there is a certain overlap between the work presented in Section 2.3.1 and their paper [126, 127], as we both independently arrived at similar conclusions from reading the manuscript by Hansen et al [1]. In acknowledgement of his contribution to our work, we included Professor Viossat as a co-author on our paper, and he helped us to refine some of our arguments in Section 2.3.1. All other work, including identification and characterisation of the role of turnover, and the fitting to patient data were carried out solely by myself.

2.2 The Mathematical Model

Tumours are heterogeneous populations of cells with differential responses to drug, indicating a degree of pre-existing resistance in most tumours [46, 146, 147]. To model this heterogeneity, we divide the tumour into two competing cell types: drug-sensitive cells, $S(t)$, and fully resistant cells, $R(t)$, and we make the following assumptions (Figure 2.2):

- The initial densities of sensitive and resistant cells are given by S_0 and R_0 , respectively.
- In isolation each population grows logistically with proliferation rates r_S and r_R , and with environmental carrying capacities K_S and K_R , respectively. That




Figure 2.2: The mathematical model. Drug-sensitive (S) and resistant cells (R) proliferate and die at rates, r_S and r_R , and d_S and d_R , respectively. Proliferating sensitive cells die at a rate d_D when exposed to drug, D . Finally, both populations compete for resources, and are assumed to have environmental carrying capacities K_S and K_R , respectively.

is, K_S and K_R represent, respectively, the densities of sensitive and resistant cells which could hypothetically be supported by the resources available in the environment (e.g. the available space). Furthermore, we assume cells die at density-independent death rates $d_S < r_S$ and $d_R < r_R$, respectively.

- The tumour cells compete with each other for resources and space according to a Lotka-Volterra competition model. This means that the presence of the competitor reduces a population's growth rate in a fashion that is linearly proportional to the competitor's population density. Here, we will assume that competition only affects the proliferation rate (for analysis of a model in which competition affects the death rate, see our publication [125]). For simplicity we will further assume that the strengths of inter- and intra-species competition are the same. While arguably these are both oversimplifications, as many tumours have been found to experience a non-linear form of density-dependence (e.g. Gompertzian [39, 148, 41]), the aim of our work was to develop qualitative understanding, for which the simple and well-known nature of the Lotka-Volterra

model was more suitable. For an analysis of more complex models we refer to the recent work by Viossat and Noble [126].

- Only actively dividing cells are killed by the drug. Many chemotherapies induce DNA damage or inhibit the division machinery, which induces apoptosis only in cells that attempt to divide [7]. We will adopt the classical Norton-Simon model [42] which assumes that cell kill increases linearly with the drug dose, $D(t)$, so that at MTD a fraction d_D of dividing cells are killed.
- Progression is driven primarily by pre-existing resistant cells, and resistance acquisition during treatment is negligible. Viossat and Noble [126] have recently shown that high acquisition rates would be required to alter the qualitative dynamics of the model, faster than what would be plausible through genetic mutation. Incorporating results on phenotypic cell plasticity (e.g. [149, 150, 151]) will be the subject of future research (see also Appendix C).

These assumptions translate into the following set of ODEs:

$$\frac{dS}{dt} = r_S \left(1 - \frac{S+R}{K_S}\right) \left(1 - \frac{2d_D}{D_{Max}}D(t)\right) S - d_S S, \quad (2.1)$$

$$\frac{dR}{dt} = r_R \left(1 - \frac{S+R}{K_R}\right) R - d_R R, \quad (2.2)$$

$$N(t) = S(t) + R(t), \quad (2.3)$$

with initial conditions:

$$N(0) = N_0, \quad S(0) = S_0, \quad \text{and} \quad R(0) = R_0,$$

where $N_0 = S_0 + R_0$. Note that the factor of 2 in the drug response accounts for the fact that if a cell dies during mitosis not only the potential daughter but also the mother are lost. We observe also that the total tumour density is fully determined by the density of sensitive and resistant cells, so that it suffices to study Equations (2.1) and (2.2) to obtain $N(t)$. Thus, in the following analysis we will primarily focus on Equations (2.1) and (2.2).

The current drug concentration, $D(t)$, is determined by the treatment schedule and the drug's PK. For simplicity, we will assume that the latter can be neglected and the concentration over time follows exactly the levels foreseen by the drug schedules. We will consider two treatment schedules:

1. CT at the MTD, D_{Max} : $D(t) = D_{Max} \forall t$.

2. AT which withdraws treatment once a 50% decrease from the initial tumour size is achieved as implemented in the AT trial by Zhang et al [97] (see also Figure 1.3c for an illustration):

$$D(t) = \begin{cases} D_{\text{Max}}, & \text{until } N(t) < 50\%N_0 \\ 0, & \text{until } N(t) = N_0 \end{cases} \quad (2.4)$$

We will compare treatment schedules based on their TTP, which is the time between the start of treatment and progression. According to the RECIST criteria [152, 15] tumour progression is defined as a 20% increase in tumour size (SLD) from the nadir (see also Section 1.1.4). However, as tumour size is expected to oscillate during AT due to the treatment changes, this definition is of limited use in evaluating AT. As such, we slightly modify this definition, and define progression as a 20% increase in tumour size from the pre-treatment baseline, akin to the criterion used in most prior AT modelling studies (e.g. [117, 115, 124]).

2.2.1 Non-dimensionalisation

In order to reduce the number of free parameters, we will non-dimensionalise Equations (2.1) - (2.3). Given that we are interested in studying the treatment dynamics if we change the characteristics of the resistant cells relative to the sensitive cells, we base our scales on the sensitive population. In addition, we note that d_D and $D(t)$ are redundant in that the same proportion of dividing cells may be killed by either increasing the drug efficacy, d_D , or the dose, $D(t)$. Thus, we will combine these two parameters into a single measure of treatment level (or drug-induced cell kill), $c(\tau)$. This yields the following set of transformations:

$$\tau = rst, \quad s = \frac{S}{K_S}, \quad r = \frac{R}{K_S}, \quad c = \frac{2d_D D}{D_{\text{Max}}}, \quad \text{and} \quad n = \frac{N}{K_S}.$$

Equations (2.1)-(2.3) then re-scale to:

$$\frac{ds}{d\tau} = (1 - s - r)(1 - c(\tau))s - \hat{d}_S s, \quad (2.5)$$

$$\frac{dr}{d\tau} = \hat{r}_R \left(1 - \frac{1}{\hat{k}}(s + r)\right)r - \hat{d}_R r, \quad (2.6)$$

$$n(\tau) = s(\tau) + r(\tau), \quad (2.7)$$

where:

$$\hat{r}_R = \frac{r_R}{r_S}, \quad \hat{d}_S = \frac{d_S}{r_S}, \quad \hat{d}_R = \frac{d_R}{r_S}, \quad \text{and} \quad \hat{k} = \frac{K_R}{K_S}.$$

Table 2.1: Parameters and their ranges used in the numerical simulations of the model.

Parameter	Description	Value	Reference
r_S	Sensitive cell proliferation rate and model time scale	0.027d^{-1}	Adopted from [97].
$\hat{r}_R := \frac{r_R}{r_S}$	Resistant cell proliferation rate (as a percentage of sensitive cell proliferation rate)	50 - 100%	Lower limit: [124]; Upper limit: no cost
$\hat{d}_S := \frac{d_S}{r_S}$, $\hat{d}_R := \frac{d_R}{r_S}$, $\hat{d}_T := \frac{d_T}{r_S}$	Cell turnover rates (as a percentage of sensitive cell proliferation rate)	0 - 50%	Lower limit: no turnover; Upper limit: see Section 2.3.2.1.
$c := \frac{2d_D D}{D_{\text{Max}}}$	Drug-induced cell kill (see text for further explanation)	1.5	Adopted from [120].
$\hat{k} := \frac{K_R}{K_S}$	Carrying capacity of resistant cells relative to sensitive cells	0.5 - 1	Lower limit: Guess; Upper limit: equal carrying capacities
$n_0 := \frac{N_0}{K_S}$	Initial tumour cell density (as a percentage of sensitive cell carrying capacity)	10 - 75%	Values within this range reported by [153].
$f_R := \frac{r_0}{n_0}$	Initial resistant cell fraction (as a percentage of initial cell density)	0.1 - 10%	Values within this range reported by [147].
$f_S := \frac{s_0}{n_0}$	Initial sensitive cell fraction (as a percentage of initial cell density)	90 - 99.9%	Determined by $1 - f_R$.

The initial conditions transform to:

$$n(0) = \frac{N_0}{K_S} := n_0, \quad s(0) = \frac{S_0}{K_S} := s_0, \quad \text{and} \quad r(0) = \frac{R_0}{K_R} := r_0,$$

and the treatment schedules are adjusted accordingly (not shown). Finally, for notational convenience we will define $f_S := s_0/n_0$ and $f_R := r_0/n_0$ as the initial fractions of sensitive and resistant cells, respectively.

2.2.2 Model parameters

Given the key role which prostate cancer has played in the development of AT we parametrise our model according to this disease. As such, we adopt the proliferation

rate for sensitive cells given in [97] ($r_S = 0.027\text{d}^{-1}$) as our time scale and the treatment level parameter $c = 1.5$ from [120]. For the other parameters we perform parameter sweeps within their biologically realistic ranges. All parameters, their definitions and their ranges, are summarised in Table 2.1.

2.2.3 Numerical methods

We solve the equations using the RK5(4) explicit Runge-Kutta scheme [154] provided in Scipy (specifically, the `SCIPY.INTEGRATE.ODE` class). AT was implemented as a sequence of intervals of length $\Delta\tau$ with CT at level c , where c is adapted between intervals according to the algorithm in Equation 2.4. In preliminary experiments we tested different values of $\Delta\tau$ and found $\Delta\tau = 1/r_S$ to provide a good trade-off between computational speed and numerical accuracy (not shown). All simulations were carried out in Python 3.6, using Scipy 1.1.0 and Numpy 1.15.1. Visualisations were produced with Pandas 0.23.4, Matplotlib 2.2.3, and Seaborn 0.9.0.

2.3 Impact of a cost in the proliferation rate, r_R

In principle, differences between sensitive and resistant cells can occur in any of the three parameters that describe the kinetics of the resistant population (r_R , d_R and K_R). For example, Gallaher et al [124] report a reduced growth rate, r_R , whereas Carrere [28] observes an increased size of the resistant cells which would correspond to a decreased carrying capacity, K_R . In order to develop a baseline understanding of the model under AT, and how this is altered by a resistance cost, we will initially focus on a cost in the proliferation rate, r_R . This is also the most common way in which a cost is modelled (e.g. [115, 124, 146]). In Section 2.3.2 we will subsequently extend our insights to consider other types of cost. As such, throughout this section we will make the following simplifying assumptions:

1. Sensitive and resistant cells share identical turnover rates, $\hat{d}_S = \hat{d}_R := \hat{d}_T$.
2. Sensitive and resistant cells have a common environmental carrying capacity, $K_S = K_R := K$, so that $\hat{k} = 1$.

Thus, in this chapter we will primarily study the following system of equations:

$$\frac{ds}{d\tau} = (1 - s - r)(1 - c(\tau))s - \hat{d}_T s, \quad (2.8)$$

$$\frac{dr}{d\tau} = \hat{r}_R(1 - s - r)r - \hat{d}_T r, \quad (2.9)$$

We make the observation that in the absence of turnover ($\hat{d}_T = 0\%$), we can rewrite Equations (2.8) & (2.9) as:

$$\begin{aligned}\frac{1}{1-s-r} \frac{ds}{d\tau} &= (1-c(\tau))s, \\ \frac{1}{1-s-r} \frac{dr}{d\tau} &= \hat{r}_R r,\end{aligned}$$

so that the term representing competition, $(1-s-r)$, essentially acts as a time scale, which will facilitate analysis. Thus, we divide our discussion in this section into two parts: Firstly, in Section 2.3.1 we will consider the dynamics in the absence of turnover, and subsequently, in Section 2.3.2 we will extend our results to the treatment response in the presence of turnover.

2.3.1 Dynamics in the absence of turnover

The AT algorithm used by Zhang et al [97] in the prostate cancer AT trial alternates between two drug doses: $c = 0$ and $c = c_{\text{Max}}$, where c_{Max} denotes the treatment effect at the maximum tolerable dose, $D = D_{\text{Max}}$. Thus, we can think of the treatment trajectory under AT as the solution of a piece-wise dynamical system, where off-treatment phases are governed by the dynamical system given by Equations (2.8) & (2.9) with $c(\tau) = 0$, whereas on-treatment phases are governed by the system defined by Equations (2.8) & (2.9) with $c(\tau) = c_{\text{Max}}$. In order to build an understanding of the dynamics of each of these systems we will, in the following, perform a phase plane analysis of Equations (2.8) & (2.9) and examine how the phase flow changes as we increase the treatment level, $c(\tau)$.

2.3.1.1 Phase plane analysis of tumour response to continuous treatment at treatment level, $c(\tau) = c^*$

We summarise the steady states of our model in Table 2.2 and show examples of phase planes for different values of c^* in Figure 2.3. As we increase the treatment level, c^* , through 1, we see that the steady state corresponding to tumour elimination transitions from being an unstable node to a saddle. As such it always has at least one unstable manifold, so that tumour elimination is possible if and only if no resistant cells are present, and $c^* > 1$ (Figure 2.3). Moreover, all other steady states lie along the line $r = 1 - s$ where at one extreme, SS2, describes a tumour composed only of sensitive cells, whereas at the other extreme, SS3, corresponds to a tumour made up completely of resistant cells (see also Figure 2.3). All other solutions along $r = 1 - s$ describe tumours composed of a mixture of sensitive and resistant cells. Due to the

Label	Description	State	Eigenvalues
SS1	<i>Tumour Elimination</i>	$s^* = 0,$ $r^* = 0$	$\lambda_1 = \mathbf{1} - \mathbf{c}^*$ $\lambda_2 = \hat{r}_R$
SS2	<i>Sensitive Tumour</i>	$s^* = 1,$ $r^* = 0$	$\lambda_1 = 0$ $\lambda_2 = \mathbf{c}^* - \mathbf{1}$
SS3	<i>Resistant Tumour</i>	$s^* = 0$ $r^* = 1$	$\lambda_1 = 0$ $\lambda_2 = -\hat{r}_R$
SS4	<i>Coexistence</i>	$r^* = 1 - s^*$	No simple form available

Table 2.2: Steady states and linear stability for the model (Equations (2.5) and (2.6)) under CT at treatment level $c(\tau) = c^*$, assuming resistance costs only in the proliferation rate ($\hat{d}_S = \hat{d}_R := \hat{d}_T$ and $\hat{k} = 1$), and no turnover ($\hat{d}_T = 0\%$). **Bolded values** indicate eigenvalues which are negative in at least part of the parameter space.

zero eigenvalues of SS2 and SS3, we can not use linear stability analysis to examine the stability of these steady states. Instead, we study the dynamics by inspecting Equations (2.8) & (2.9) and their associated phase flow (Figure 2.3). From this we can see that between $c^* = 0$ and $c^* = 1$, the tumour will keep growing until it reaches a point on SS4 and saturates (Figure 2.3). The greater the value of c^* , the larger the proportion of resistant cells at equilibrium. When $c^* = 1$, $ds/d\tau$ becomes exactly 0 everywhere (Figure 2.3; see also Equation 2.8). As c^* is increased above 1, the sign of $ds/d\tau$ switches (Figure 2.3). Moreover, the lower section of the line SS4 becomes unstable, and we observe the formation of heteroclinic orbits connecting the unstable portions of SS4 to the stable portions of SS4 (Figure 2.3).

We can formalise these observations by deriving a closed-form expression for the solution in $s - r$ space²

Theorem 1. *Let $s(\tau)$ and $r(\tau)$ be the normalised densities of sensitive and resistant cells, starting from initial densities s_0 and r_0 , respectively. Then in the absence of turnover ($\hat{d}_D = 0$) and under continuous treatment at effective dose $c(\tau) = c^*$ we have:*

$$\frac{s}{r^\alpha} = \frac{s_0}{r_0^\alpha} = \text{constant}, \quad (2.10)$$

for all time, τ , where:

$$\alpha = \frac{1 - c^*}{\hat{r}_R}.$$

²We acknowledge the help of Professor Viossat in deriving this result.

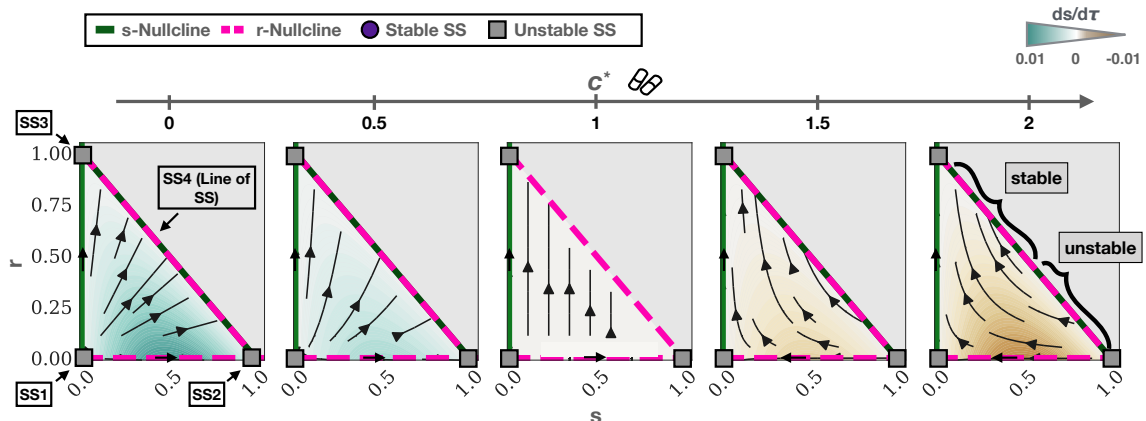


Figure 2.3: Phase plane analysis of the model (Equations (2.5) and (2.6)) under CT with increasing levels of treatment level, c^* (other parameters as in Table 2.1). We assume no turnover ($\hat{d}_T = 0\%$) and no cost of resistance ($\hat{r}_R = 100\%$). SS stands for steady state. Grey zones correspond to tumours which are greater than their carrying capacity and, thus, not biologically realistic. There is a continuum of steady states (SS4) given by $s + r = 1$. For $c^* < 1$ the tumour continues to grow until it reaches carrying capacity on SS4. For $c^* > 1$, sections of SS4 become unstable and we see the formation of heteroclinic orbits between unstable and the stable section of SS4 (see text for details). These orbits represent tumours which are, initially, largely sensitive so that they will experience tumour shrinkage when treatment begins, but will subsequently recur due to the outgrowth of resistance.

Proof. Comparing the rate of growth of sensitive cells, s (Equation (2.8)), to that of resistant cells, r (Equation (2.9)), we find that:

$$\begin{aligned} \frac{ds}{d\tau} &= \frac{(1-s-r)(1-c^*)s}{\hat{r}_R(1-s-r)r} \\ &= \frac{(1-c^*)s}{\hat{r}_R r}, \end{aligned}$$

which we can integrate with respect to τ . This gives:

$$\ln(s) = \alpha \ln(r) + C,$$

where $\alpha = \frac{1-c^*}{\hat{r}_R}$ and C is the integration constant. Applying the initial conditions completes the proof. \square

Based on this we can now assert the following observations: Assuming the tumour is initially composed of a mixture of sensitive and resistant cells ($s_0, r_0 > 0$), we have:

Observation 1. The final tumour composition is given by (s^*, r^*) , where s^* is the solution of the equation: $s + \left(\frac{s}{C_0}\right)^{1/\alpha} = 1$, where $C_0 = s_0/r_0^\alpha$; and $r^* = 1 - s^*$.

Proof. Equation (2.6) has steady states $r^* = 0$ or $s^* + r^* = 1$. Given our assumption that $r_0 > 0$, we conclude we must have $s^* + r^* = 1$. In combination with Equation (2.10) we can solve for s^* which yields the above-stated result. \square

Observation 2. Neither the sensitive nor the resistant population will ever go extinct (except asymptotically).

Proof. For $c^* \leq 1$, inspection of Equations (2.8) & (2.9) shows that everywhere within the feasible region ($s + r \leq 1$) we have $ds/d\tau > 0$ and $dr/d\tau > 0$. Thus, neither population can decrease in size.

Similarly, if $c^* > 1$, then we still have $dr/d\tau > 0$ everywhere, so that the resistant population will not shrink. As for the sensitive population, we observe from Equation (2.6) that $r(\tau) \leq 1, \forall \tau \geq 0$. Using Equation (2.10) and noting that $c^* > 1$ implies that $\alpha < 0$, we conclude that $s = \frac{s_0}{r_0^\alpha} r^\alpha \geq \frac{s_0}{r_0^\alpha} > 0$. \square

Biologically, the transition through $c^* = 1$ represents a transition from when the drug only reduces the growth rate of the tumour to when it actually decreases the size of the sensitive population. Thus, unless $c^* > 1$ we can not actually achieve a decline in the tumour size following treatment. The heteroclinic orbits represent tumours which initially shrink in size but subsequently recur due to the emergence of resistance. Moreover, as will become clear in the remainder of this chapter, Observation 2 has important implications for the tumour's treatment response. Firstly, it implies that in the absence of turnover we can never eliminate all sensitive cells, because the growth of the resistant cells saturates the environment which blocks the division of sensitive cells and, thus, protects them from drug. Secondly, and more subtly, Observation 2 implies that in the absence of turnover, the resistant population will not diminish in size despite competition by the sensitive cells and a resistance cost. As we will see, this has important ramifications for AT.

2.3.1.2 AT does not *per se* require a resistance cost

Previous work has shown that both the abundance of resistance (f_R) and the proximity to carrying capacity (n_0), which reflects the intensity of resource competition, determine the benefit of AT over CT [59, 115, 124, 122]. To test whether this holds true also in the current model, we simulate treatment under each approach for 9 different tumours representing three different levels of *a priori* resistance and three different levels of proximity to carrying capacity ($(\hat{r}_R, \hat{d}_T) = (100\%, 0\%)$). We find that the benefit of AT is greatest when there is little resistance ($f_R = 0.1\%$) and the

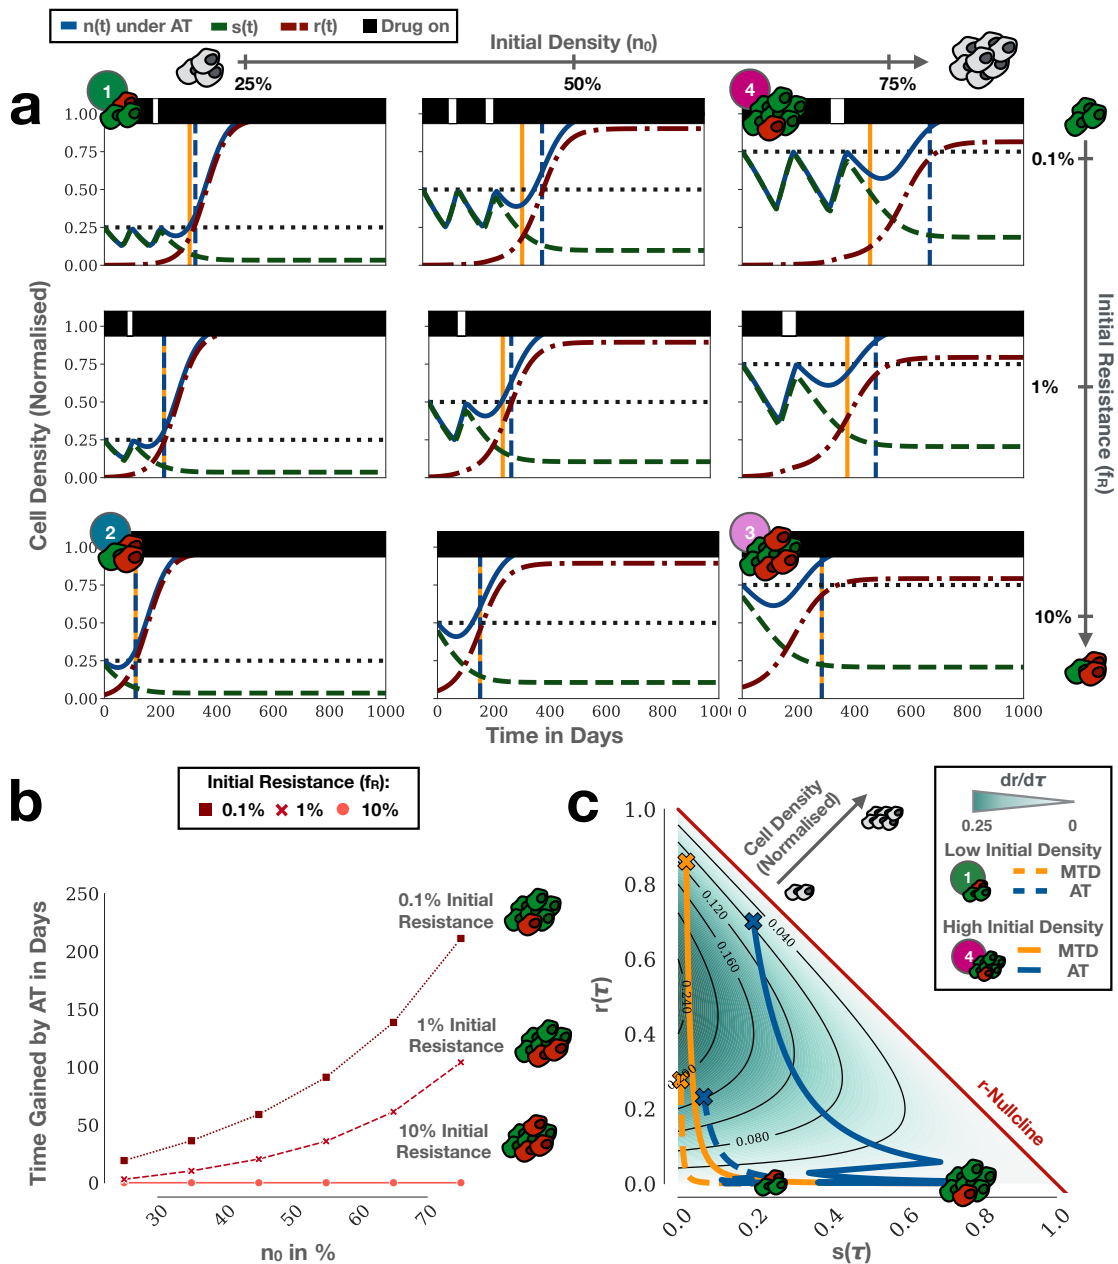


Figure 2.4: AT can extend TTP compared to CT even in the absence of a cost of resistance. We are assuming no turnover here ($\hat{d}_T = 0\%$). (a) Simulations of AT for a cohort of tumours with different initial compositions. Vertical lines mark the time of progression of CT (yellow, solid) and AT (blue, dashed), respectively. Black bars indicate that the treatment is on. (b) Gain in TTP by AT compared to CT as a function of initial proximity to carrying capacity and abundance of resistance. (c) $dr/d\tau$ as a function of s and r , together with treatment trajectories for Tumours 1 ($(n_0, f_R) = (25\%, 0.1\%)$) & 4 ($(n_0, f_R) = (75\%, 0.1\%)$) from (a). Crosses indicate progression. This illustrates how AT extends TTP by reducing $dr/d\tau$ via competition, and demonstrates that for certain tumours AT can extend TTP even if no cost of resistance is present.

tumour is close to K ($n_0 = 75\%$; $TTP_{AT} - TTP_{CT} = 211d$; Figure 2.4a). To further characterise this relationship, we extend our analysis to 18 tumours and plot the gain in TTP under AT as a function of the tumour composition parameters in Figure 2.4b. We find that if the abundance of resistance is high, then AT provides no benefit over CT, independent of how close the lesion is to its carrying capacity (Figure 2.4b). In contrast, if resistance is low, then AT prolongs TTP in a non-linear manner dependent on both the abundance of resistance and the ratio n_0 . We conclude that AT can improve TTP, even in the absence of a cost, and that this gain is greatest the closer the tumour is to carrying capacity and the lower the initial abundance of resistance.

2.3.1.3 Treatment vacations provide a benefit only if intra-tumoral competition is strong

Moreover, we find that each of the two characteristics of the tumour's initial composition has a distinct impact on the treatment dynamics. As we increase the initial abundance of resistance from 0.1% to 10%, we decrease the number of completed AT cycles (Figure 2.4a). In the most extreme case, at 10% initial resistance, treatment can not decrease the tumour burden sufficiently to trigger any treatment withdrawal in the three tumours. In contrast, increasing the level of growth saturation in a tumour (proximity to carrying capacity) does not alter the AT cycle number, but does increase the benefit delivered by each cycle. For example, even though all tumours with 1% initial resistance complete one AT cycle, a meaningful benefit in TTP is only achieved in the case when the tumour is 75% saturated (Figure 2.4a). The reason for this is that the competition exerted by sensitive cells only has significant impact on the growth of the resistant cells if the tumour is close to K . This implies that because of the discrete thresholds built into the AT algorithm, we require not just competition, but also that the tumour burden can be decreased sufficiently to trigger treatment withdrawal before we can hope to see a benefit of AT.

2.3.1.4 AT extends TTP by minimising the resistant population growth rate

How can we explain the benefit of AT over CT in the absence of a resistance cost? To address this question, we will adopt an argument which was recently proposed by Hansen et al [1] in the study of antibiotic resistance in bacteria³. The authors analysed a 2-population ODE model which was almost identical to ours, except that they additionally included mutation of sensitive to resistant cells. They then asked

³We acknowledge the assistance of Professor Viossat in deriving this result.

when it would be beneficial to maintain sensitive bacteria to suppress resistant bacteria, and when it might be better to minimise the bacterial burden, even if this causes competitive release (although unlike us they did not study any explicit dosing algorithm to achieve control). To address this question, they assumed that the drug is sufficiently effective, so that drug-sensitive cells can be easily removed whenever there are too many ($c^* > 1$). This then implies that progression is driven only by the expansion of the resistant population, and it suffices to study the governing equation of $r(\tau)$ (Equation (2.6)) to determine TTP. The more a treatment strategy can inhibit the resistant population growth rate, $dr/d\tau$, whilst also maintaining control of the tumour size overall, the longer TTP. We can re-write Equation (2.6) as follows:

$$\frac{1}{\hat{r}_R} \frac{dr}{d\tau}(s, r, \tau) = \underbrace{\left(1 - \frac{\hat{d}_T}{\hat{r}_R}\right) r}_{\text{Growth}} - \underbrace{\frac{r}{\hat{k}} r}_{\text{Intra-Specific Competition}} - \underbrace{\frac{s}{\hat{k}} r}_{\text{Inter-Specific Competition}}. \quad (2.11)$$

While we cannot alter the growth term with treatment, we can change intra- and inter-specific competition. In particular, we can see from Equation (2.11) that the larger we maintain s , the more we can slow down $dr/d\tau$. Moreover, this holds true even in the absence of a cost.

More formally, we can say that if $s_{CT}(t)$ and $s_{AT}(t)$ denote the density of sensitive cells after t subsequent days of treatment under CT and AT, respectively, then we have that $s_{CT}(t) \leq s_{AT}(t)$ because CT does not provide sensitive cells with an opportunity to recover. Assuming that the tumour is still far from progression ($r \ll s$) so that the resistant cells primarily compete with sensitive cells, the resistant growth rate is greater for CT ($\left.\frac{dr}{d\tau}\right|_{CT}$) than for AT ($\left.\frac{dr}{d\tau}\right|_{AT}$):

$$\begin{aligned} \left.\frac{dr}{d\tau}\right|_{CT} &\approx \left(\hat{r}_R - \hat{d}_T\right) r - \frac{\hat{r}_R s_{CT}}{\hat{k}} r \\ &\geq \left(\hat{r}_R - \hat{d}_T\right) r - \frac{\hat{r}_R s_{AT}}{\hat{k}} r \approx \left.\frac{dr}{d\tau}\right|_{AT}, \end{aligned}$$

This implies that while a cost extends overall TTP, AT will be superior even in the absence of a resistance cost. Moreover, this benefit increases with increased initial tumour density as s_{AT} can be maintained large for longer. In addition, this implies that the most effective AT would maintain the tumour at its original size for as long as possible in order to maximise the effect of competition.

To illustrate this, we plot $dr/d\tau$ from Equation (2.6) as a function of the tumour composition $(s(\tau), r(\tau))$ (Figure 2.4c). In this representation a tumour lesion is thought of as a point in a two-dimensional space, where its x-position represents the

current density of sensitive cells $s(\tau)$, and its y-position the current density of resistant cells, $r(\tau)$. Each point is coloured according to the resistant population growth rate. This representation clearly illustrates how high tumour densities (see inset grey arrow) are generally associated with lower resistant growth rates. Furthermore, we show trajectories of CT (yellow) and AT (blue) for two tumours corresponding to Tumours 1 & 4 from Figure 2.4a. As can be seen, CT trajectories tend to traverse regions of high resistant growth (dark green shading). In contrast, these same regions are avoided under AT regimens, especially for Tumour 4 (Figure 2.4c).

To sum up, this implies that if we have perfect control over the sensitive cells, if there is no mutation from sensitive to resistant cells, and if CT gains no benefit over AT from intra-specific competition, then AT will always be non-inferior to CT in delaying progression. This is a profound insight. It ascertains that AT can work under a broad range of conditions. For example, we can conclude from this that if we changed the logistic growth law in our model to a Gompertzian one, then we may see different quantitative predictions, but the qualitative conclusions would remain the same (see also Viossat and Noble [126]). At the same time, it also identifies three assumptions which, if not met, may result in a situation in which AT can perform more poorly than CT (perfect control, no mutation, intra-specific competition). Their study will be an important aspect of the subsequent chapters of this thesis.

2.3.1.5 In the absence of turnover, cost of resistance increases time gained by AT only for low levels of resistance

Next, we examine whether a cost of resistance enhances the benefit of AT. We still assume that no turnover is present but now assume a 30% decrease in proliferation rate of the resistant population ($\hat{r}_R = 70\%$) and simulate CT and AT for the four corner cases from Figure 2.4a. We choose a value of 30% as this lies at about the midpoint of the assumed range of plausible cost values. As expected, TTP_{AT} increases in all four cases (Figure 2.5a; compare vertical blue dashed line in top and bottom row). Moreover, the effect is much more pronounced when the tumour is close to carrying capacity and the initial fraction of resistance is small (Tumour 4; Figure 2.5a). However, TTP_{CT} increases as well, so that significant gains compared to CT are only seen for Tumour 4 (difference in TTP extended from 211d to 330d; Figure 2.5a). To further quantify this observation we compute the difference between TTP_{AT} and TTP_{CT} for each of the four tumours in Figure 2.5a for costs ranging from 0% to 50%. This shows that in the absence of turnover, a resistance cost significantly

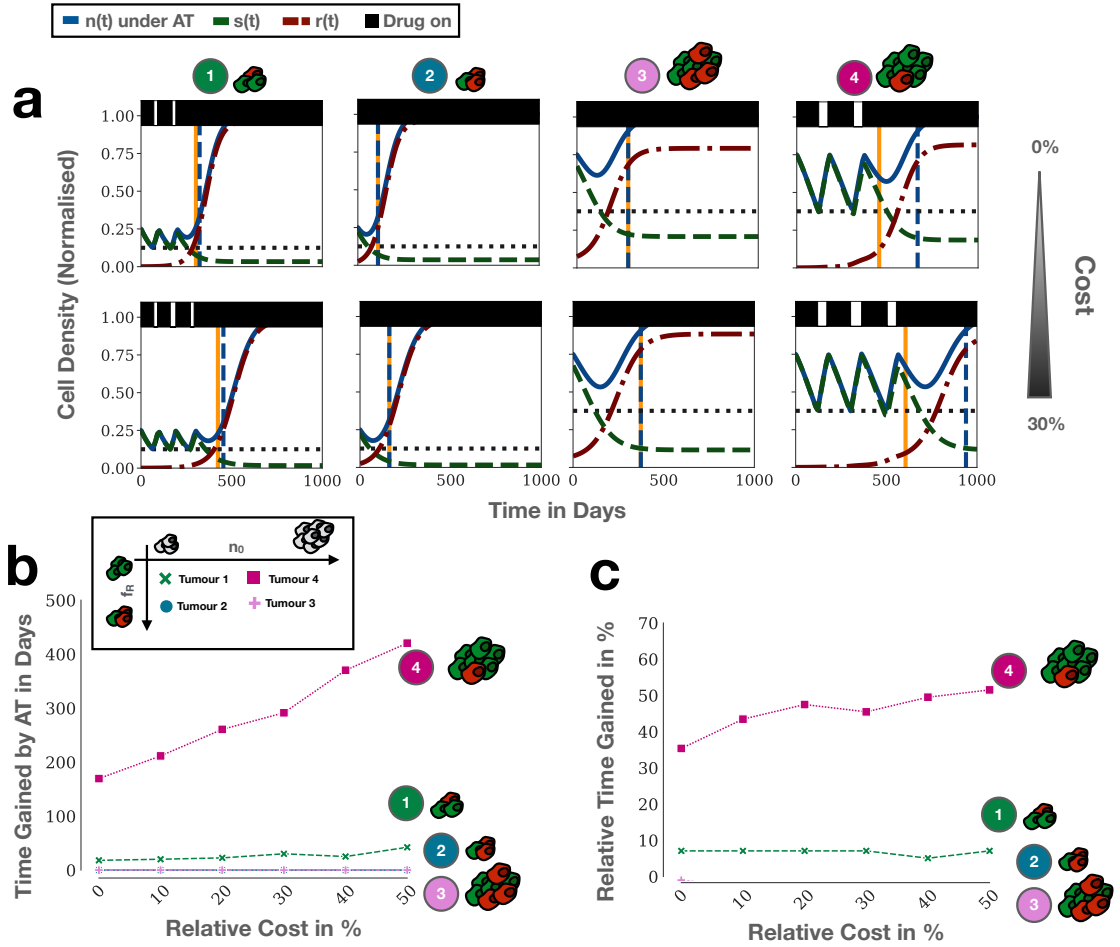


Figure 2.5: In the absence of turnover ($\hat{d}_T = 0\%$) a cost of resistance enhances AT only when resistance is rare and tumours are close to carrying capacity. **(a)** Simulations of Tumours 1-4 from Figure 2.4a with and without a 30% cost of resistance ($\hat{r}_R = 70\%$; Tumours 1 & 4 given in Figure 2.4; Tumour 2: $(n_0, f_R) = (25\%, 10\%)$; Tumour 3: $(n_0, f_R) = (25\%, 10\%)$). **(b)** & **(c)** Absolute ($TTP_{AT} - TTP_{CT}$) and relative $((TTP_{AT} - TTP_{CT})/TTP_{CT})$ gain in TTP by AT compared to CT for different magnitudes of resistant costs for Tumours 1-4. Only Tumour 4 sees a significant gain.

increases the benefit of AT only for Tumour 4, that is, only if resistance is rare and the tumour is close to carrying capacity (Figure 2.5b & c).

2.3.2 Dynamics in the presence of turnover

So far we have assumed that death of tumour cells is negligible. However, whilst tumour cells are known for their excessive proliferation rates, they are also subject to significant resource starvation and immune predation. In addition, homeostatic human tissues continuously renew, so that there still is a degree of natural turnover

also in tumours [36]. Although it is less often discussed, and it is harder to quantify, a number of authors have therefore argued that cell turnover may play an important role in tumours [34, 36, 155]. Our analysis will further support this claim.

2.3.2.1 What is the rate of cell death in human tumours?

Pathological Type	T_V (in days) (95% CI)	T_{Pot} (in days)	Estimated Growth Fraction (in %)	Cell Loss Factor, ϕ (in %)	Estimate for \dot{d}_T (=GF $\times\phi$)
Embryonal tumours	27 (22-33)	1.66	90	94	84.4%
Hemato-sarcoma	29 (23-37)	1.7	90	94	84.4%
Mesenchymal Sarcomas	41 (35-50)	13.2	11	68	7.48%
Squamous cell carcinoma	58 (48-70)	6.0	25	90	22.5%
Adeno-carcinoma	83 (72-96)	23.8	6	71	4.3%

Table 2.3: Examples of growth rates and cell loss factors measured in human tumours (taken from the meta-analysis by Malaise and colleagues in [36]). GF: Growth fraction.

A number of studies have set out to address this question using a method developed by Steel in 1960s [34, 156, 35]. The idea is to estimate the theoretically possible (“potential”) doubling time, T_{Pot} , achieved in the absence of any cell death, and the actually observed volumetric doubling time of a tumour, T_V . By comparing these values one can compute the “cell loss factor”, ϕ , defined as $\phi = 1 - T_{Pot}/T_V$, which represents the proportion of cells being born to cells dying or leaving the tumour. The greater the value of ϕ , the greater the loss of cells due to cell death, predation, or emigration [34, 156, 35].

T_V is estimated from x-ray, CT or MRI images. T_{Pot} is obtained by measuring the fraction of cells in S-phase in the tumour via halogenated pyrimidine, bromo- or iodo-deoxyuridine labelling. Cells carrying out DNA synthesis incorporate these agents into their DNA, and can subsequently be identified using histological techniques or flow cytometry. The proportion of labelled cells is termed the labelling index (LI). To obtain T_{Pot} , the LI is adjusted to account for proliferating cells which are not currently in S-phase and is multiplied by an estimate of the cell cycle time, based on the time required for DNA synthesis (T_S). This gives $T_{Pot} = \lambda T_S / LI$, where λ

denotes the proportion of the cell cycle spent in the S-phase. For a detailed review of the process of estimating T_{Pot} see [157].

In Table 2.3 we show estimates for the cell loss factor for a range of human tumours. These are all well in excess of 50% - an observation which has been corroborated by a large number of studies over the years (see [156, 36, 158] for extensive reviews). This implies that at any point in time the number of cells produced in the tumour closely matches the number of cells lost. To translate estimates for ϕ to values of our model parameter \hat{d}_T we need to note that ϕ compares the cell turnover rate to the effective proliferation rate of the population (which takes into account the actual fraction of proliferating cells), not the proliferation rate. To illustrate this, consider a logistic model of a growing tumour. Let $N(t)$ denote the total number of cells and assume that cells divide and die at rates r_T and d_T , respectively, where the proliferation rate decreases linearly as the tumour approaches its carrying capacity, K . Then: $dN/dt = r_T(1 - N/K)N - d_TN$. If we were to measure this tumour at time t^* we would observe a growth fraction of $(1 - N(t^*)/K)$ and estimate (making the assumption of the Steel method that this growth fraction remains constant), that $T_{\text{Pot}} = \ln(2)/r_T(1 - N(t^*)/K)$. The actual volumetric doubling time (again assuming the growth fraction remains constant) would be $T_V = \ln(2)/(r_T(1 - N(t^*)/K) - d_T)$. Thus, $\phi = d_T/(r_T(1 - N(t^*)/K))$. To obtain $\hat{d}_T := d_T/r_T$, we therefore need to scale the observed values of ϕ by the observed growth fraction, GF: $\hat{d}_T = \text{GF} \times \phi$. This yields values between 4% and 81%, with rates for solid tumours (the settings in which AT has been mostly studied) clustered below 25% (Table 2.3). Given the variability between cancer types, and the back-of-the-envelope nature of our calculations, we choose to consider values of up to 50%. Future research should investigate turnover rates in tumours in more detail to obtain more accurate estimates.

2.3.2.2 Phase plane analysis

In order to build an understanding of how turnover changes the dynamics of our model, we will again begin with a phase plane analysis under constant treatment level, c^* . We summarise the steady states and their linear stability in Table 2.4. Initially, let us assume that there is no cost of resistance ($\hat{r}_R = 100\%$). In Figure 2.6 we show a bifurcation analysis of the non-trivial steady states for s and r . If $c^* = 0$, then we again observe the line of steady states formed by the overlap of the non-trivial s - and r -nullclines (Table 2.4 and Figure 2.6a & b). However, this time the tumour saturates below the carrying capacity ($s^*, r^* < 1$) and, if the treatment

Label	Description	State	Eigenvalues
SS1	<i>Tumour Elimination</i>	$s^* = 0,$ $r^* = 0$	$\lambda_1 = \mathbf{1} - \mathbf{c}^* - \hat{\mathbf{d}}_T$ $\lambda_2 = \hat{r}_R - \hat{d}_T$
SS2	<i>Sensitive Tumour</i>	$s^* = 1 - \frac{\hat{d}_T}{1 - c^*},$ $r^* = 0$	$\lambda_1 = \hat{\mathbf{d}}_T \left(\frac{\hat{r}_R}{1 - c^*} - \mathbf{1} \right)$ $\lambda_2 = (c^* - \mathbf{1}) \left(\mathbf{1} - \frac{\hat{d}_T}{1 - c^*} \right)$
SS3	<i>Resistant Tumour</i>	$s^* = 0$ $r^* = 1 - \frac{\hat{d}_T}{\hat{r}_R}$	$\lambda_1 = \hat{\mathbf{d}}_T \left(\frac{1 - c^*}{\hat{r}_R} - \mathbf{1} \right)$ $\lambda_2 = \hat{\mathbf{d}}_T - \hat{r}_R$
SS4	<i>Coexistence</i>	$r^* = 1 - \frac{2\hat{d}_T}{1 + \hat{r}_R - c^*} - s^*$	No simple form available

Table 2.4: Steady states and linear stability for the model (Equations (2.5) and (2.6)), assuming cellular turnover and resistance costs only in the proliferation rate ($\hat{d}_S = \hat{d}_R := \hat{d}_T$ and $\hat{k} = 1$). For discussions of the singularities at $c^* = 1$ and $c^* = 1 + \hat{r}_R$ see the main text. **Bolded values** mark eigenvalues which are negative in at least part of parameter space.

level, c^* , is increased to some value $0 < c^* < 1$, then the s -nullcline, given by:

$$s(\tau) = 1 - \frac{\hat{d}_T}{1 - c^*} - r(\tau),$$

is translated downwards (Figure 2.6a & b-c). The result is that SS4 ceases to exist for $c^* > 0$. Instead, SS3 becomes stable ($\lambda_1, \lambda_2 < 0$) whereas SS2 becomes unstable ($\lambda_1 > 0$; Figure 2.6c - f). Thus, in the presence of turnover the tumour evolves to become fully resistant even if treated at a very low dose.

When $c^* = 1 - \hat{d}_T$, the sensitive steady state (SS2) merges with the trivial steady state (SS1), so that the resistant steady state (SS3) remains the only attracting and biologically feasible steady state (Table 2.4 and Figure 2.6d - f). If $c^* > 1$, then SS2 takes on a value above 1 which converges to 1 from above as $c^* \rightarrow \infty$ (see Table 2.4 and Figure 2.6f). Whilst positive, this state is not biologically realistic as it implies that the population can grow beyond the environmental carrying capacity. Indeed, for any initial tumour size less than the carrying capacity, it is trivial to show mathematically that it stays bounded above by the carrying capacity.

Next, we consider the impact of a cost of resistance in this setting ($\hat{d}_T < \hat{r}_R < 100\%$). From Table 2.4 we see that for $c^* < 1 - \hat{r}_R$, SS2 is stable whereas SS3 is unstable (note that in this case $(1 - c^*)/\hat{r}_R > 1$). Thus, in the presence of a cost, low dose treatment can shrink the tumour to a fully sensitive, stable steady state (Figure 2.7a). However, if the treatment level is intensified beyond $1 - \hat{r}_R$, SS2 again becomes unstable, SS3 becomes stable (Figure 2.7a) and the tumour will turn fully resistant over time (Figure 2.7d - f).

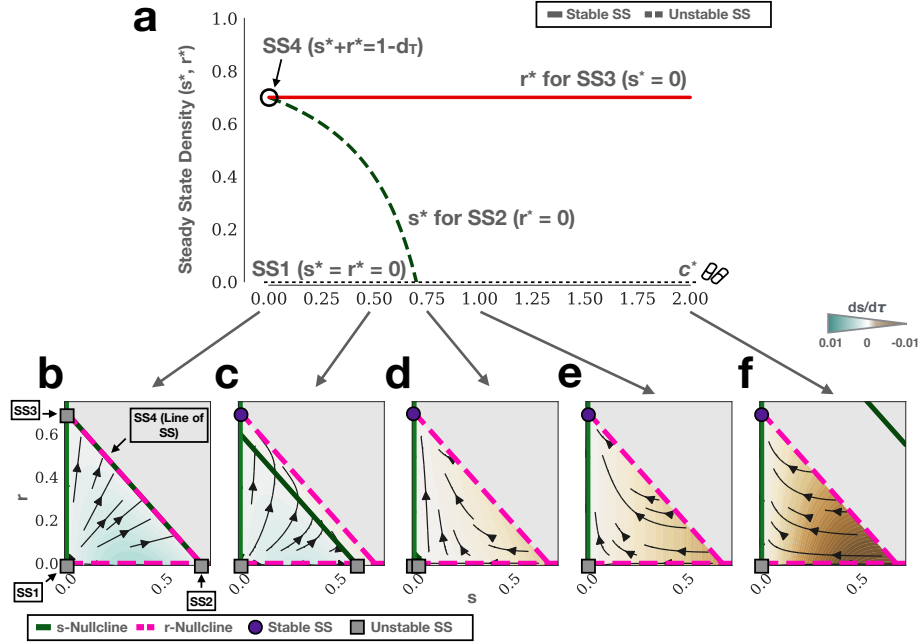


Figure 2.6: Bifurcation analysis of the model (Equations (2.5) and (2.6)) for increasing levels of treatment, c^* , assuming turnover, but no cost of resistance ($\hat{r}_R = 100\%$). (a) Steady states, and their stability, as a function of c^* (shown is the density of the non-zero population in each case). Colour indicates whether the line shows the steady state density of sensitive (green) or resistant cells (red), respectively. (b)-(f) Phase planes corresponding to specified values of c^* ((b) $c^* = 0$, (c) $c^* = 0.5$, (d) $c^* = 0.7$, (e) $c^* = 1$, (f) $c^* = 2$). In the absence of treatment both non-zero nullclines collapse to give a continuum of steady states (SS4 in (a) & (b)). When $c^* > 0$ the sensitive steady state (SS2) becomes unstable and eventually biologically infeasible ((a) & (c)-(f)), so that the tumour evolves to be fully resistant. The tumour-free state (SS1) is always unstable. Parameters: $\hat{r}_R = 100\%$, $d_T = 30\%$, $\hat{k} = 1$.

2.3.2.3 Turnover mediates the impact of a cost of resistance

The steady state analysis showed that turnover causes growth to saturate before the tumour reaches the environmental carrying capacity, so that there is continuous cycling of cells in the tumour. This has two important implications for treatment: Firstly, it increases the drug kill, and thereby the selection for resistance under treatment. Figure 2.6 illustrates how the sensitive cells are slowly killed off until the tumour is completely resistant, even under a very low dose of treatment.

Secondly, turnover also increases selection against resistance in the absence of treatment. To illustrate this, we simulate a competition experiment in which we seed sensitive and resistant cells at equal proportions ($f_R = 0.5$) and monitor their abundance over time in the absence of drug (Figure 2.8). Without turnover, both

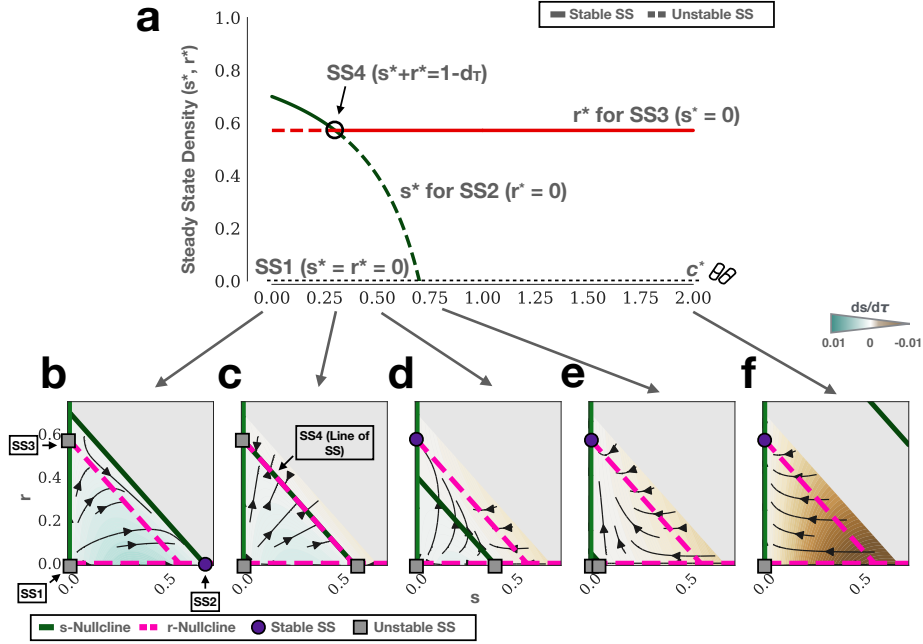


Figure 2.7: Bifurcation analysis of the model (Equations (2.5) and (2.6)) for increasing levels of treatment level, c^* , assuming turnover, and a cost of resistance ($\hat{r}_R = 70\%$). (a) Steady states, and their stability as a function of c^* (shown is the density of the non-zero population in each case). Colour denotes the dominant population. (b)-(f) Phase planes corresponding to specified values of c^* ((b) $c^* = 0$, (c) $c^* = 0.3$, (d) $c^* = 0.5$, (e) $c^* = 0.7$, (f) $c^* = 2$). For $c^* < 1 - \hat{r}_R$, the sensitive steady state (SS2) is the attracting steady state (a). The resistant cells are outcompeted by the sensitive cells and the tumour becomes fully sensitive over time (b). When $c^* = 1 - \hat{r}_R$, the non-trivial s - and r -nullclines overlap to give a continuum of steady states (c). For $c^* > 1 - \hat{r}_R$, the sensitive steady state (SS2) becomes unstable and eventually biologically infeasible ((d)-(f)), so that the tumour evolves to be fully resistant. Parameters: $\hat{r}_R = 70\%$, $d_T = 30\%$, $\hat{k} = 1$.

populations grow until they make up a fixed proportion of the tumour given in the limit by $(s, r) = (s^*, 1 - s^*)$, where s^* satisfies the conservation equation (Equation (2.10); Figure 2.8a). Importantly, while a cost of resistance reduces the number of resistant cells, these cells never go extinct (Figure 2.8a). In contrast, if we introduce turnover ($\hat{d}_T > 0\%$), resistant cells paying a cost go extinct over time (Figure 2.8b & c; see also Figure 2.7). Turnover is, thus, a key modulator of the impact of a resistance cost.

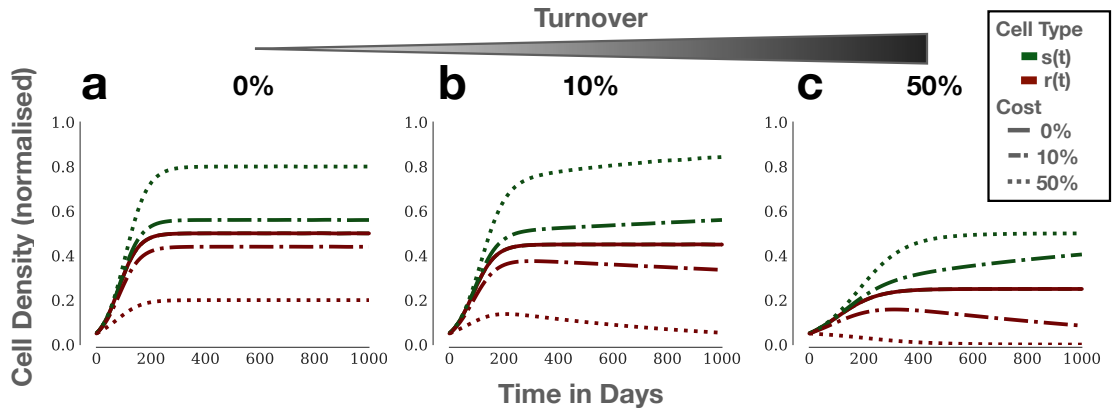


Figure 2.8: Turnover modulates the impact of a resistance cost. Model simulations of a competition experiment between sensitive and resistant cells seeded at equal proportions ($f_R = 0.5$), and grown in the absence of drug ($c^* = 0$). (a) In the absence of turnover ($\hat{d}_T = 0\%$), even a significant cost does not result in the elimination of the resistant cells. (b) & (c) In contrast, if we increase the natural death rate, then resistant cells paying a cost are more and more quickly outcompeted by the sensitive cells ((b) $\hat{d}_T = 10\%$; (c) $\hat{d}_T = 50\%$).

2.3.2.4 Tumours with high turnover may show better response to AT than tumours with low turnover

How do these insights translate to AT? To address this question, we compute the time gained by AT as a function of turnover and cost for Tumour 1 (Figure 2.9). This shows that when turnover is low, even a large cost of resistance does not guarantee that AT will significantly extend TTP (compare Cases i and ii in Figure 2.9a). In contrast, if turnover is high, then AT generates significant benefits even if the cost is small, or completely absent (Case i vs iii in Figure 2.9a). We conclude that high turnover may help to control a tumour with AT.

Figure 2.9a also shows an area of parameter space where the tumour becomes indefinitely controllable in the model, so that CT and AT can maintain it from progressing forever. This is because in the presence of turnover, the final tumour size under CT is limited by $n(t) = 1 - \frac{\hat{d}_T}{\hat{r}_R} := k_{\text{EFF}}$, where we are assuming that the treatment effect is sufficiently strong to outweigh any resistance cost ($c^* > 1 - \hat{r}_R$; Table 2.4). Thus, if $k_{\text{EFF}} < 1.2n_0$, then a tumour can not progress under CT nor AT according to our definition of progression as a 20% increase from pre-treatment size. To illustrate this we show an example in Figures 2.10a & b. This kind of indefinite tumour control is likely unrealistic, as we would expect the tumour's environmental carrying capacity not to remain constant over time, but to slowly increase due to,

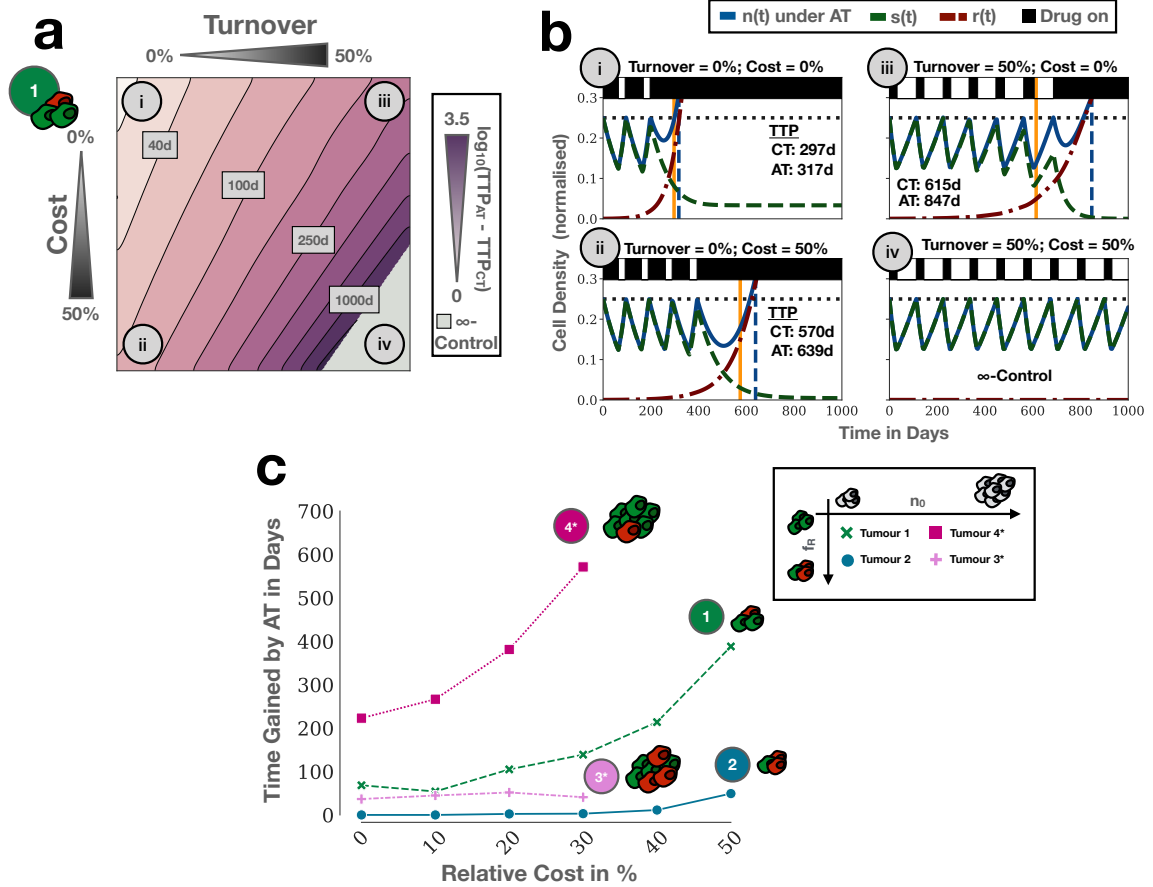


Figure 2.9: Turnover increases the time gained by AT. **(a)** TTP gained by AT as a function of cost and turnover for Tumour 1. Turnover increases the benefit of AT and amplifies the effect of a cost. **(b)** Simulations of the treatment dynamics for four combinations of cost and turnover corresponding to the four case studies highlighted in **(a)**. **(c)** Turnover increases TTP gained by AT ($TTP_{AT} - TTP_{CT}$) both in the presence and absence of a resistance cost for a range of different tumour compositions (Tumours 1 and 2 as in Figure 2.4a; Tumour 3*: $(n_0, f_R) = (50\%, 10\%)$; Tumour 4*: $(n_0, f_R) = (50\%, 0.1\%)$). Above a cost of 30% Tumours 3* and 4* become indefinitely controllable and so no TTP can be obtained.

for example, angiogenesis. Thus, in reality the tumours in this domain would still eventually progress. Nevertheless, this shows mathematically what we would have to do to achieve indefinite tumour control: We have to change the system dynamics so that the tumour is above the r -nullcline. This may be achieved either by increasing the density of sensitive cells, or by changing the position of the r -nullcline.

Finally, we examine how these results generalise as we change the initial tumour density and resistance fraction. We repeat the cost-gain relationship analysis from Figure 2.9b with a turnover rate of $\hat{d}_T = 25\%$. As Tumours 3 & 4 are indefinitely

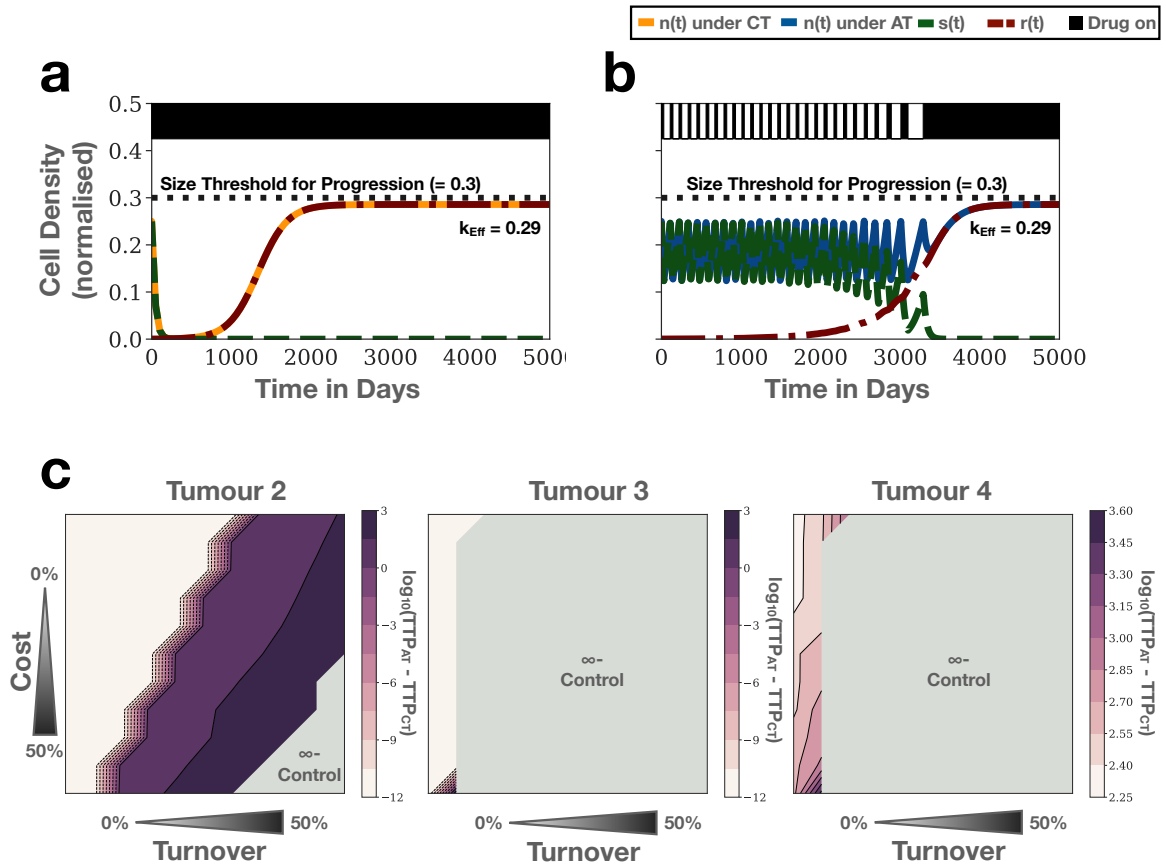


Figure 2.10: As a consequence of the assumption that carrying capacity is constant, the model allows for indefinite tumour control when $(1 - \hat{d}_R / \hat{r}_R) < 1.2n_0$. (a) Example simulation illustrating indefinite control under CT for Tumour 1. While the tumour becomes completely resistant it can not grow above the 20% size increase required for progression, as the effective carrying capacity of the resistant population is too low. Parameters: $n_0 = 25\%$, $f_R = 0.1\%$, $\hat{r}_R = 70\%$, $\hat{d}_R = 50\%$. (b) Same tumour as in (a), but treated with AT. In this case also the tumour can never meet the threshold for progression. (c) The gain by AT as a function of turnover and cost for Tumours 2-4. As Tumour 2 is far from carrying capacity ($n_0 = 25\%$) and has a high resistance portion, strong turnover and/or big resistance costs are required to see significant benefits of AT. Conversely, since they are close to K ($n_0 = 75\%$) Tumours 3 & 4 are indefinitely controllable for a large parameter regime.

controllable in this parameter regime (see Figure 2.10c), we replace them by tumours with a slightly lower initial density ($n_0 = 50\%$), denoted by Tumours 3* & 4*, respectively. We find that all cases except for Tumour 2, which is far from carrying capacity and has a high resistance fraction, benefit from AT even in the absence of a cost (Figure 2.9d). In particular, note that Tumour 3* does not benefit from AT at all in the absence of turnover (Figure 2.4a), but gains 37 days when $\hat{d}_T = 25\%$.

Moreover, in all cases a cost of resistance increases the gain of AT. This shows that turnover relaxes the need for the tumour to be close to environmental carrying capacity for AT to be successful, and that it amplifies the gains obtained even from small resistance costs. However, it also shows that the initial resistance fraction remains an important determinant of TTP_{AT} .

2.3.2.5 Resource availability, turnover and cost combine to define the competition in the tumour

Why does turnover increase the benefit of AT? In Section 2.3.1.4 we showed that AT extends TTP by slowing the growth of the resistant population through competition. As can be seen in Figure 2.4c, this competition is modulated by the proximity of the tumour to environmental carrying capacity. The closer the tumour is to K , the smaller $dr/d\tau$. To investigate the impact of turnover on competition, we rewrite Equation (2.2) as:

$$\frac{dR}{dt} = (r_R - d_R) \left(1 - \underbrace{\frac{1}{(1 - d_R/r_R)} \frac{S + R}{K}}_{C(S, R, K, d_R, r_R)} \right) R,$$

where we have chosen to consider the dimensional form of the equation to explicitly discriminate the role of each parameter.

The function $C(S, R, K, d_T, r_R)$ describes the combined growth inhibition of the resistant population through intra- and inter-specific competition. A value of $C = 0$ corresponds to no inhibition, whereas a value of $C = 1$ results in the complete growth arrest of the resistant population. We see that the strength of competition is determined by three factors: i) The available resources, K ; ii) the total cell density, $S + R$; and iii) the ratio of turnover to proliferation, d_R/r_R , known as the cell replacement rate. Thus, an increased cell replacement rate - due to either increased turnover or a proliferative cost - results in stronger competition. Alternatively, one can say that the resistant population's growth is limited not by the environmental carrying capacity, K , but by the effective carrying capacity $K_{\text{Eff}} = (1 - d_R/r_R)K$.

Mathematically, we can think of this using the phase plane representation from Figure 2.4c. The closer a tumour is to the r -nullcline, the smaller $dr/d\tau$. The r -nullcline of the full model (Equations (2.5)-(2.6)) is given by:

$$r(\tau) = \underbrace{\left(1 - \frac{\hat{d}_R}{\hat{r}_R} \right)}_{k_{r;\text{Eff}}} \hat{k} - s(\tau), \quad (2.12)$$

where $k_{\text{T,Eff}}$ denotes the resistant population's effective carrying capacity. We see that turnover translates the r -nullcline downwards, thereby reducing the distance of a tumour from the nullcline. In the absence of turnover, a resistance cost merely re-scales time and the $dr/d\tau$ surface. In its presence a cost also shifts the r -nullcline, further reducing the distance of a tumour to the nullcline and extending the time for which a tumour may be controlled (see also Figure A.1).

2.4 Extension to other types of cost

So far we have assumed that the resistance cost manifests itself in the proliferation rate of the cells. However, it is plausible that it might manifest itself also in one of the other population parameters. In this section we will extend our analysis to consider the implications of a cost in the death rate (d_R) and the carrying capacity (K_R). We will begin by considering the impact of each parameter on the effective carrying capacity of the resistant population, before subsequently comparing the impact on treatment dynamics.

2.4.1 Impact on the effective carrying capacity

In the previous section we saw that the r -nullcline provides a useful framework for understanding the impact of a cost on the resistant population. Thus, we will begin by comparing how d_R and K_R affect this nullcline. Equation (2.12) shows that costs in d_R and K_R also act by decreasing the r -intercept, that is the resistant population's effective carrying capacity, $k_{\text{T,Eff}}$ (recall that $\hat{d}_R := \frac{d_R}{r_S}$ and $\hat{k} := \frac{K_R}{K_S}$). As such, any of the three types of cost will bring the resistant population closer to its effective carrying capacity, which will increase the benefit of AT. However, due to the different functional relationship of each parameter with $k_{\text{T,Eff}}$, the magnitude of the impact caused by a cost in each differs. With the following theorem we establish the relative impact of each. That being said, we will subsequently find that this order does not quite hold true in practice, as the treatment dynamics are not solely governed by the proximity to $k_{\text{T,Eff}}$.

Theorem 2. *Let $\sigma \in [0, 1)$ be the relative cost experienced by the resistant population. Further, denote by $r_r^* = \left(1 - \frac{\hat{d}_R}{(1-\sigma)\hat{r}_R}\right)\hat{k}$, $r_d^* = \left(1 - \frac{(1+\sigma)\hat{d}_R}{\hat{r}_R}\right)\hat{k}$, and $r_K^* = \left(1 - \frac{\hat{d}_R}{\hat{r}_R}\right)(1-\sigma)\hat{k}$ the population equilibria under a relative cost $\sigma \in [0, 1]$ in r_R , d_R , and K_R , respectively. Then:*

1. The impact of a cost in r_R on the equilibrium value of r will always be greater than the impact of a cost in d_R . That is, $r_r^* < r_d^*, \forall \sigma \in (0, 1)$.
2. Provided $\frac{\hat{d}_R}{\hat{r}_R} < \frac{\sigma}{1+\sigma}$, the impact of a cost in K_R on the equilibrium value of r will be greater than the impact of a cost in d_R ($r_K^* < r_r^*$). Otherwise a cost in d_R will have a greater impact.
3. The impact of a cost in K_R on the equilibrium value of r will be greater than the impact of a cost in r_R ($r_K^* > r_r^*$) provided $\frac{\hat{d}_R}{\hat{r}_R} < \frac{1-\sigma}{2-\sigma}$. Otherwise a cost in r_R will have a greater impact.

Proof. Claim 1: Comparing r_r^* and r_d^* we find:

$$\begin{aligned}
& r_r^* < r_d^* \\
& \Leftrightarrow \left(1 - \frac{\hat{d}_R}{(1-\sigma)\hat{r}_R}\right) \hat{k} < \left(1 - \frac{(1+\sigma)\hat{d}_R}{\hat{r}_R}\right) \hat{k} \\
& \Leftrightarrow (1+\sigma) < \frac{1}{1-\sigma} \\
& \Leftrightarrow \sigma^2 > 0,
\end{aligned}$$

so that $r_r^* > r_d^*, \forall \sigma \in (0, 1)$.

Claim 2: Similarly, comparing r_K^* and r_d^* :

$$\begin{aligned}
& r_K^* < r_d^* \\
& \Leftrightarrow \left(1 - \frac{\hat{d}_R}{\hat{r}_R}\right) (1-\sigma)\hat{k} < \left(1 - \frac{(1+\sigma)\hat{d}_R}{\hat{r}_R}\right) \hat{k} \\
& \Leftrightarrow -\sigma(1+\sigma)\hat{d}_R - \hat{r}_R(1+\sigma) + \sigma\hat{r}_R(1+\sigma) > -\hat{r}_R \\
& \Leftrightarrow \frac{\hat{d}_R}{\hat{r}_R} < \frac{\sigma}{1+\sigma},
\end{aligned}$$

where we have skipped some steps for brevity.

Claim 3: Finally, comparing r_K^* and r_r^* :

$$\begin{aligned}
& r_K^* < r_r^* \\
& \Leftrightarrow \left(1 - \frac{\hat{d}_R}{\hat{r}_R}\right) (1-\sigma)\hat{k} < \left(1 - \frac{\hat{d}_R}{(1-\sigma)\hat{r}_R}\right) \hat{k} \\
& \Leftrightarrow -(1-\sigma)^2 \frac{\hat{d}_R}{\hat{r}_R} - \sigma(1-\sigma) < -\frac{\hat{d}_R}{\hat{r}_R} \\
& \Leftrightarrow \frac{\hat{d}_R}{\hat{r}_R} < \frac{1-\sigma}{2-\sigma},
\end{aligned}$$

where we again have skipped some steps for brevity. □

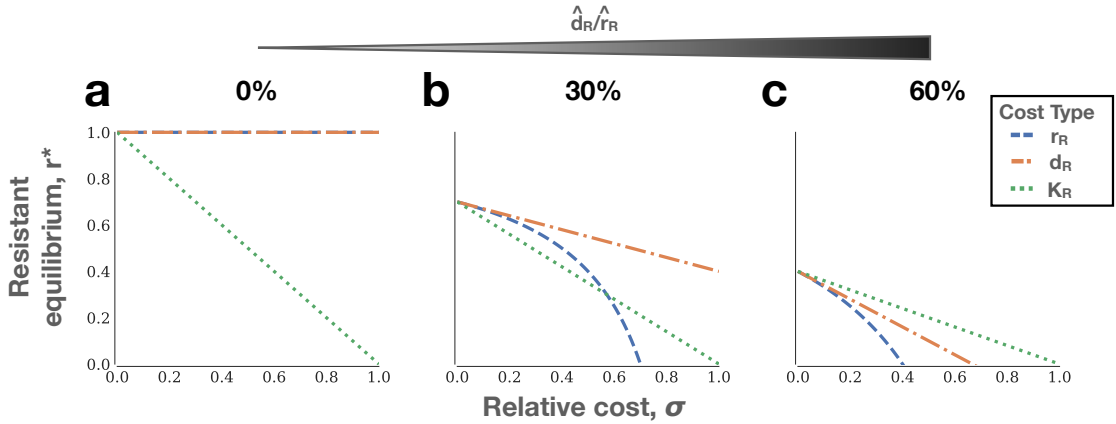


Figure 2.11: Comparison of the impact of different manifestations of resistance costs on the resistant population equilibrium, r^* . The smaller the value of r^* , the longer the tumour can be controlled. Shown are the effects of a cost in r_R ($r_r^*(\sigma)$), d_R ($r_d^*(\sigma)$), and K_R ($r_K^*(\sigma)$) as functions of a relative cost, $\sigma \in [0, 1)$. Note that when $r^* \leq 0$ the resistant population will always go extinct. (a) When there is no turnover, only a cost in K_R has an impact on r^* . (b) For values of $\hat{d}_R/\hat{r}_R < 0.5$, there are two possible cases. For $\sigma \leq 1 - \frac{1}{\hat{r}_R/\hat{d}_R - 1}$, a cost in K_R has the greatest impact. Otherwise a cost in r_R has a greater impact. (c) For $\hat{d}_R/\hat{r}_R > 0.5$, a cost in r_R always has the greatest impact, followed by costs in d_R and K_R .

There are two important conclusions to draw from Theorem 2. Firstly, a cost in the proliferation rate will always have a greater impact than the same relative increase in the death rate (unless $\hat{d}_R = 0$). Secondly, the relative importance of each of the three parameters depends on the ratio \hat{d}_R/\hat{r}_R . In Figure 2.11 we provide a visual illustration of Theorem 2, by plotting r_r^* , r_d^* , and r_K^* as functions of σ for three values of \hat{d}_R/\hat{r}_R (Figure 2.11a). When $\hat{d}_R/\hat{r}_R = 0$, only a cost in K_R has an impact on the effective carrying capacity (Figure 2.11a). In contrast, when $\hat{d}_R/\hat{r}_R > 0.5$ (so that $\frac{\hat{d}_R}{\hat{r}_R} > \frac{\sigma}{1+\sigma}$ and $\frac{\hat{d}_R}{\hat{r}_R} > \frac{1-\sigma}{2-\sigma}$ for all $\sigma \in [0, 1)$), a cost in K_R will have the least impact - that is the largest value of r^* (Figure 2.11c). Finally, for intermediate replacement rates ($0 < \hat{d}_R/\hat{r}_R \leq 0.5$), there are two regimes (Figure 2.11b). Small costs will have the greatest impact if they are in K_R . If the costs are larger, then they will have more impact in r_R . The transition occurs when the $r_r^*(\sigma)$ and $r_K^*(\sigma)$ lines intersect at $\frac{\hat{d}_R}{\hat{r}_R} = \frac{1-\sigma}{2-\sigma}$ (Figure 2.11b). Based on our discussion in Section 2.3.2.1, it seems most likely that in patients $\hat{d}_R/\hat{r}_R < 0.5$, so that our calculations predict that costs in K_R will have the greatest impact, followed by costs in r_R and d_R .

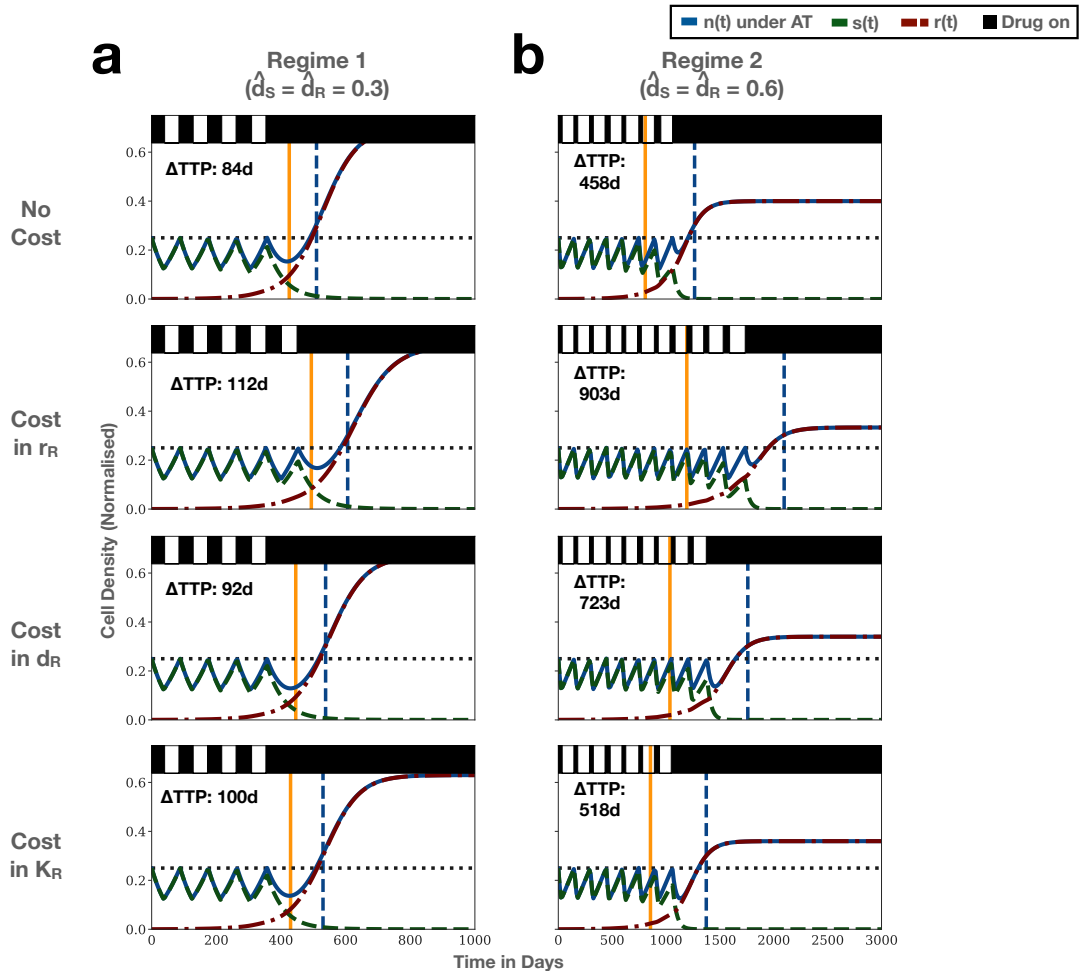


Figure 2.12: Simulations illustrating the impact of a 10% resistance cost in each of r_R , d_R and K_R on the AT treatment dynamics for Tumour 1 ($(n_0, f_R) = (25\%, 0.1\%)$). (a) Simulations corresponding to Figure 2.11b, in which a cost in K_R has the greatest effect on $k_{T,\text{Eff}}$. However, we see that here, due to the small value of n_0 , a cost in r_R has the biggest impact on TTP. (b) Simulations corresponding to Figure 2.11c. Here, the results match those expected from Theorem 2 with, again, a cost in r_R being the most impactful. This illustrates that the comparison of different types of costs also depends on the density of sensitive cells in the tumour. Note that in the simulations of a cost in death rate, d_R is adjusted to take on values 10% greater than in the cases when there is no cost in death rate (so 0.33 or 0.66).

2.4.2 Comparison of the AT treatment dynamics

How do the different possible manifestations of a cost affect treatment dynamics? In Figure 2.12, we show the AT treatment dynamics for a cost of $\sigma = 10\%$ in each of r_R , d_R and K_R for Tumour 1 ($(n_0, f_R) = (25\%, 0.1\%)$). Firstly, we consider the regime $\hat{d}_R/\hat{r}_R = 0.3 < 0.5$ so that a cost in K_R will have the greatest impact on $k_{T,\text{Eff}}$.

Interestingly, we find that - contrary to our expectations from Theorem 2 - a cost in r_R has the largest impact on TTP (Figure 2.12a). As we can see from Figure 2.12a the resistant population takes longer to expand and more AT cycles can be completed when there is a cost in r_R compared to when there is a cost in K_R (4 vs 5 cycles). To explain this, we rewrite Equation (2.6) as:

$$\frac{1}{\hat{r}_R} \frac{dr}{dt} = \underbrace{\left(1 - \frac{\hat{d}_R}{\hat{r}_R}\right)}_{\text{Density-Independent Expansion}} - \underbrace{\frac{s}{\hat{k}}}_{\text{Density-Dependent Inhibition}} \quad (2.13)$$

where we again are assuming that intra-specific competition can be neglected until close to progression ($r \ll s$). We see that the benefit from the reduced carrying capacity (smaller \hat{k}) depends on the density of sensitive cells. If the tumour is far from carrying capacity, as is the case in Figure 2.12 ($s \leq 0.25$), then the benefit is small as the competition is not very strong. In contrast, costs in r_R (and d_R) have an impact *independent* of the proximity to carrying capacity.

In the second case, we assume that the replacement rate is greater than 0.5 ($\hat{d}_S = \hat{d}_R = 0.6$). While this is probably not biologically realistic, we include this case for completeness. We find that here the order in which cost extends TTP coincides with what would be expected from Theorem 2 (Figure 2.12b). This is because the costs in r_R and d_R reduce both the density-dependent, and independent, growth rates more than the cost in K_R (Figure 2.12b).

This analysis further demonstrates that in order to understand the impact of a resistance cost we have to consider not only the cost itself, but also the context in which it occurs (here death rate and initial proximity to carrying capacity). It also suggests two ways in which treatment can seek to maximise the impact of any resistance cost (regardless of its type). Firstly, we may decrease K_R . Secondly, we may aim to increase \hat{d}_R/\hat{r}_R . These insights may be helpful in designing multi-drug AT approaches in the future.

2.5 Application to intermittent androgen deprivation therapy in prostate cancer

The key result of the modelling in the previous sections was that cellular turnover is an important factor in AT. In particular, we predicted that higher turnover facilitates control of drug-resistance and that if a resistance cost is present, then its benefit depends not on its absolute value, but on its value relative to the turnover rate. In

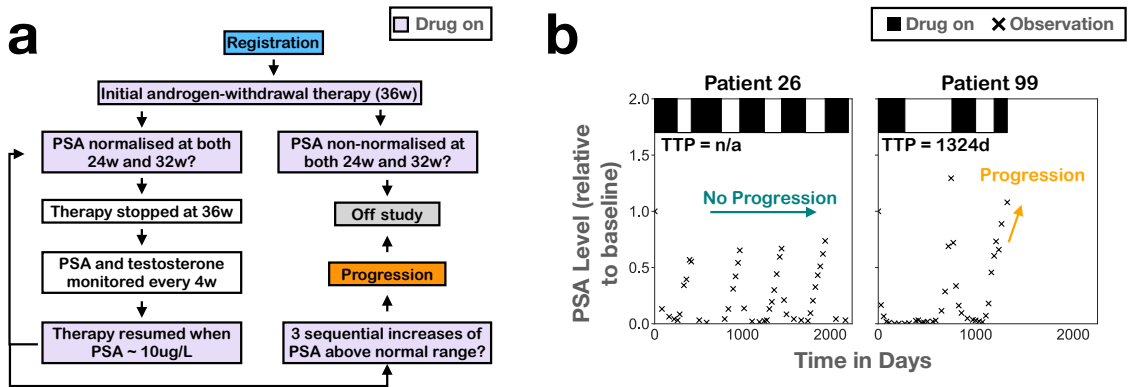


Figure 2.13: Treatment protocol and example trajectories from patients in the trial by Bruchovsky et al [109]. (a) Intermittent treatment algorithm used in the study. Re-drawn and modified with permission from Bruchovsky et al [109] (Copyright © 2006 American Cancer Society). (b) Exemplary patient trajectories from a patient continuously responding for over 5 years, and one patient with progression after 3.5 years.

the following section we will return to our assumption that the fitness cost manifests itself in the cell’s proliferation rate and aim to test this hypothesis. To do so, we will analyse longitudinal tumour response data from 67 patients undergoing intermittent androgen deprivation therapy for recurrent, locally advanced prostate cancer. The data were collected as part of a Phase II trial by Bruchovsky et al [109] and were subsequently made publicly available. We will fit our model (Equations (2.8) & (2.9)) to the serum levels of PSA (recall that this stands for Prostate Specific Antigen) measured in these patients over time. We will show that both cost and turnover are required as model parameters in order to explain the observed patient dynamics. Moreover, some of the patients progress during the trial and, consistent with our hypothesis, we find that these patients are not characterised by low resistance costs or low turnover alone, but by a combination of both. Together these results provide evidence for our modelling predictions, and illustrate how our insights can help to improve our understanding of resistance management in a specific disease setting.

2.5.1 The data

Prostate cancer is the most common cancer in men, and the second most deadly [159]. It is treated by surgery and radiation to eliminate primary and metastatic lesions, followed by hormonal treatment to suppress cancer growth in a systemic fashion. Many prostate cancers are initially dependent on androgen-driven growth stimulation, and

inhibition of this pathway slows their expansion. However, androgen-deprivation treatment is rarely curative due to the emergence of androgen-independent populations, and represents a significant burden on patients' finances and quality-of-life [109, 160].

In an attempt to address these issues, in the 1990s Bruchovsky and colleagues [161, 109] proposed to cycle treatment, with the aim of alleviating side-effects and better controlling cancer stem cells (as such, in fact, based on arguments quite similar to those for AT). Following promising results in a mouse model [161] Bruchovsky et al launched the "Canadian Prospective Phase II Trial of Intermittent Androgen Suppression for Men in Biochemical Recurrence after Radiotherapy for Locally Advanced Prostate Cancer" which was carried out at four Canadian centres between 1995 and 2001. Its aim was to test the feasibility of treatment cycling, its safety, and its impact on patients' quality-of-life. The study enrolled a total of 109 patients and followed them until the end of the trial period or until disease progression, using serum PSA levels to monitor disease burden. Patients were treated with cyproterone acetate and leuprolide acetate according to the algorithm shown in Figure 2.13a. Progression was defined as a series of three sequential increases of PSA above $4.0\mu\text{g}/\text{L}$ despite castrate levels of serum testosterone [109]. Example trajectories of two patients without and with progression are shown in Figure 2.13b.

There are two key differences between this schedule and the one used by Zhang et al [97] in the AT clinical trial: 1) The lead-in phase, and ii) use of an absolute, instead of a relative, threshold for treatment modulation. The effect of the latter is that in many patients the PSA level is reduced much further than the 50% level in the AT study (Figure 2.13b). Furthermore, some [97, 123] have argued that the lead-in phase eliminated most of the sensitive cells and accelerated competitive release. Nevertheless, most patients completed at least one cycle (93/109 patients) with the average number of cycles completed being 2.3 ± 0.1 cycles (median, 3.0 cycles; range, 1-5 cycles). In total, 12 patients progressed during the trial, and a further 6 continued to respond but developed a metastasis [109]. While overall survival was not one of the end points of this study, it was evaluated in subsequent Phase III trials. These concluded that intermittent treatment provided neither an advantage, nor a disadvantage, in survival, but did improve quality-of-life [110, 111].

2.5.2 Data processing and data fitting

The data were downloaded from <http://www.nicholasbruchovsky.com/clinicalResearch.html> on July 15th 2020. Only those data from patients who

qualified for the intermittent arm were available ($n=72$). From these patients we excluded those who developed a metastasis during treatment ($n=5$), so to avoid potentially confounding effects from a change in the number of lesions. This yielded a total of 67 patients for analysis.

Subsequently, we fitted the model (Equations (2.8)-(2.9)) to each patient by minimising the root mean squared error (RMSE) between the normalised PSA measurements, $\hat{y}(t)/\hat{y}(0)$, where $\hat{y}(t)$ is the PSA measurement at time t , and the normalised predicted cell density, $n(t)/n_0$. The values for r_S and \hat{d}_D were fixed as in Table 2.1, and the remaining four parameters (\hat{r}_R , \hat{d}_T , n_0 , and f_R) were allowed to vary freely between the following bounds: $\hat{r}_R \in [0, 1]$, $\hat{d}_T \in [0, 1]$, $n_0 \in [0.1, 1]$, and $f_R \in [1 \times 10^{-5}, 0.25]$. Fitting was done using the Levenberg-Marquardt algorithm, and repeated 25 times for each patient starting from different initial parameter combinations, with the best fit according to the Akaike Information Criterion (AIC) taken forward for analysis. Fitting was done using the lmfit package in Python [162].

2.5.3 The model is able to recapitulate treatment dynamics

In Figure 2.14 we show the best fit for each patient as determined by AIC. We find that even though our model was constructed as a conceptual tool, it can recapitulate individual patient dynamics for the majority of patients, and that it can describe patients who continuously respond, as well as those who eventually progress. Furthermore, we find that for each patient the optimal fit is well defined, as we do not observe evidence for multiple optima amongst the 25 independent runs for each patient and because model selection with a different model selection criterion (Bayesian Information Criterion (BIC)) yields the same optimum as selection via AIC in all patients.

That being said, while for some patients, such as Patient 41, the model describes the quantitative change in PSA over time almost exactly, in others the fit is more qualitative (e.g. Patient 55; Figure 2.14). Moreover, in 15 patients, marked by a grey background in Figure 2.14, we find that the model either fails to describe one or more cycles (e.g. Patient 77 or Patient 105), or predicts resistance where none was observed (Patient 104 or Patient 6). Investigating why the fits are poor for these patients was outside the scope of this work, however, based on a visual examination it appears that in patients for whom the model fit well the cycle timing remained fairly constant over time, whereas in patients for whom the model struggled there was a rapid change in the speed of cycling (e.g. Patient 77). Because of their poor fits we excluded these 15 patients from further analysis.

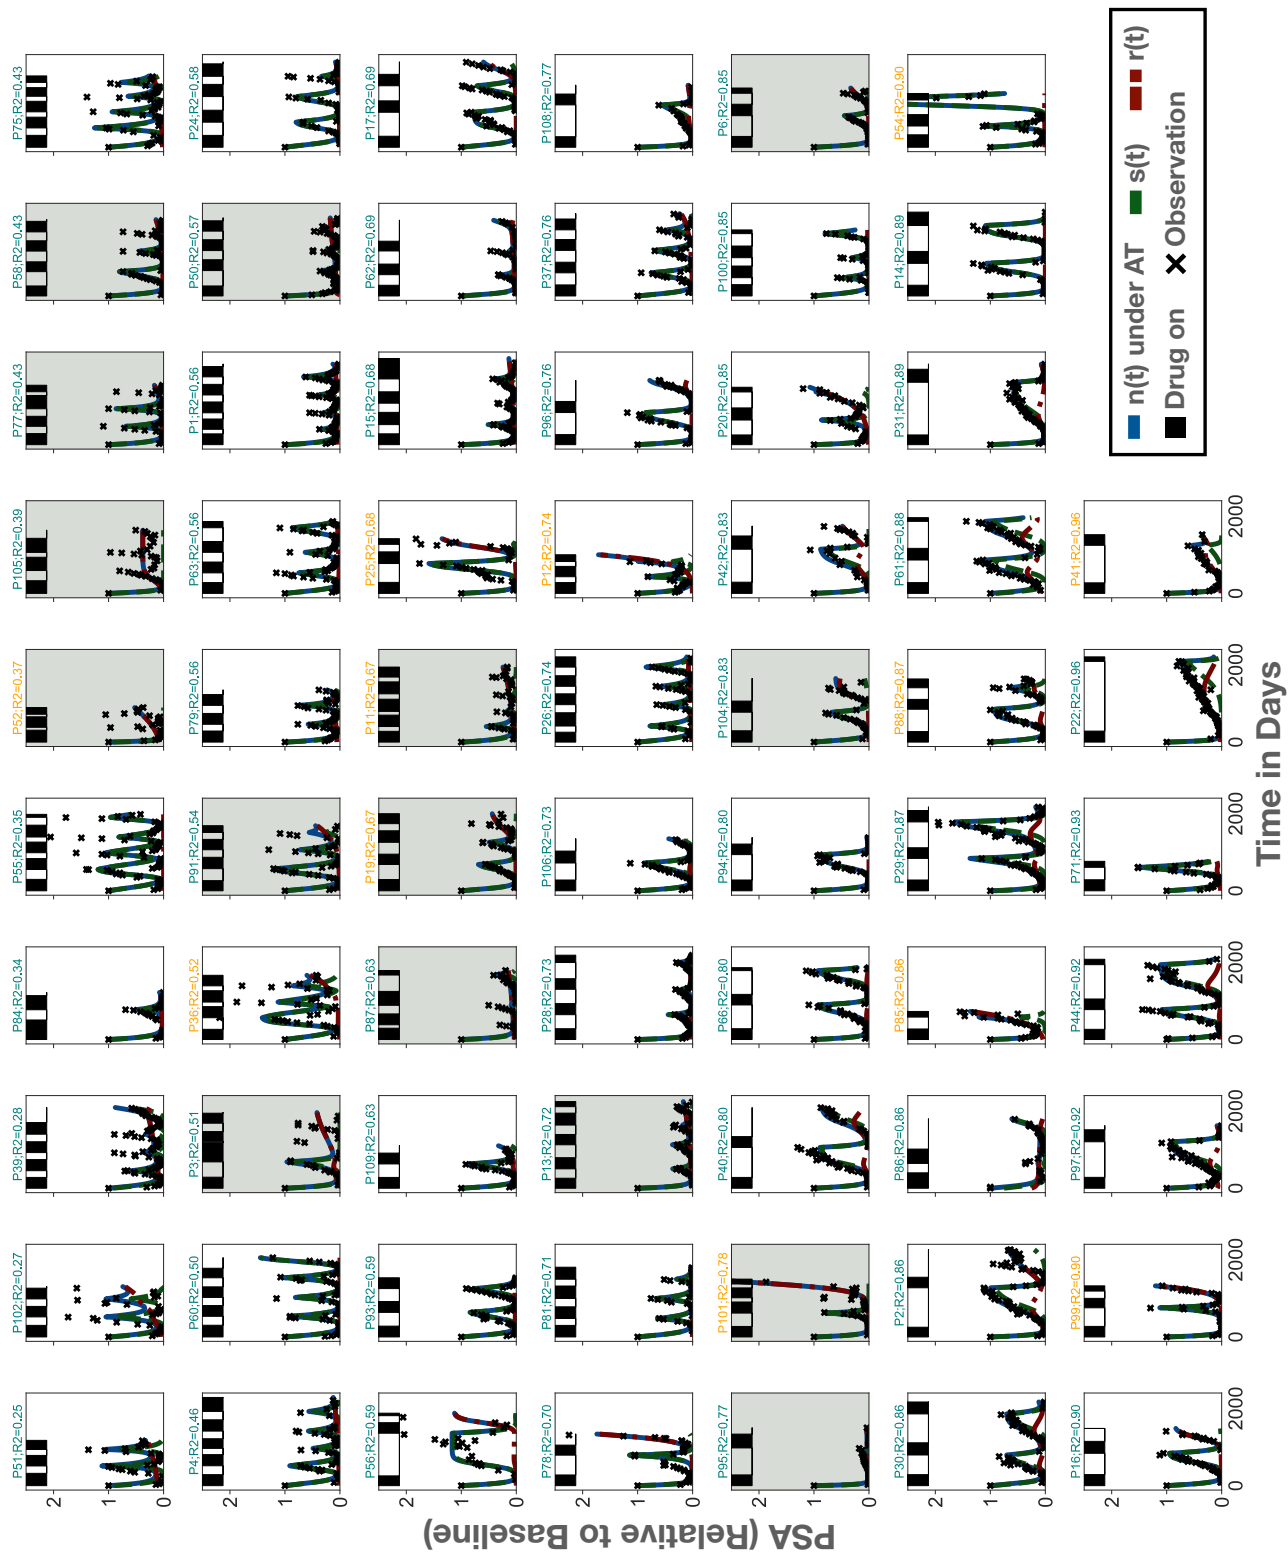


Figure 2.14: Overview of all 67 patient fits arranged by the r^2 value of the best model fit. Title colour indicates whether a patient progressed (orange) or not (green). Patients who were excluded from further analysis due to poor model fits are marked with a grey background.

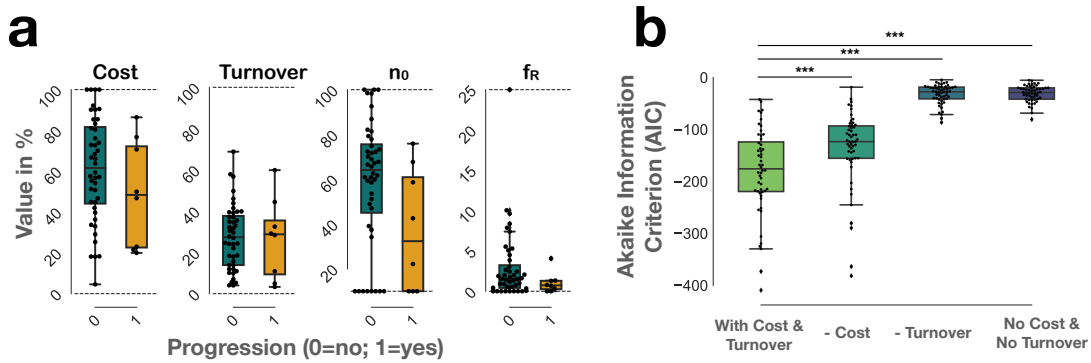


Figure 2.15: The role of cost and turnover in describing patient dynamics. **(a)** Distribution of fitted parameters for patients with and without progression. Dashed lines mark the limits imposed on the parameters during fitting. **(b)** Removing turnover or cost as a variable in the model significantly decreases the goodness-of-fit across the patient cohort, indicating both variables are important in explaining response. Stars denote p-values of less than 1% in paired student t-tests. This suggests that both cost and turnover are important for describing treatment response in these patients.

2.5.4 Cost and turnover are required for describing patient dynamics

Next we examined what role cost and turnover play in the response dynamics. In Figure 2.15a we plot the distributions of the estimated parameters for the 52 patients analysed. In most patients the estimate for both cost and turnover is far from 0, which suggests that both parameters are important for describing the treatment dynamics.

To further corroborate this finding, we performed a systematic model comparison analysis in which we removed first cost from the model, then turnover, and finally both. In each case we fitted the reduced model to the 52 patients carried forward for analysis using the same protocol as before. It is clear that removing a degree-of-freedom will reduce the fit of the model. So we used AIC to perform the model comparison for each patient, which penalises the extra complexity in the more complicated models. A difference in AIC of more than 2 is typically seen as evidence that the model with the smaller AIC is a better descriptor of the data [163]. In Figure 2.15b we plot the AIC values for all four models across the patient cohort. This shows that for the majority of patients removing cost and turnover from the model significantly reduces the fit, more than would be expected from simply removing a degree-of-freedom. The only exceptions are Patients 56 and 102 for whom the model without cost fits as well as the more complex four parameter model. Consistent with this observation, we find that dropping cost from the model increases AIC less than removing turnover, suggesting that while both play a role the latter is more important

than the former. Repeating the comparison using the BIC yields identical results (not shown). Overall, this supports our hypothesis that both resistance costs and turnover are important for explaining the time-course of response to treatment in patients.

2.5.5 Neither cost nor turnover in isolation correlate with outcome

Having established a role of cost and turnover in describing treatment response, we next asked whether they could predict outcome. To do so we compared the estimated values for both parameters for patients with and without progression (Figure 2.15a). Interestingly, we do not detect a significant difference for neither cost (Mann Whitney U test, $U = 134$, p-value > 0.05), nor turnover (Mann Whitney U test, $U = 164$, p-value > 0.05). The fact that cost in isolation is not predictive by itself (although there is a trend towards a difference) is consistent with our theoretical analyses in Section 2.3, in which we showed that a cost is not necessary nor sufficient for being able to manage resistance. However, the finding that there is no difference in turnover is contrary to our expectations. Based on our modelling in Section 2.3 we would have expected to see higher turnover in patients without progression. A possible explanation for why our prediction does not hold is that tumours with higher turnover may also harbour more resistance as these tumours will have passed through more generations before detection [155]. However, testing for a correlation between turnover and f_R we do not find evidence for this explanation either (Pearson's Correlation Coefficient: -0.038 , p-value > 0.05).

For completeness, we also examine differences in the other two parameters (see also Figure 2.15a). This yields a statistically significant difference in n_0 (Mann Whitney U test, $U = 104.5$, p-value > 0.035), but not in f_R (Mann Whitney U test, $U = 121$, p-value > 0.05). The fact that patients without progression have a higher estimate for n_0 is consistent with the result from Section 2.3 that tumours are controllable for longer the closer they are to carrying capacity. However, one should be somewhat careful with this interpretation, as it is unclear what carrying capacity means in this context. PSA is a serum marker reflecting the systemic disease burden to which multiple tumour lesions may be contributing. As such it is probably more helpful to interpret this result in terms of the nature of the growth curve that it produces. A large value of n_0 in the model corresponds to tumour trajectories which do not grow much above their initial size and/or show sub-exponential growth. Conversely, small values of n_0 suggest that such tumours do not exhibit signs of growth inhibition.

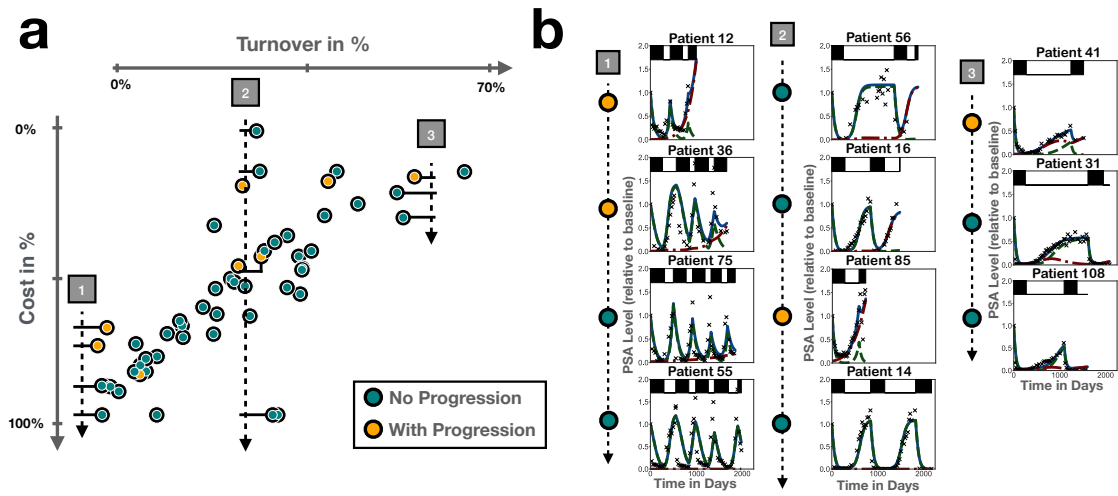


Figure 2.16: The importance of combining cost and turnover to predict outcome of intermittent therapy in the Bruchovsky et al trial [109]. (a) Plot of the estimates for turnover and cost (each point represents one patient; colour indicates whether or not a patient progressed on the trial) showing clustering of fits in a defined region of parameter space consistent with the study inclusion criteria. Moreover, patients who progress (yellow) tend to fall into the upper half space, indicating that these patients are defined not by small values of cost or turnover but by a small combination of the two. (b) Examples of patient fits along three transects through the space in (a), showing how turnover and cost separate patients with different cycle dynamics, and patients with and without disease progression.

2.5.6 Context matters: Cost and turnover in combination predict outcome

The key result of our modelling analysis in Section 2.3 was the observation that the impact of a resistance cost is modulated by the turnover in the tumour. Given that neither cost nor turnover was predictive in isolation, we thus examined their interaction. This reveals a strong negative correlation, with most patients clustering closely along a line in parameter space (Pearson’s correlation coefficient: -0.77 , p -value < 0.01 ; Figure 2.16a). Comparing this observation with our theoretical predictions of how cost and turnover affect the treatment dynamics in Figures 2.9a & b we propose the following explanation. We hypothesise that this trade-off, which is not imposed but emerges from the fitting, reflects the study’s inclusion criteria. Firstly, no patient falls into the region in which both turnover and cost, and so response, are small, because such patients would not have achieved PSA level normalisation in the lead-in phase required for study inclusion. Similarly, no patients are seen in the bottom right corner where our model predicts long term tumour control, consistent with the fact

that all patients in the study were refractory after initial therapy.

Moreover, when we colour patients by whether or not they experience disease progression, we observe that patients with progression (orange) tend to fall along the upper edge of the patient cluster (Figure 2.16a). This means that patients who progress have a smaller combination of cost and turnover than patients who do not progress. Consistent with this observation we find that if we compare the sum of turnover and cost between the two groups, this difference is highly statistically significant (Mann Whitney U test, $U = 82$, p-value < 0.01). Furthermore, when we examine Patient 56 and Patient 16, who contradict this pattern as they are responders but far above the upper edge of the cluster, we see that for neither is it actually known whether they are still responsive to treatment (top two points along Transect 2 in Figure 2.16). Both have just concluded an off-treatment phase and have not yet re-started treatment. In fact for both of these patients the last off-phase has been shorter than the previous one, so that it is plausible that they might have developed resistance (Figure 2.16b).

In addition, when we examine the treatment dynamics at different points in the cost-turnover space, a distinct pattern emerges (Figure 2.16b). As we move left to right along the turnover axis patients are sorted in an unsupervised fashion according to their cycling speed. As we move top to bottom along the cost axis, patients are increasingly likely to respond for longer. We see that depending on their turnover value, patients may cycle at different speeds, and that their TTP is determined by the cost of resistance. This shows that multiple factors have to be considered to understand response to intermittent treatment, and would explain why Bruchovsky et al [109] did not find that the cycling speed per se was predictive of outcome.

To sum up, while there is no evidence in these data that higher turnover by itself is correlated with longer resistance control, our results provide strong support for the notion that the interaction between cost and turnover is important. Our analysis suggests that patients who progress are characterised not by a lack of a resistance cost or slow turnover, but by a value of the combination of the two that is too small. For example, if turnover is small patients will progress even if the cost is high, whereas if turnover is high, a much lower cost may be required to see long term response (Figure 2.16a). Moreover, we find that the cost-turnover combination of a patient also correlates with their cycling dynamics, so that different areas of parameter space correspond to different types of cycling dynamics. This observation may help to better understand the biology underlying the different speeds of cycling observed in patients, and we will return to it in the next chapter.




Figure 2.17: *Three strategies to improve tumour control by increasing competition, illustrated on Patient 99 who progressed on the trial (Figure 2.13B). Parameters are those obtained from fitting ($n_0 = 10\%$, $f_R = 1.25 \times 10^{-3}\%$, $\hat{r}_R = 79\%$, $\hat{d}_T = 44\%$). Baseline: AT with drug withdrawal after a 99% tumour burden reduction, mimicking the Bruchovsky protocol. Strategy 1: AT with drug withdrawal at 50%. Strategy 2: secondary drug, such as an anti-angiogenic drug, to reduce K (here by 50%). Strategy 3: secondary drug, such as a low dose chemotherapy, to increase turnover (here by 20%)*

2.6 Improving AT by amplifying competition

To conclude this chapter, we will discuss how, based on our insights, one may be able to improve AT. Specifically, we deduce that there are three ways in which one may be able to amplify intra-tumoral competition:

1. Reduction of the AT threshold: Firstly, one can lower the threshold for treatment withdrawal, so that treatment is withdrawn earlier and more frequently, assuming the increased tumour burden is acceptable. This conclusion is based on Section 2.3.1.4 which showed that in our model the competitive suppression of resistance is maximised the closer the tumour is to its carrying capacity.
2. Targeting the environmental carrying capacity, K : Secondly, one can target the environmental carrying capacity, for example by administering an anti-angiogenic drug to reduce the supply of nutrients.
3. Increasing cell turnover: Finally, one can increase cell turnover to reduce the growth of resistant cells and amplify the strength of selection against them. This may be done, for example, by adding low-dose chemotherapy.

In Figure 2.17 we illustrate each strategy on Patient 99 who relapses during intermittent androgen deprivation treatment (see Figure 2.13B). Interestingly, while it is Strategy 1 which optimises the gain of AT over CT, it is Strategy 3 which maximises absolute TTP. This shows that each strategy has a unique effect on the treatment response, and motivates further research into multi-drug AT protocols.




Figure 2.18: The interplay between resistance costs and turnover in AT. Sensitive cells slow the growth of resistant cells through competition, but are killed by treatment. A cost of resistance reduces the resistant growth rate. Turnover modulates the speed of selection against drug-resistance by amplifying the effects of competition and resistance costs.

2.7 Discussion

The aim of this chapter was to develop a theoretical understanding of AT, and gain insights into the circumstances under which a resistance cost is required, and under which circumstances it is not. To do so, we developed a simple two-population Lotka-Volterra competition model and treated it according to the AT algorithm from Zhang et al [97]. We showed that a resistance cost is not necessary for AT to extend TTP, and used the work by Hansen et al [1] to provide an explanation for why this is the case: If we assume that sensitive cells competitively inhibit resistant cells, and there is no harm from the larger tumour burden, then the presence of sensitive cells will always slow down the growth of the resistant cells. This point is trivial, but it has a crucial implication which is not as clear in the more complex models analysed previously: As long as we assume no harm from the increased tumour burden, and that growth is deterministic, AT will perform no worse than CT. This point is also discussed in more detail by the recent works by Viossat and Noble [126] and Hansen and Read [129].

That being said, how much we actually gain depends on the following four factors: i) the initial fraction of resistant cells (f_R), ii) the proximity of the tumour to environmental carrying capacity (n_0), iii) the resistant cell proliferation rate (r_R), and iv) the rate of cellular turnover (d_T). While the importance of the first two factors has been previously recognised [115, 1, 124], it has not been clear what exactly it means for a tumour to be close to its carrying capacity. Carrying capacity is defined as the

population size at which growth saturates [164]. As such it is a multi-faceted concept which may be defined at different spatial scales, and which may vary over time and space. For example, one can define a systemic carrying capacity for the whole body, or local carrying capacities for individual tumour lesions. Moreover, carrying capacity is determined by space and resource availability, which will change as the tumour recruits vasculature and stromal support, or gains the ability to invade surrounding tissue. In addition, it is determined by the ratio of birth to death in the tumour. Working with the popular $r - K$ form of the logistic growth equation this fact is easily overlooked, and results in the misleading perception that r and K are independent. A key contribution of this chapter is to highlight that not only environmental constraints, but also cell proliferation and death rates, determine the impact of intratumoral competition. The shorter a cell's life span, the fewer opportunities it will have to divide, and the greater will be the impact of interference through intra- and inter-specific competition. Changes in proliferation or death rates, due to mutation, immune predation, or treatment, will alter the tumour's carrying capacity and so change the competition between tumour cells. The relationship between r and K has been an area of intense debate in ecology, centred in particular around the so-called $r - K$ selection hypothesis [164]. Realisation of their inter-dependence has been an important step in consolidating ecological and evolutionary theory [165, 166, 167], and may similarly help to improve our understanding of tumour evolution and treatment.

An important implication of this dynamic view of carrying capacity is that tissue turnover plays a role in AT (Figure 2.18). Specifically, our modelling indicates that higher turnover facilitates the control of drug-resistance and that if a resistance cost is present, then its benefit depends not just on its absolute value, but on its value relative to the turnover rate. We tested this hypothesis by fitting our model to response data from prostate cancer patients undergoing intermittent androgen deprivation therapy. This showed that both cost and turnover are required to explain the treatment dynamics, as their removal significantly reduces goodness-of-fit, even when a penalty for the additional degrees-of-freedom is applied. Interestingly, and contrary to our expectations, we did not find a difference in turnover when we tested for differences between patients who progressed and those who did not. This may be due to tumours with higher turnover also acquiring more resistance as these tumours will have passed through more generations before detection [155], although we did not detect such a correlation ourselves. However, when we examined both cost and turnover together, we saw that patients who progressed tended to be defined by a smaller combination of cost and turnover than patients who did not progress. While

these results should be viewed with some caution, as neither cost nor turnover were measured directly, they are consistent with our hypothesis, and show how our insights may help to understand treatment response.

One reason why it was difficult to test our hypothesis more directly was that relatively little is known about turnover rates in human tumours. Most available data were collected between the 1960s and 1990s based on a method pioneered by Steel [156]. While the measures are indirect and have to be interpreted with caution, most of these studies found that tumours grew significantly slower than expected based on cell proliferation rates, suggesting that the rate of cell death closely matches that of cell production [156, 36, 168]. In light of our findings we advocate that more focus should be put on quantifying the turnover rates in tumours. Not only may information on turnover rates help to identify suitable candidates for AT, it may also be important in interpreting data from experimental model systems. While *in vitro* no resistance cost may be observed, tumour control may still be possible *in vivo* or in patients as turnover, via greater nutrient deprivation, and immune suppression, will intensify competition, and amplify even small fitness differences.

While we focussed our attention on the impact of a cost in the proliferation rate, it is clear that it may also manifest itself in other ways. We showed that our analysis easily extends also to costs in the turnover rate and carrying capacity, where either costs in the proliferation rate or costs in the carrying capacity will have the greatest impact depending on the cell replacement rate.

In aiming to keep our model simple we made a number of important simplifications. Firstly, similar to previous studies [115, 1], we assumed that the tumour is not curable. However, including a resistance cost and greater turnover also results in fewer cells at the nadir during CT. This implies that the tumours in which AT will be most effective are those in which CT just misses being curative, an observation also made elsewhere [59, 124]. Thus, the very long gains predicted under AT have to be viewed with some skepticism, since in a proportion of these patients CT would have potentially cured the tumour. That being said, AT is intended for advanced settings in which the available treatment is rarely curative, which suggests that such cases will be rare [98, 169]. Moreover, it seems plausible that if one finds that AT controls the tumour very well, one could then decide to switch to a curative approach.

A further important caveat of our study is that we did not account for the potential health-risks associated with the increased tumour burden under AT. As a result, our model predicts that the longest TTP can be achieved by keeping the tumour as close to its initial size as possible. This is true for most studies in this area

[60, 115, 122, 114, 120]. However, a higher tumour level increases the potential for *de novo* resistance acquisition and metastasis, and under certain circumstances these costs may outweigh the benefits of AT [1, 126, 170]. Investigating the risks of AT is an important direction of future research.

Finally, in this chapter we neglected the impact of space. However, previous work [115, 124] has demonstrated that the tumour's spatial architecture is an important factor in AT. For example, if resistant cells can be "trapped" by surrounding sensitive cells, then the tumour may be controlled for a long time [115, 124], as resistant cells are close to their local carrying capacity. In the next chapter we will address this shortcoming, and examine whether and how spatial effects impact our conclusions.

Chapter 3

The impact of space and stochasticity on intra-tumoral competition and adaptive therapy

3.1 Introduction

In the previous chapter we analysed a simple ODE model of AT to develop a basic understanding of its guiding principles. However, in doing so, we assumed perfect mixing of cells, and neglected the impact of spatial constraints on the growth and interaction of cells. Furthermore, by using a deterministic modelling framework we neglected the potential role of stochastic events. The aim of this chapter is to address these issues. We will re-formulate our ODE model as a 2-D, agent-based model (ABM) and test to what extent our conclusions from the previous chapter translate, and where the two models differ.¹

3.1.1 Spatial modelling of AT

While a number of modelling studies have been carried out in order to elucidate the eco-evolutionary principles of AT [60, 115, 122, 124, 120, 132], so far only two of these have considered the role of space [115, 124]. Bacevic et al [115] developed a 2-D hybrid cellular automaton model, in which each cell occupies a site on an equi-spaced grid. Cells in this model proliferate and die at a rate which depends on the local oxygen concentration, which is modelled by a PDE. Moreover, cells compete for space, in that division is only possible if there is a vacant space in the cell's neighbourhood. Finally,

¹A publication based on this chapter, titled "*Spatial structure impacts adaptive therapy by shaping intra-tumoral competition*", is in preparation (see [171] for a pre-print). Both the work and the writing of this publication were carried out by myself under supervision of my supervisors, and guidance from our co-authors.

the authors also model pushing of cells by allowing a cell to displace its neighbours in order to generate space for proliferation. They divide the tumour into drug-sensitive and resistant cells and calibrate their model using *in vitro* data from a colorectal cancer cell line undergoing CDKi treatment. Subsequently, Bacevic et al [115] use their model to investigate the importance of spatial structure for controlling resistance with competitive inhibition. While they use continuous low dose, and not adaptive scheduling, their results are nevertheless relevant to AT, as the mechanism through which resistance is controlled is competition. Bacevic et al [115] show that if resistant cells are initially near the centre of the tumour, where competition for space is strong, then control via low dose treatment is much more likely to be successful than if cells are located near the periphery of the tumour, or seeded randomly. Moreover, to test their prediction that spatial constraints on resistant cell growth are important, they carry out *in vitro* experiments in which they treat a mixture of sensitive and resistant cells either at MTD or at low dose. In accordance with their theoretical predictions they find that low dose therapy results in a smaller final resistance fraction than MTD treatment in a more constrained, 3-D *in vitro* spheroid, but not in a relatively unconstrained, 2-D *in vitro* cell culture system.

The second study which has examined the role of space in AT is the work by Gallaher et al [124] who study a 2-D off-lattice model of tumour growth. Cells are modelled as discs which repel each other with a force that depends linearly on the distance between the cell centres, and whose movement is dampened through friction. Cells divide after they have completed a cell cycle, where progression through the cell cycle is interrupted if local crowding exceeds a certain threshold. Gallaher et al [124] do not include natural turnover or a nutrient in their model so that cells only die when exposed to drug. Drug kills dividing cells where a cell's sensitivity to drug is modelled as a continuous trait. Finally, the authors assume a resistance cost, so that cell cycle time is inversely related to the cell's drug sensitivity. In their study, Gallaher et al [124] compare continuous MTD treatment with the two adaptive strategies trialled by Enriquez-Navas et al [113] in a breast cancer mouse model. The first of these strategies modulates the dose of the treatment according to the tumour's response, whereas the second always applies MTD but skips doses if the tumour responds well, akin to the algorithm used in the prostate cancer clinical trial [97]. Gallaher et al [124] corroborate the findings of Bacevic et al [115] that AT is most likely to be successful if the resistant cells are spatially constrained at the centre of the tumour. Furthermore, they show that both cell migration and genetic mutation represent challenges for resistance management because they allow resistant cells to move to, or emerge at,

the edge of the tumour, thereby weakening spatial constraints. Moreover, they find that the initial degree and heterogeneity of drug sensitivity in the tumour determine which of the three treatment strategies will be most successful, an observation which we will discuss in more detail in the next chapter.

In summary, these studies have shown that the spatial distribution of resistant cells and their ability to escape crowding through pushing or movement determine how well they may be controlled via competitive inhibition. As such, this work helps to provide an interpretation of the result from non-spatial models that the “initial proximity to carrying capacity” (n_0 in Chapter 2) is an important factor in AT. Furthermore, they corroborate the role of the initial resistance fraction, and the observation that AT may be beneficial under a wide range of conditions. However, because both these studies involve rather complex assumptions about cell physics and metabolism, it is not clear which aspects of their findings are due to the spatially explicit nature of the models, and which due to these extra assumptions. It is also not known how these results compare quantitatively to those obtained by non-spatial models. Moreover, neither of these papers has examined the role of turnover, which we showed in the last chapter to have great impact on the response to AT.

The aim of this chapter is to address these questions. We develop a 2-D, on-lattice, ABM using the same set of assumptions as for the ODE model studied in Chapter 2. In doing so, we hope to bridge the non-spatial and spatial AT modelling literature, and tackle the following three questions: Firstly, we seek to examine whether turnover remains a facilitator of AT if space and stochastic effects are taken into account. In particular, we will make use of the fact that in the ABM we can monitor cell-cell interactions which allows us, for the first time, to quantify competitive inhibition during AT, and characterise the impact of resistance costs and turnover on competition. Secondly, we will perform a quantitative comparison of the spatial and non-spatial model in order to test to what degree predictions on the benefit of AT agree. Together, these analyses will allow us to dissect more clearly how space affects AT. Finally, we will re-visit the treatment responses of the prostate cancer patients on the intermittent androgen-deprivation trial by Bruchovsky et al [109]. By examining the data through the lens of the ABM we will corroborate our findings from Chapter 2, and show that there may also be differences in the tumour spatial architectures between patients. Overall, this chapter helps to provide a more detailed understanding of the spatial competition between sensitive and resistant cells during AT, and shows that both the spatial architecture of the tumour, and stochastic effects, may significantly affect treatment outcomes.

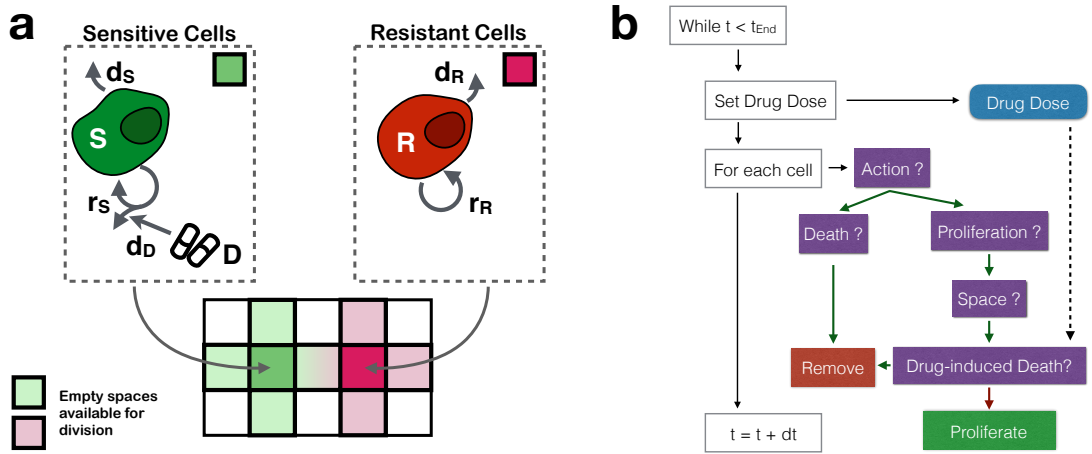


Figure 3.1: The agent-based tumour model: (a) The tumour is modelled as a mixture of drug-sensitive and resistant cells, where each cell occupies a square on a 2-D equispaced lattice. Cells divide and die at constant rates. Daughter cells are placed in empty squares in a cell’s von Neumann neighbourhood. Drug will kill dividing sensitive cells at a rate d_D . (b) Flow diagram of the simulation algorithm.

3.2 The computational model

3.2.1 Model assumptions and parameters

We will simulate the spatio-temporal evolution of the tumour using a 2-D, on-lattice, ABM. In this model each cell occupies a single site in an $l \times l$ square lattice, and behaves according to a set of rules which are defined on the level of the individual cell. We choose this type of model over alternative approaches such as PDEs or stochastic differential equations as it will allow us to explicitly study cell-cell interactions such as competition for space. Moreover, both previous spatial models in this area have been ABMs, so that this way the comparisons we draw between the spatial and the non-spatial model will be more easily translatable to the existing literature.

This is also the reason why we choose to model tumour growth in 2-D and not 3-D. In addition, this choice has the benefit that it reduces the computational costs of the simulations, while remaining biologically relevant as it can be seen to model the dynamics in an *in vitro*, flat-bottom cell-culture flask. Moreover, a 2-D model can, in a sense, be seen as an upper bound for the potential impact of space as it maximises the spatial separation between cells. Baker and Simpson [172], for example, have shown that a 3-D on-lattice model of a birth-death process is in greater agreement with the corresponding mean-field equations than the equivalent 2-D model. This is

due to the greater scope for cell-cell interaction and mixing in 3-D [172].

In order to allow a comparison with our results from Chapter 2 we carry over most of our model assumptions and parameters. As such we assume that the tumour is composed of a mixture of drug-sensitive and resistant cells, each of which occupies a single site in an $l \times l$ square lattice with no-flux boundary conditions. Further, we assume that cells behave according to the following rules (Figure 3.1a):

- Initially, there are a total of N_0 cancer cells spread randomly in the tissue, of which a fraction f_R is resistant.
- Sensitive and resistant cells attempt to divide at constant rates r_S and r_R (in units: day^{-1}), respectively. If there is at least one empty site in the cell's von Neumann neighbourhood, then the cell will divide and the daughter will be placed randomly in one of the empty spaces in the neighbourhood.
- Cells die at a constant rate δ_T (in units: d^{-1}). For notational convenience, we will express this rate relative to the sensitive cell growth rate, so that $d_T = r_S/\delta_T$. In addition, we will make the simplifying assumption that both sensitive and resistant cells die at the same rate, $d_S = d_R = d_T$.
- Movement of cells is neglected.
- The domain is sufficiently small so that drug concentration $D(t) \in [0, D_{\text{Max}}]$ is assumed homogeneous throughout the tissue.
- A sensitive cell which is currently undergoing mitosis - that is, it has attempted division and has space available in its neighbourhood - is killed by drug with probability $d_D D(t)$, where $d_D < D_{\text{Max}}$.
- Dead cells are immediately removed from the domain.

We denote the number of cells in each population at time t by $S(t)$ and $R(t)$, and the total number by $N(t) = S(t) + R(t)$, respectively. We parameterise the model using the same parameter values for cell proliferation and death rates as in Chapter 2. A summary of all parameters is shown in Table 3.1.

AT is aimed for advanced disease settings in which the tumour has invaded surrounding tissues and has also likely metastasised. To model such invasive disease we assume that the tumour cells are initially dispersed randomly throughout the domain, where empty space can be thought of as representing normal tissue or tumour stroma. The other reason we make this choice of initial condition is that it will allow us to

Table 3.1: Summary of mathematical model variables and parameters.

Parameter	Description	Range	Comment
l	Grid size (Total number of sites: $l \times l$)	100	
dt	Simulation time step	1d	
$S(t)$	Number of sensitive cells	0.0 - l^2	
$R(t)$	Number of resistant cells	0.0 - l^2	
$N(t)$	Total number of cells	0.0 - l^2	Determined by $N(t) = S(t) + R(t)$
r_S	Sensitive cell proliferation rate	0.027d^{-1}	Adopted from [97].
c_R	Resistance cost (resistant proliferation rate: $r_R = (1 - c_R)r_S$)	0 - 50%	Range as in Chapter 2.
d_T	Cell death rate (as a percentage of r_S)	0 - 50%	Range as in Chapter 2.
d_D	Drug-induced cell kill probability of sensitive cells at $D(t) = D_{\text{Max}}$	0.75	Adopted from [120].
$n_0 = \frac{N_0}{l^2}$	Initial cell density (as a percentage of total carrying capacity)	25 - 75%	Parameter sweep; values within this range reported by [153].
$f_R := \frac{R_0}{N_0}$	Initial resistant cell fraction (as a percentage of initial cell density)	0.1 - 10%	Range as in Chapter 2.
$f_S := \frac{S_0}{N_0}$	Initial sensitive cell fraction (as a percentage of initial cell density)	90 - 99.9%	Determined by $f_S = 1 - f_R$.

understand whether AT can still work even if resistant cells are not initially fully encircled by sensitive cells, as was the case in the work by Bacevic et al [115] and Gallaher et al [124].

Finally, we will again consider the same two treatment schedules as in the previous chapter:

1. CT at MTD, D_{Max} : $D(t) = D_{Max} \forall t$.
2. AT as implemented in the Zhang et al [97] prostate cancer clinical trial: Treatment is withdrawn once a 50% decrease from the initial tumour size is achieved, and is reinstated if the original tumour size (N_0) is reached:

$$D(t) = \begin{cases} D_{Max}, & \text{until } N(t) < 50\%N_0 \\ 0, & \text{until } N(t) = N_0 \end{cases} \quad (3.1)$$

3.2.2 Stochastic simulation algorithm

We simulate our model using the following fixed time-step, stochastic simulation algorithm (see also Figure 3.1b):

```

while  $t \leq t_{End}$  do
  for all cells in the population do
    Select parameter set ( $k = \text{sensitive or resistant}$ )
    Let  $z_1 \sim \mathcal{U}(0, 1)$ 
    if  $z_1 \leq (r_k + d_k) dt$  then ▷ Check if an event occurs
      Let  $z_2 \sim \mathcal{U}(0, 1)$ 
      if  $z_2 \leq \frac{r_k}{r_k + d_k}$  then ▷ Check if cell divides
        if cell has space in neighbourhood then
          Let  $z_3 \sim \mathcal{U}(0, 1)$ 
          Let  $\sigma = 1$  if the cell is sensitive, otherwise  $\sigma = 0$ 
          if  $z_3 \leq \sigma d_D D$  then ▷ Drug-induced death
            Cell dies from drug
          else ▷ Cell division
            Cell divides into randomly chosen spot
            in neighbourhood
        else ▷ Check if cell dies
          Cell dies
      else
         $t \leftarrow t + dt$ 

```

Here, $\mathcal{U}(0, 1)$ denotes the uniform distribution on $[0, 1]$, and t_{End} the end time of the simulation in days. The model is implemented in Java 1.8 using the Hybrid Automata Library (HAL) [173]. Data analysis was carried out in Python 3.6, using Pandas 0.23.4, Matplotlib 2.2.3, Seaborn 0.9.0, and openCV 3.4.9. To test the algorithm’s convergence and implementation we carried out a convergence analysis. In addition, we performed a consistency analysis to determine how many stochastic replicates would be required to obtain a representative sample of outcomes [174, 175]. The details of these studies can be found in Appendix B. Based on these analyses we chose to use a time step of $dt = 1\text{d}$, and a sample size of $n = 250$ for plots in which we show treatment dynamics, and a value of $n = 1000$ for all quantitative analyses. The reason we use the smaller sample size for the time-series plots is to reduce the computational burden on the plotting routine.

3.3 Comparison of AT and CT in the ABM

The analysis of the ODE model in Chapter 2 provided us with an understanding of the basic principles underlying AT. We will now extend our understanding by testing to what extent our conclusions still hold if space and stochastic effects are taken into account. We begin our discussion by performing a similar analysis to that in Chapter 2. This shows that both models qualitatively agree in their behaviour but that, due to the stochastic nature of the ABM, it is possible in this case for AT to result in a shorter TTP than CT. Moreover, we will use the fact that in the ABM we can track individual cell-cell interactions to study competitive inhibition during AT. This allows us to gain a deeper understanding of precisely how different factors change intra-tumoral competition.

3.3.1 AT can be superior to CT despite random initial seeding of cells

In Figure 3.2 we show an example of the treatment responses observed in the ABM. We assume that there are initially $N_0 = 7,500$ cells present of which $R_0 = 7$ are resistant, corresponding to an initial density of $n_0 = 75\%$ and resistance fraction of $f_R = 0.1\%$. To compare the two treatment strategies we carry out matched simulations in which the same “tumour”, as defined by its parameters, the initial spatial distribution of cells and the random number seed for the stochastic simulations, is treated once with CT and once with AT. In order to examine the distributions of outcomes that are possible for a given initial tumour composition (n_0, f_R) we carry

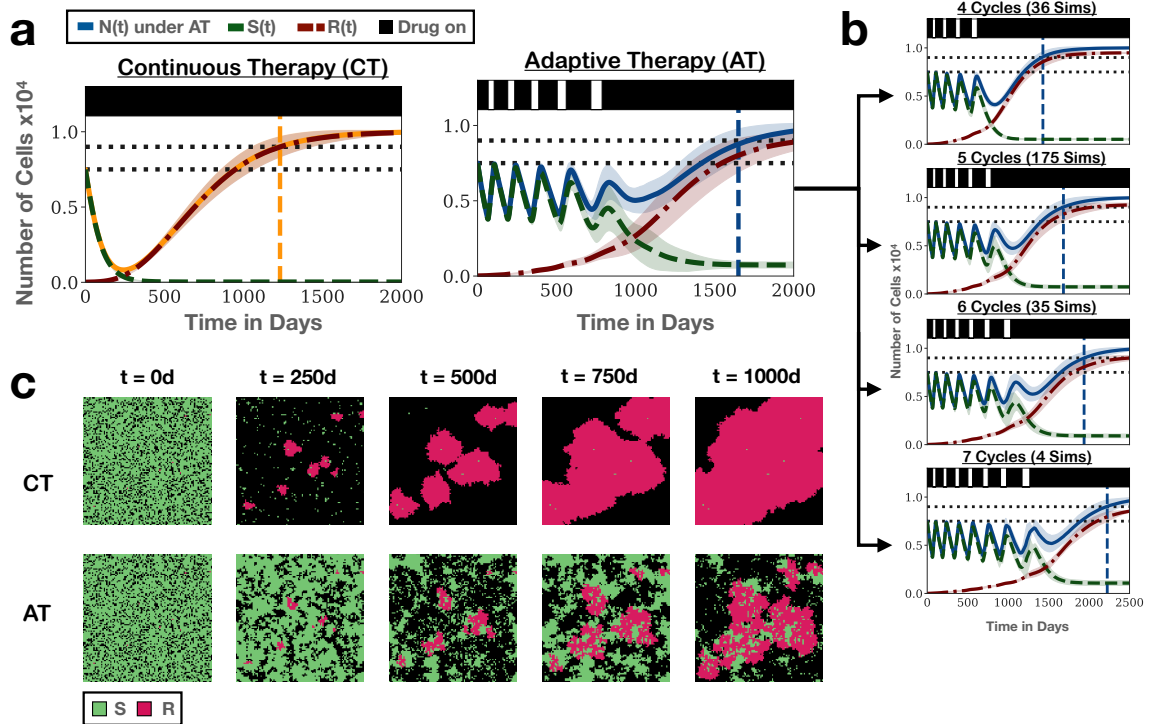


Figure 3.2: Example of the treatment dynamics under CT and AT in the ABM $((n_0, f_R, c_R, d_T) = (75\%, 0.1\%, 0\%, 0\%))$. (a) Sensitive, resistant and total cell numbers over time for CT and AT across 250 independent simulations. Shown are the mean and standard deviation (shading) of the tumour cell numbers. Black bars illustrate the treatment schedule with the most common number of cycles. Horizontal dotted lines show the initial cell number, and the cell number at progression. Vertical lines and associated shading mark the mean, and the 1st and 3rd quartile of the distribution of TTP. (b) Simulations grouped by the total number of AT cycles undergone. Despite having identical model parameters there is significant variability in the response to AT between replicates, due to differences in the initial seeding of cells and the stochastic nature of the simulations. (c) Snapshots from one of the replicates, illustrating how AT delays the competitive release of resistant tumour nests (this replicate undergoes 5 cycles).

out 250 independent replicates, each of which uses a different random number seed for the initial spatial cell configuration and for the stochastic simulation. Similar to Chapter 2, progression is defined as a 20% increase in total tumour cell number relative to the start of treatment.

In the simulations in Figure 3.2a we assume neither a resistance cost nor turnover. Moreover, cells are seeded randomly so that resistant cells are not necessarily fully surrounded by sensitive cells. Nevertheless, we find that AT extends TTP by 15.3m (95% CI: [14.81m, 15.86m]). This corroborates our result from Chapter 2 that AT can

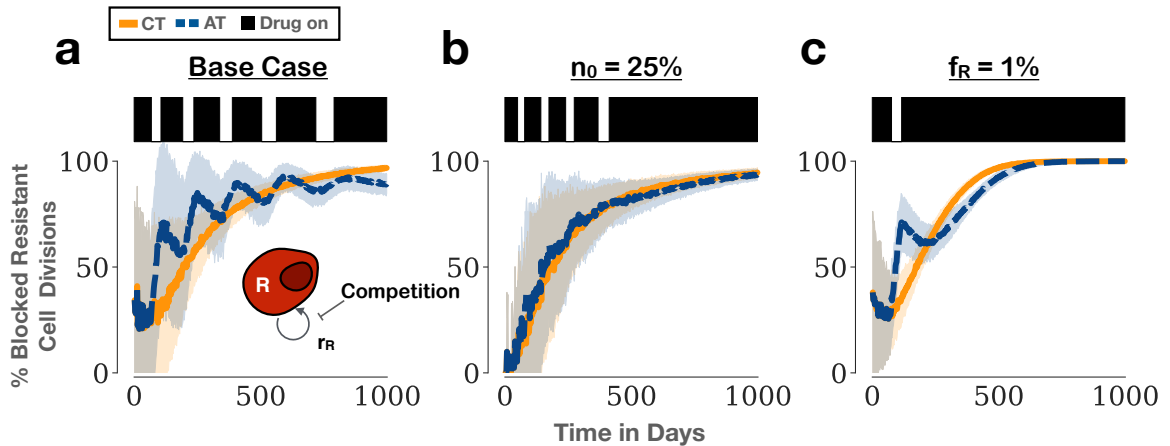


Figure 3.3: Quantification of competitive inhibition during AT. (a) Fraction of blocked resistant cell divisions as a function of time for CT and AT for the simulations in Figure 3.2. This shows how drug holidays help to slow the emergence of resistance by increasing the blocking of resistant cell divisions. Lines mark the mean value and shading the standard deviation. (b) Impact of reducing the initial proximity of the tumour to carrying capacity (n_0) on competitive inhibition ($(n_0, f_R, c_R, d_T) = (25\%, 0.1\%, 0\%, 0\%)$; 250 replicates). While there are still drug holidays, these have little effect on resistant cell growth. (c) Impact of increasing the resistant cell fraction (f_R) relative to (a) ($(n_0, f_R, c_R, d_T) = (75\%, 1\%, 0\%, 0\%)$; 250 replicates). This results in only a single drug holiday, but this holiday nevertheless decreases the resistant cell growth rate.

improve on CT under very general conditions, assuming resistance is rare and the tumour is close to its carrying capacity. That being said, we also see significant variability in the response to AT between replicates. While in some cases treatment goes through 4 cycles, in other cases we observe 7 cycles, with corresponding differences in TTP (Figure 3.2b). This indicates that there is potentially significant inter-patient variability in response to AT.

3.3.2 Quantification of competitive inhibition of resistant cells during treatment

While competition is thought to be a driving mechanism behind AT, it is challenging to quantify its role in real tumours. This is because it is difficult to rule out confounding factors such as paracrine interactions or immune predation, and because no simple biomarkers exist for measuring the strength of competition. Within the ABM we have precise control over experimental conditions and we can measure quantities which are not accessible in real life. As such, the ABM provides a unique oppor-

tunity to study competitive inhibition during AT. Firstly, in Figure 3.2c we show representative images from one of the simulations for CT and AT, respectively. We can see how CT very rapidly eliminates all sensitive cells, which subsequently allows for the expansion of the resistant nests. Conversely, we can see how AT maintains sensitive cells for longer. Importantly, while only a small proportion of sensitive cells are actually adjacent to resistant cells, the presence of these cells does visibly inhibit the expansion of the resistant colonies.

To further quantify this, we measure what proportion of attempted resistant cell divisions is blocked by the presence of another cell at different time points during treatment. During CT this fraction is initially low and gradually increases as the resistant tumour nests become larger, and cells at the centres of these nests are blocked by cells at the edge (Figure 3.2c & Figure 3.3a). While this gradual increase due to intra-specific competition also occurs during AT, we additionally observe rapid intensification of inhibition whenever treatment is withdrawn (Figure 3.2c & Figure 3.3a). This results in a level of growth inhibition which is consistently higher than that during CT until late in the course of treatment, and demonstrates how AT leverages inter-specific competition.

3.3.3 High tumour cell density and low initial resistance fraction maximise competition from sensitive cells

One of the basic results about AT, which has been described by us in the previous chapter, and others [59, 124, 126], is that the initial proximity of the tumour to carrying capacity (n_0) and the initial fraction of resistant cells (f_R) are important determinants for the success of AT. To test whether this is also the case for the current spatial model, we perform a parameter sweep over different values of n_0 and f_R . This confirms that a close proximity of the tumour to its carrying capacity and a small resistance fraction maximise the benefit of AT also in the current model (Figure 3.4 and Figure 3.5a). Moreover, we see that f_R primarily modulates the number of AT cycles and n_0 determines the benefit gained from each cycle, similar to the ODE model (Figure 3.4a). That being said, the distinction here is not as clear as in the ODE model as an increase of n_0 also increases the average cycle number (Figure 3.4a; top row).

When we examine intra-tumoral competition we find that, consistent with the competitive inhibition hypothesis of AT, large values of n_0 and small values of f_R maximise the difference in the average rate of blocking of resistant cell divisions between CT and AT (Figure 3.5b). In addition, we observe that n_0 and f_R have

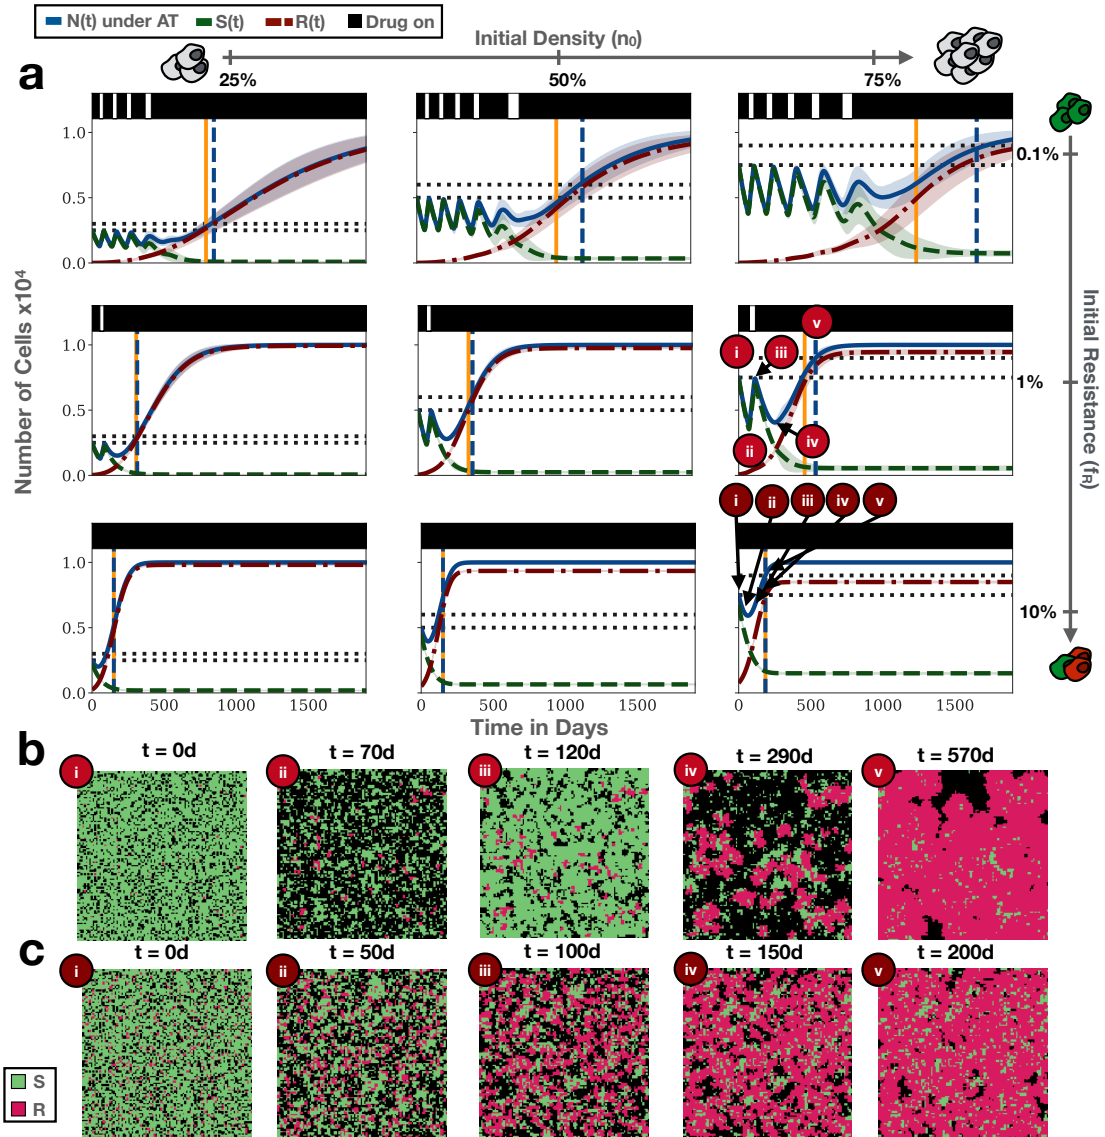


Figure 3.4: Benefit of AT over CT as a function of the initial proximity to carrying capacity, n_0 , and the initial resistance fraction, f_R . Resistance cost and turnover are assumed to be 0. (a) Example simulations at different points in n_0 - f_R space (250 replicates per parameter combination). Legend as in Figure 3.2a. (b) & (c) Snapshots from simulations in (a) showing how the dynamics changes as the resistance level is increased. Parameters: (b): $(n_0, f_R) = (75\%, 1\%)$, (c): $(n_0, f_R) = (75\%, 10\%)$. For simulations with an even smaller resistance level, see Figure 3.2c.

distinct effects on the competition experienced by resistant cells during AT. When we reduce n_0 , the tumour still cycles, but the impact of each cycle on the reduction of the resistant cell growth rate is greatly diminished (compare Figures 3.3a & b; for a plot of the cell numbers over time see Figure 3.4a). Conversely, if we increase the initial resistance fraction, treatment fails after only a single cycle (compare Figures

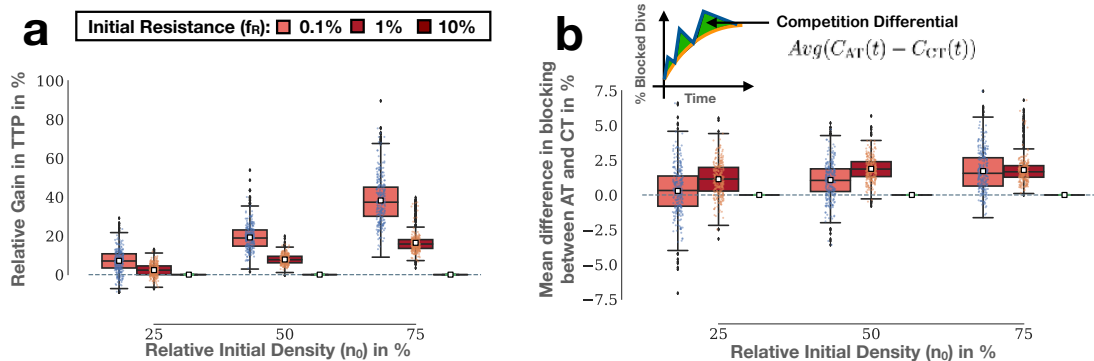


Figure 3.5: Quantification of the effect of the initial resistance fraction and proximity to carrying capacity on AT in the simulations in Figure 3.4. **(a)** TTP gained by AT relative to CT as a function of n_0 and f_R . Negative values denote cases in which CT achieves a longer TTP (1000 independent replicates per condition). This shows that also in this model the benefit of AT is maximal if tumours are close to their carrying capacity and resistance is rare. **(b)** Effect of n_0 and f_R on competitive inhibition. Shown is the mean difference in the rate at which resistant cell divisions are blocked by the two protocols in the simulations in (a), given by $1/TTP_{CT} \int_{t=0}^{t=TTP_{CT}} C_{AT(t)} - C_{CT(t)} dt$. $C_{CT(t)}$ and $C_{AT(t)}$ denote the rate of blocking under AT and CT as illustrated in Figure 3.3 (i.e., the fraction of attempted resistant cell divisions that have been unsuccessful due to the presence of another cell). This shows that high density and low resistance fractions also maximise the difference in growth inhibition due to competition between the two protocols (1000 independent replicates per condition).

3.3a & c; for a plot of the cell numbers over time see Figure 3.4a). However, this single cycle still induces blocking of resistant cells.

These results show that while both the case of small values of n_0 and the case of large values of f_R result in circumstances in which AT has little benefit over CT they do so for different reasons. If n_0 is small, then resistant cells are only mildly constrained in their growth by sensitive cells, so that maintaining the latter does not have much of an effect. In contrast, if f_R is large, then it will take less time for the resistant cells to make up the majority of the tumour and cause progression, but maintaining sensitive cells may still help to slow down this process. In particular, in this case further treatment de-escalation can help to extend control (not shown).

To sum up, this section has corroborated our result from Chapter 2 that neither resistance costs nor turnover are necessary *per se* for AT to be superior over CT. However, we see again that the tumour has to be close to its carrying capacity and resistance has to be rare. If these conditions are not met, then the benefit of AT may be small, and in some cases CT may even be superior (Figure 3.5a). This raises

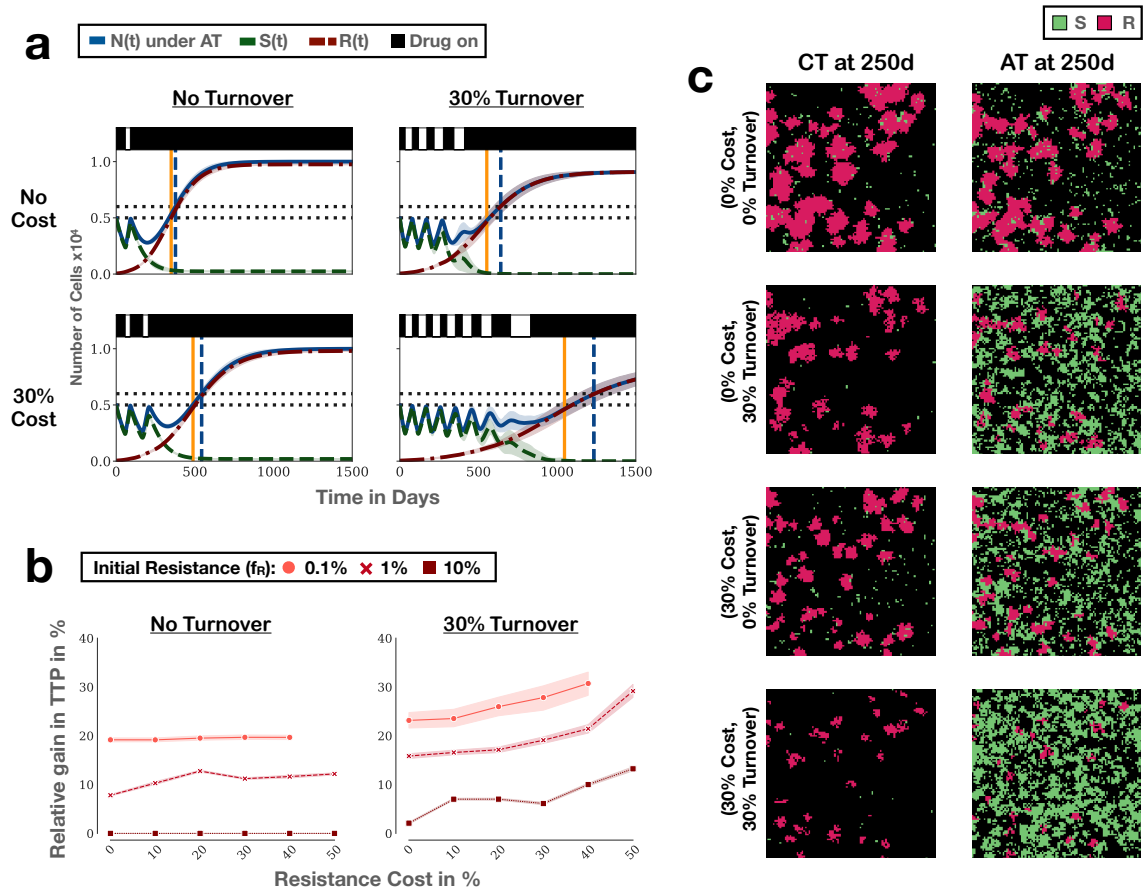


Figure 3.6: Turnover benefits AT and modulates the impact of resistant costs also in the ABM. (a) Simulations illustrating the role of resistance costs and turnover on the treatment response to AT ($(n_0, f_R) = (50\%, 1\%)$; $n = 250$ independent replicates). Legend as in Figure 3.2. (b) Relative benefit of AT over CT as a function of resistance cost, resistance fraction, and turnover ($n_0 = 50\%$). This illustrates that turnover modulates AT and increases the impact of a resistance cost. Shown are the mean (lines and points) and the 95% confidence intervals (CIs; shading) from 1000 replicates per condition. (c) Snapshots at time $t = 250\text{d}$ from one of the replicates in (a). All 8 simulations were started from the same initial conditions, and use the same random number seeds.

the question of whether in such cases a resistance cost or turnover may increase the benefit of AT.

3.3.4 Turnover also increases the benefit of AT in the ABM

A key aim of this chapter was to investigate whether the result from Chapter 2 that turnover is a facilitator of AT and modulates the impact of resistance costs also holds true if space is taken into account. To investigate this, we show a set of

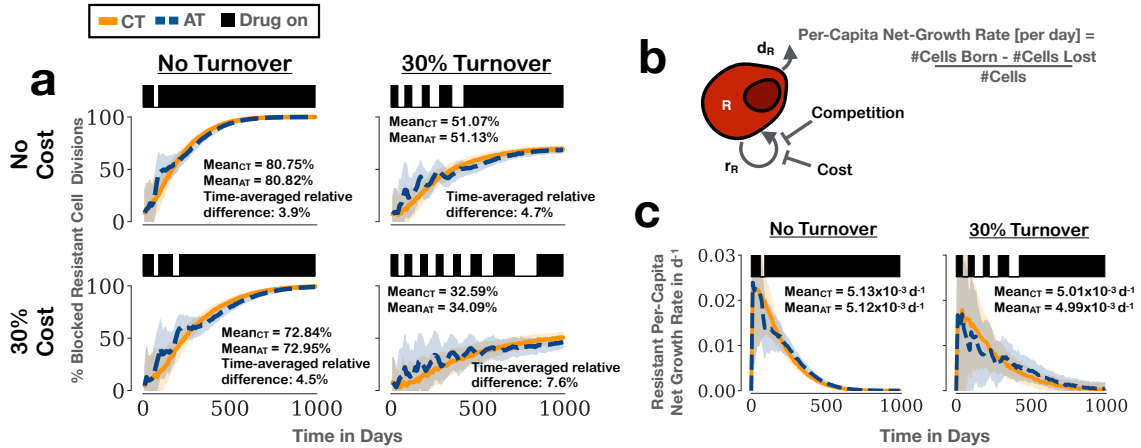


Figure 3.7: Impact of resistance costs and turnover on competitive inhibition. (a) Blocking of resistant cell divisions for the simulations in Figure 3.6a ($(n_0, f_R) = (50\%, 1\%)$; 250 independent simulations per condition). Lines denote mean values and shading the standard deviation. (b) The effect of competition on resistance controllability is modulated by turnover. If turnover is high, then even moderate inhibition of cell proliferation may result in large reductions in net-growth rate. (c) Per-capita growth rate of resistant cells as functions of time, corresponding to the simulations in Figure 3.6a. This illustrates how turnover reduces net-growth in a density independent fashion and helps to amplify the effects of competition on the resistant population’s net-growth rate.

example simulations in Figure 3.6a, in which we vary both cost and turnover. We find that, consistent with our results from Chapter 2, addition of turnover increases the benefit of AT. Furthermore, we observe again that a resistance cost slows down the progression of the tumour under both therapies but in the absence of turnover it only mildly increases the benefit of AT (Figure 3.6a). In contrast, in the presence of turnover, resistance costs significantly increase the benefit of AT. Performing a parameter sweep to map out this relationship in more detail corroborates these results (Figure 3.6b). We conclude that in the spatial model increased turnover is also predicted to result in more effective resistance control with AT. That being said, we can see in Figure 3.6b that turnover also increases the variability in outcomes, which suggests the observed benefit of AT may be a lot more, or a lot less, than expected.

3.3.5 Turnover amplifies the impact of competitive inhibition on resistant cell growth

To better understand why turnover facilitates the control of the drug-resistant population, we examine its impact on the tumour’s spatial architecture. In Figure 3.6c we show snapshots from a simulation after 250 days of treatment with CT or AT

under different conditions. We see that the addition of turnover results in fewer and smaller independent resistant cell nests under both treatments, and a larger remaining sensitive cell population at 250 days under AT. In addition, we see gaps within the nests where cells have recently died. Conversely, the addition of a resistance cost also results in smaller resistant cell colonies, but their number does not change as no cell dies. When both cost and turnover are present their effects are combined and we see the fewest and smallest colonies, and the largest remaining sensitive population under AT. This suggests that cost and turnover reduce the number of independent nests from which resistance can emerge, and slow the growth of resistant cells. Furthermore this slower growth allows for the drug to be withdrawn more often which, in turn, helps to maintain a greater number of sensitive cells for longer (compare to Figure 3.6a). This then results in additional inhibition of resistant cell growth.

Next we asked how cost and turnover affect intra-tumoral competition. To do so, we quantified the competition experienced by resistant cells in the presence, or absence, of turnover. Surprisingly, we find that both cost and turnover *decrease* the average rate of blocking of resistant cell divisions (Figure 3.7a). This is because with a cost it takes longer for cells to fill up the domain, and cell death opens up space for division, especially at the centre of colonies where previously no division was possible. That being said, when we consider the difference in blocking between AT and CT, we see that both cost and turnover increase the relative difference in competitive inhibition between the two strategies. Thus, while the absolute levels of competitive inhibition might be smaller, cost and turnover allow AT to exploit inter-specific competition for longer, which yields a greater benefit of AT.

The results in Figure 3.7a raise a question: if the absolute rate of blocking is smaller with cost or turnover, why do both increase the overall TTP? The answer to this is that progression in this model is driven by the net-growth rate of the resistant cells which depends on the birth and the death rate of the cells (Figure 3.7b). Competition for space reduces the birth rate in our model, but only once a certain density is reached. In contrast, resistance costs and cell death reduce the net-growth rate independent of density. To illustrate this, we show the resistant cell net-growth rate in the simulations in the presence and absence of turnover in Figure 3.7c. We can see that in the presence of turnover the mean net-growth under AT is consistently lower than that under CT if turnover is present. Only later in the course of treatment does it drop to comparable levels in the two treatment arms. In addition, we observe that in the presence of turnover, drug holidays result in more pronounced drops in the per-capita growth rate than in its absence. This suggests that

competitive inhibition has a greater impact on resistant cell growth in the presence of turnover, and illustrates how turnover increases the impact of competition as it reduces the number of opportunities a cell has to divide in its life time.

To sum up, cost and turnover affect AT through the interplay of the following three factors: Firstly, they reduce the net-resistant growth rate and the number of independent nests from which resistance arises, which results in slower overall resistance growth and longer control. Secondly, this slower growth allows more frequent drug withdrawal which helps to maintain more sensitive cells and allows AT to exploit inter-specific competition for longer. This yields a greater relative competitive inhibition under AT compared to CT. Thirdly, turnover amplifies the impact blocking has on resistant cell growth. Overall, these results corroborate our hypothesis that the rate of tumour cell death is an important factor and help to shed light on how it affects intra-tumoral competition.

3.3.6 AT can be inferior to CT

So far we have focussed on the benefit of AT. However, an equally, if not potentially more, important question is whether AT can result in inferior outcome compared to CT. We showed in Chapter 2 that this is not possible in the ODE model as the presence of sensitive cells will always slow progression (Section 2.3.1.4). In the “worst” case, AT will be non-superior to CT. That being said, Figure 3.5a shows that in the ABM this is not the case. Here it is possible that a lesion will progress *faster* under AT than if it had been treated with CT. In the following we will study this problem in more detail in order to understand why AT fails in these cases.

3.3.6.1 How often does AT yield inferior results?

To begin with, we quantified how frequently AT yields inferior results in our simulations. To do so, we categorised the outcome of a set of matched simulations into one of five categories. Firstly, if the tumour progresses under CT within the simulation time frame (10 years) then there are three possible outcomes: i) AT is superior to CT, ii) AT and CT progress at the same time, and iii) AT is inferior to CT. Secondly, if CT does not progress, then we can have that either iv) the tumour is cured, or v) the tumour has not progressed under either approach, so the outcome is not yet known.

In Figure 3.8a we show the frequency of these outcomes for different parameter combinations. The initial tumour composition here was chosen to be the one with the

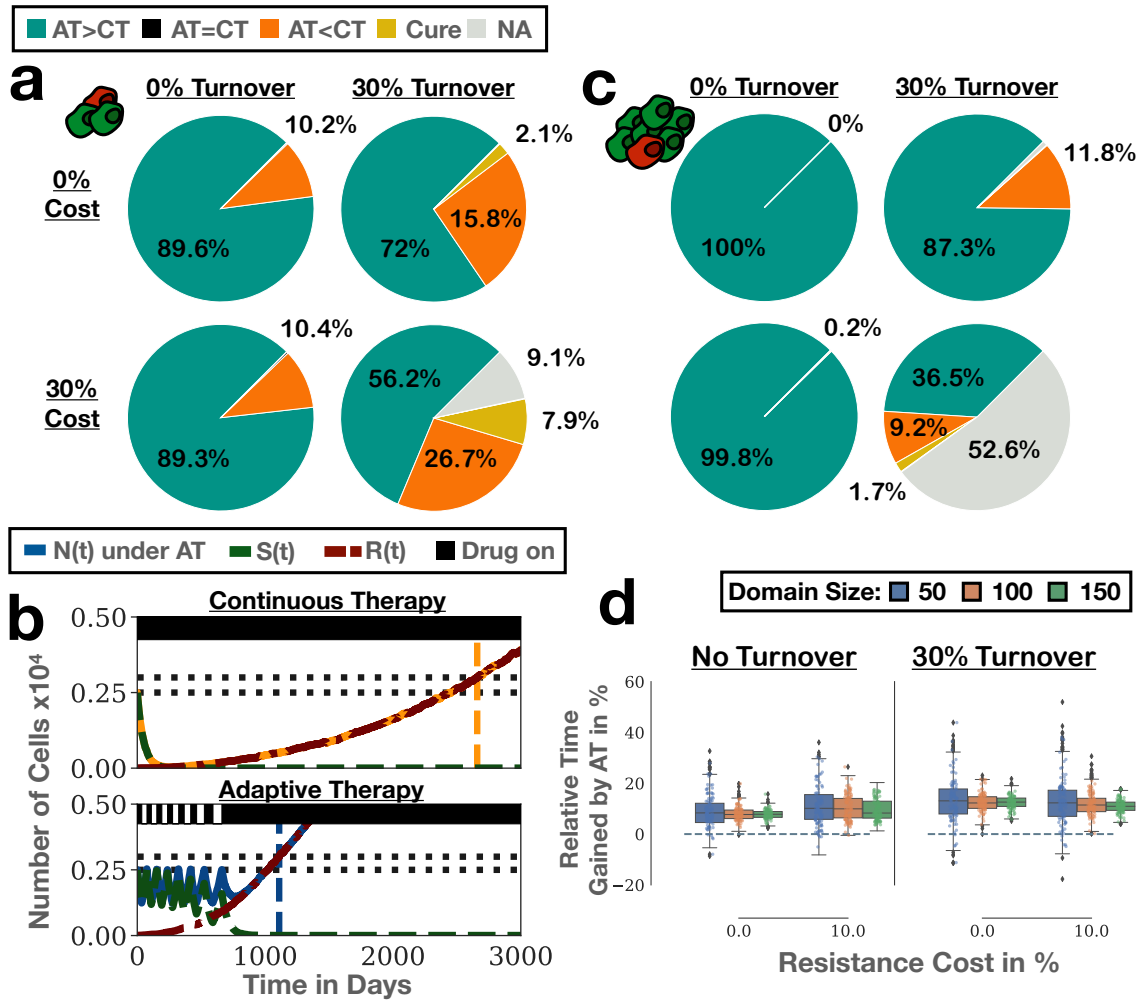


Figure 3.8: Distribution of possible outcomes in the model simulations. While cures are rare, AT can yield inferior TTP compared to CT in a non-insignificant fraction of simulations. (a) Outcome distribution for different values of cost and turnover for an initial tumour composition of $(n_0, f_R) = (25\%, 0.1\%)$ ($n = 1000$ replicates per condition). (b) Example simulation of a case in which the tumour progresses faster under AT than under CT ($(n_0, f_R, c_R, d_T) = (25\%, 0.1\%, 0\%, 30\%)$; single simulation). (c) Outcome distribution for the same resistance fraction but a higher initial density compared to (a) $(n_0, f_R) = (50\%, 0.1\%)$. With more initial resistant cells present, the probability of inferior outcomes decreases. (d) Distribution of the time gained for the tumour in (a) for different domain sizes. As the domain size increases the variation in outcomes decreases, and inferior outcomes become less frequent. This suggests that the failures of AT in this model are due to stochastic events.

largest fraction of simulations with inferior results for AT in Figure 3.8a ($(n_0, f_R) = (25\%, 0.1\%)$). We make three observations from this. Firstly, we see that, overall, cure is rare. Only if there is a considerable rate of turnover and resistance costs, do cures account for a noticeable fraction of the outcomes. Secondly, while in the majority of

simulations AT is superior to CT, there is a consistent fraction of between 10% and 27% of simulations in which it is not. Importantly, in this case the drug holidays prescribed by the adaptive protocol *shorten* TTP compared to if the patient had received CT. An example of this is shown in Figure 3.8b. Finally, we observe that the fraction of simulations in which AT is non-superior increases as we add cost and turnover (Figure 3.8a). This implies that while cost and turnover amplify the average benefit of AT over CT, they also increase the variance in outcome and the potential for detrimental results.

In addition, the likelihood of inferior outcomes also depends on the initial tumour composition. For example, in a tumour with a higher initial cell density, inferior outcomes for AT become much less likely (Figure 3.8c), and the same holds true for larger initial resistance fractions (not shown). To sum up, the probability for AT to worsen a patient's outcome is highest when the initial tumour cell density and resistance fraction are small, and cost and turnover rates are high.

3.3.6.2 The frequency of inferior outcomes decreases with the number of cells in the simulation, suggesting stochasticity causes failure

The observation that AT can result in inferior outcomes in a non-insignificant proportion of simulations is important. Even though in principle the argument from the previous chapter that sensitive cells only slow the growth of resistant cells holds true also for the current model, and AT is in fact non-inferior to CT on average, this new observation implies that there are circumstances in which AT can still result in poorer outcome. The ABM incorporates not only the effects of space but also those of stochastic events. The fact that the initial tumour composition in which AT is most likely to be non-superior is also the one with the smallest initial number of resistant cells and smallest average benefit of AT suggests that stochasticity plays a role in driving the failures. To test this we repeated the simulations in Figure 3.8a with smaller and larger domain sizes. We find that the variance in outcomes decreases as the domain size increases, suggesting that the failures of AT observed in our simulations are primarily driven by stochastic effects (Figure 3.8d). We conclude that when cell numbers are large the intuition from the ODE model about the non-inferiority of AT holds up, but if initial resistant cell numbers are small, then it is possible for AT to yield worse outcomes because of stochastic events.

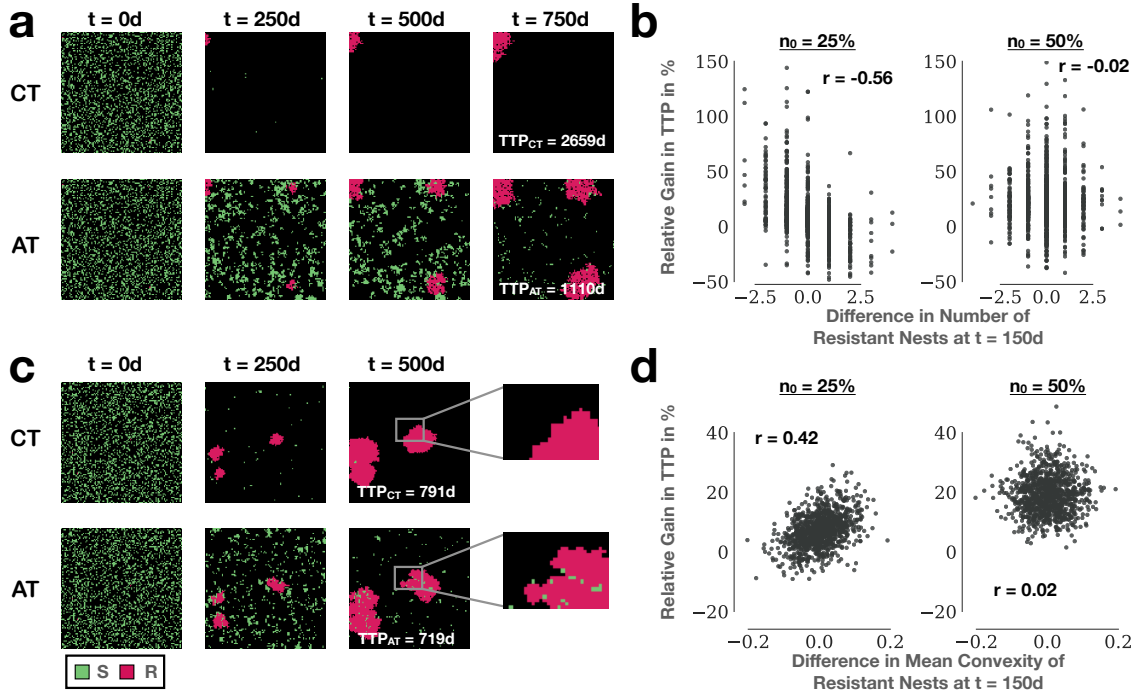


Figure 3.9: Two ways in which AT can cause faster progression than CT in the model. (a) In the presence of turnover, resistant nests can randomly go extinct in one but not the other treatment branch. As a result of the larger number of nests, AT progresses more quickly $((n_0, f_R, c_R, d_T) = (25\%, 0.1\%, 0\%, 30\%))$; corresponds to Figure 3.8b). (b) Correlation between the difference in the number of resistant nests after 150d of treatment and the gain of AT (r is Pearson's correlation coefficient). This shows that random extinctions are an important factor in explaining failure at high turnover for low, but not for high, initial cell density $((f_R, c_R, d_T) = (0.1\%, 0\%, 30\%), n = 1000$ replicates). (c) In the absence of turnover, certain failures in AT can be explained by the fact that AT induces a more branched growth morphology $((n_0, f_R, c_R, d_T) = (25\%, 0.1\%, 0\%, 0\%))$. (d) Correlation between the difference in the average resistant nest convexity after 150d of treatment and the gain of AT (convexity computed using Equation (3.2); r is Pearson's correlation coefficient). A positive difference means that resistant nests under AT are smoother (more convex) than under CT and vice versa. This shows that nest shape is a factor in explaining failure in the absence of turnover at low, but again not at high, initial cell density $((f_R, c_R, d_T) = (0.1\%, 0\%, 0\%), n = 1000$ replicates).

3.3.6.3 Differences in resistant cell nest numbers due to random extinction can cause failure of AT

The finding that stochasticity plays an important role in the observed failures raises the question of whether these failures are biologically relevant. Cell numbers in human tumours are orders of magnitude larger than those in our simulations, and our model is a highly simplified representation of the dynamics within a tumour lesion. In order

to understand whether the failures predicted by our model are plausible in practice we characterised the mechanisms by which they occurred.

We identify two main causes for inferior performance of AT. Firstly, in the presence of turnover, we find that it is possible that individual nests go extinct under CT but not AT. We show an example of this in Figure 3.9a, which corresponds to the simulation in Figure 3.8b. As can be seen, under CT all but one resistant nest go extinct early during treatment, so that progression is driven by this single nest, which is additionally constrained by being located in the corner of the domain. This explains the very slow re-growth we see in Figure 3.8b. In contrast, during AT two additional nests survive which allow for much faster overall expansion of the population and explains why AT progresses much earlier than CT in this case.

In order to quantify to what extent this difference can explain the observed failures, we counted the numbers of resistant nests after 150 days for both treatment protocols, and assessed the correlation between the difference in this number and the relative gain of AT. To do so, we used the `openCV` [176] library in Python to read-in the images from our simulations and identify the number of independent colonies via the `connectedComponentsWithStats()` function. In Figure 3.9b we show the results of this analysis for 1000 simulations with two initial tumour compositions ($n_0 = 25\%$ and 50%), corresponding to Figures 3.8a and 3.8c. For the smaller initial density value we see a clear correlation between the differences in nest numbers and outcome (Pearson’s correlation coefficient, $r = -0.56$, p-value < 0.01), showing that it is one of the factors determining outcome. However, there remains a large proportion of unexplained variance ($r^2 = 0.26$), and in the case of $n_0 = 50\%$ there is no evidence for a correlation at all (Pearson’s correlation coefficient, $r = 0.02$, p-value = 0.5). This indicates that other factors are at play as well.

3.3.6.4 AT induces a more branched growth pattern which can accelerate progression

A second factor which we identified as playing a role in the failure of AT is the tumour morphology induced by AT. We illustrate this in Figure 3.9c. Here, the initial tumour composition is the same as in Figure 3.9a but there is no turnover, so that random extinction can not be the reason for the failure of AT (note that the two simulations also have different random number seeds). Instead, we observe that the obstruction from sensitive cells induces a branching growth pattern under AT as resistant cells are forced to grow around obstructing sensitive islets (see the central resistance nest at $t = 250\text{d}$ and $t = 500\text{d}$ during AT). While this slows down resistant cell growth

initially, once the sensitive cells are killed off by drug the more branched structure allows the resistant colony to expand more rapidly as a larger fraction of cells has access to space for division.

In order to test to what degree this can explain the performance of AT we compared the average “ruggedness” of resistant nests at 150d after treatment initiation with either AT or CT. To quantify how branched a nest was we measured its convexity, given by:

$$\text{Convexity} = \frac{\text{Perimeter of Convex Hull}}{\text{Perimeter}}, \quad (3.2)$$

where the convex hull is the minimal convex region which contains the resistant nest, and convexity means that the line segment joining any two points in the region is fully contained within that region. To compute a nest’s convexity we used first the `connectedComponentsWithStats()` function in `openCV` [176] to identify individual nests, followed by the `convexHull()` function to find their convex hull. Finally, we calculated each nest’s perimeter and used Equation 3.2 to obtain the nest’s convexity.

In Figure 3.9d we show the results of this analysis which yields a positive correlation between the relative gain of AT and the difference in the mean nest convexity (Pearson’s correlation coefficient, $r = 0.42$, p-value < 0.01). This means that instances in which nests under AT are more branched than nests under CT are associated with poorer performance of AT. That being said, again this factor only explains part of the observed variance in outcomes ($r^2 = 0.17$), and it has little explanatory power at higher n_0 values (Pearson’s correlation coefficient, $r = 0.02$, p-value < 0.46 ; Figure 3.9d).

To sum up, we have found that due to the stochastic nature of our simulations it is possible that AT results in inferior outcomes compared to CT in the ABM. This occurs even though AT is beneficial on average, and even though the simulations share the same model parameters. These findings agree with those made recently by Hansen and Read [128] using a non-spatial stochastic model and suggest that there may potentially be significant inter-patient variability in AT responses. Moreover, while the highly simplified nature of our model means that the predicted frequencies of different outcomes are unlikely to be accurate, the mechanistic nature of our model allowed us to propose two plausible mechanisms by which failure may occur. In particular the concept that AT may result in a different tumour morphology than CT has, to the best of our knowledge, not been discussed before and warrants further investigation in the future.

3.4 Comparison between the ODE model and the ABM

The previous section has shown that the ABM behaves similarly to the ODE model in many ways. This is consistent with the fact that prior ODE and ABM studies of AT have yielded qualitatively similar conclusions. However, our results have also indicated that there may be differences as we observe, for example, that AT can perform worse than CT, which is not possible in the ODE model from Chapter 2. To identify the differences between the two models and develop a more detailed understanding of how space impacts AT, we will now carry out a quantitative comparison of the two models. To do so, we will compare matched simulations with the same parameters, so that both models make identical assumptions about the rate of cell proliferation, death, and drug kill. This way any discrepancies we observe will be due to difference in the interactions between cells in the two models. We will organise our discussion in the same fashion as in the previous section and will first discuss the role of the initial tumour composition, followed by a discussion of the impact of cost and turnover.

3.4.1 The impact of the initial tumour composition

3.4.1.1 The ABM predicts longer TTP but less relative benefit of AT

In Figure 3.10a we show matched simulations of the two models for different initial resistance fractions. While these qualitatively agree that AT prolongs TTP over CT, there are important quantitative differences. Firstly, we see that when f_R is small ($f_R = 0.1\%$) the TTP of both CT and AT is almost twice as long in the ABM as in the ODE model. In contrast, if f_R is large, then progression occurs more rapidly in the ABM. To further quantify this observation, in Figure 3.10b we compare the predicted TTP_{AT} for different values of f_R and for different initial densities, n_0 . This corroborates our finding that the ABM predicts faster progression when the f_R is large, and shorter TTPs when f_R is small. Furthermore, we can see that increasing the initial density increases the TTP in both models, but has a relatively stronger effect in the ODE model compared to the ABM (note how the curves bend towards the right in Figure 3.10b). For the corresponding simulations the interested reader may compare Figure 3.4a to Figure 2.4a from Chapter 2.

A second observation we make in Figure 3.10a is that even though the absolute TTP of AT may be larger in the ABM, its benefit is relatively smaller. We quantify this in Figure 3.10c, where we compare the relative time gained by AT in both models.

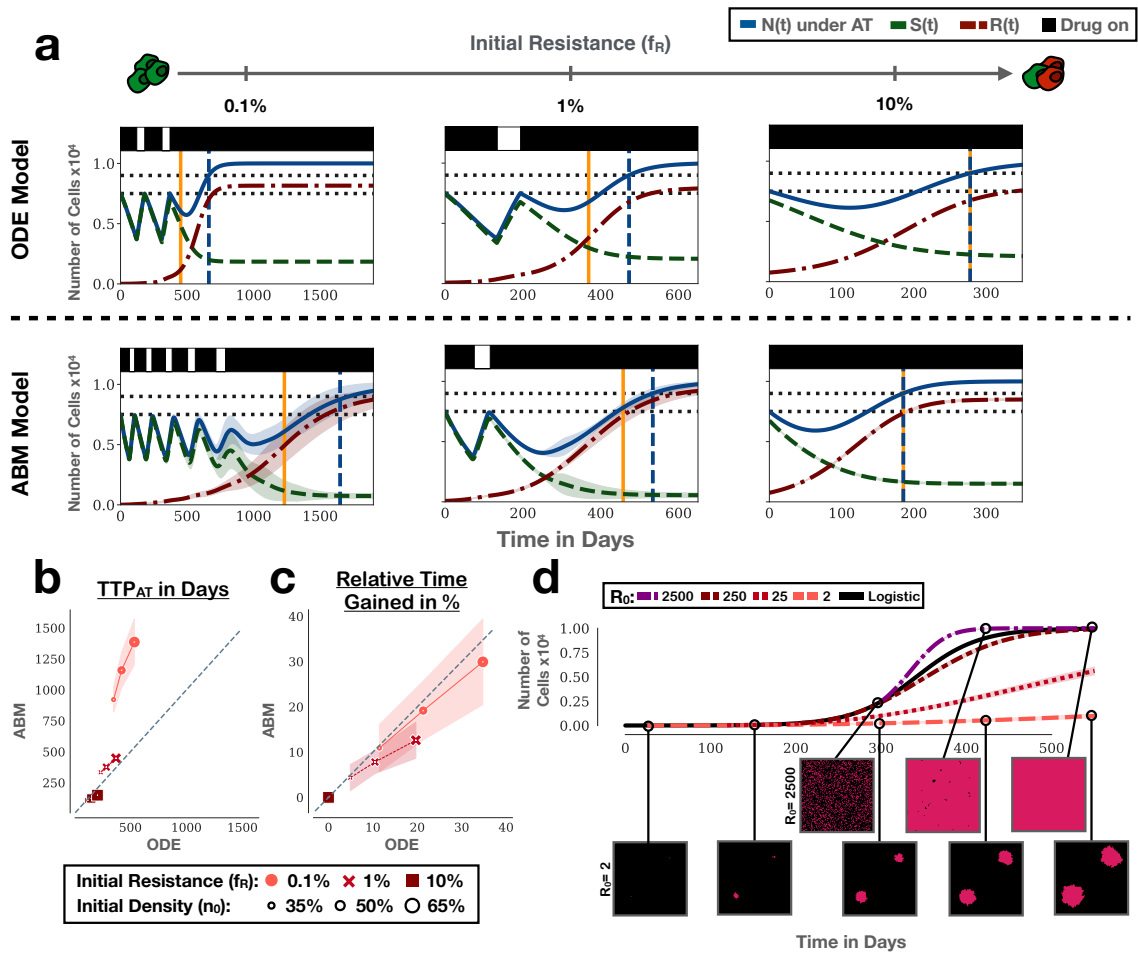


Figure 3.10: Comparison of the treatment predictions of the ODE model and the ABM for different initial tumour compositions. We are assuming no cost or turnover. Lines and shading denote the mean and standard deviation, respectively. (a) Matched ODE and ABM simulations for different initial resistance fractions ($n_0 = 75\%$; all model parameters identical). For the ABM the mean and standard deviation of $n = 250$ independent replicates are shown. (b) Comparison of the TTP under AT in the ODE model and the ABM. The ABM predicts faster progression when initial resistant cell numbers are high, and slower progression when these numbers are low ($n = 1000$ replicates of the ABM). (c) Comparison of the relative time gained by AT ($(TTP_{AT} - TTP_{CT})/TTP_{CT}$) predicted by the two models. The ABM tends to predict a smaller benefit than the ODE model. (d) An important factor causing the differences between the models is the resistant population growth kinetics. Depending on the initial cell number (R_0) the resistant population will grow sub- or super-logistically (logistic growth: black solid line). Resistant cells were seeded and grown in isolation and without drug ($n = 250$ replicates). To allow direct comparison with logistic growth, the agent-based curves were shifted to start at the time at which the logistic model reached the corresponding starting number of cells.

This confirms that on average the relative benefit is smaller in the ABM, and that this discrepancy increases the closer the tumour is initially to its carrying capacity.

3.4.1.2 The ABM predicts faster cycling dynamics than the ODE model

A third difference we identify between the two models is in the predicted speed of cycling. This is best illustrated by the case $f_R = 0.1\%$ where we observe that in the time in which the ODE model predicts completion of two cycles, the ABM predicts completion of almost four cycles (Figure 3.10a). To understand whether these differences are due to difference in the on-drug or off-drug dynamics, or both, we consider the case $f_R = 1\%$ in Figure 3.10a. While both models predict completion of only a single cycle, we can see that in the ABM both the on-treatment and off-treatment times are shorter than in the ODE model. In addition, we can see that the final composition of the tumour in the two models is different, with the ODE model predicting a higher remaining fraction of sensitive cells than the ABM.

3.4.1.3 The number of independent starting colonies determines the resistant population's growth kinetics

How can we explain the observed differences? Similar to the ODE model, progression in the ABM is caused by the resistant population, as drug-sensitive cells can always be removed by drug. So, to understand why the progression dynamics differs between the two models we compare the growth kinetics of the resistant population under the three conditions in Figure 3.10a. We observe that when f_R is small the resistant population in the ABM grows much more slowly than in the ODE model, while the converse holds when the initial resistance fraction is large ($f_R = 10\%$).

To explain these differences in growth kinetics we examine the spatial architecture of the tumour in the ABM. We observe that when the initial resistance fraction is low ($f_R = 0.1\%$), the resistant population emerges from only a handful of initial nests (Figure 3.2c). In contrast, when the resistant fraction is large ($f_R = 10\%$), progression is driven by many small, independent colonies (Figure 3.6c). This suggests that it is the initial number of independent resistant cell nests that causes the differences between the two models. When there are initially few resistant cells present, the overall growth dynamics is slow as only the resistant cells at the rims of the expanding colonies have access to space for division. In contrast, if there are many independent resistant nests, then growth is much faster as many more cells have the potential for division (large surface-to-volume ratio).

To test this hypothesis we compared the growth kinetics resulting from four different starting populations (Figure 3.10c). For simplicity, we consider just the resistant population in this case and do not include sensitive cells or the drug in these simulations. Our results show that when initiated from two cells, the resistant population grows much more slowly than predicted by logistic growth. As the number of cells, and so the number of independent nests and the surface to volume ratio, is increased, the growth of the population speeds up until it exceeds that of logistic growth (Figure 3.10c). This explains why progression occurs sooner in the ABM when the initial resistance fraction is large, but later when it is small. It also explains why the predicted benefit of closer proximity to carrying capacity is greater in the ODE model: while it helps to increase competitive inhibition of resistant cells, a larger value of n_0 also implies the presence of more independent resistant nests which allows for faster growth.

3.4.1.4 Differences in the predicted sensitive cell growth fraction explain discrepancies in predicted cycling behaviour

A similar mechanism also explains the differences in the cycling dynamics, and the final tumour composition. In both models, drug only kills dividing sensitive cells. Because of the random seeding this fraction is again higher in the ABM than in the ODE model. For example, consider the starting conditions shown in Figure 3.2c, when $n_0 = 75\%$. In the ODE model it is assumed that only $1 - \frac{S(t)+R(t)}{K} = 25\%$ of cells have an opportunity to divide. However, in the initial snapshot shown in Figure 3.2c 69% of cells have space available for division. This explains why not only the drug response, but also subsequent re-growth, is faster in the ABM, and why more sensitive cells are killed overall.

3.4.2 The impact of resistance costs and turnover

3.4.2.1 The ODE model predicts a larger benefit from resistance costs and turnover than the ABM

Having studied how the predictions of the two models differ depending on the initial tumour composition, we next turn to consider the roles of cost and turnover. In Figure 3.11a we show a series of matched simulations in which we gradually add first a resistance cost and subsequently turnover. In the absence of either factor, both models predict a single AT cycle and a small benefit in TTP for AT. In addition, consistent with our previous observations, the ABM predicts a longer overall TTP than the ODE model. When we add a cost but assume that turnover is zero both models

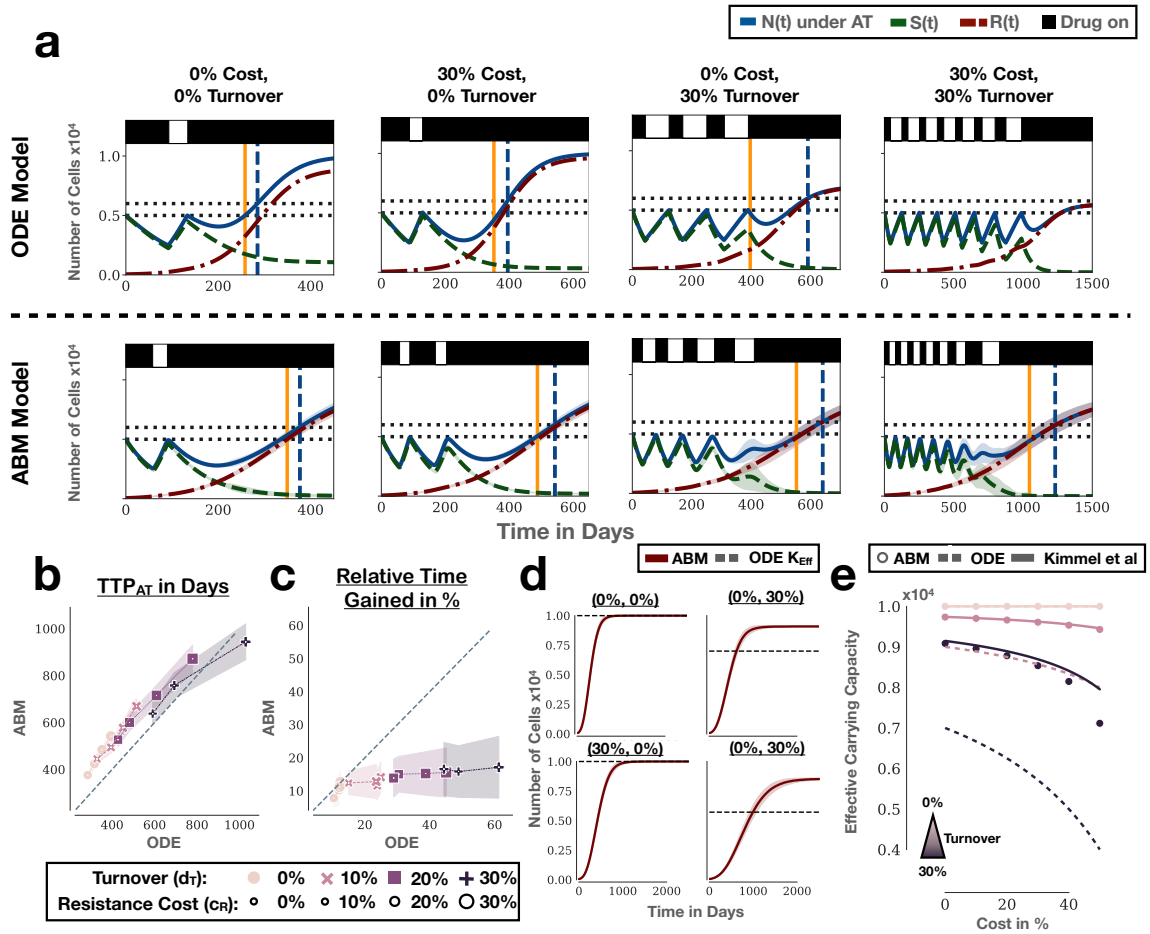


Figure 3.11: Comparison of the treatment predictions of the ODE and the ABM for different values of cost and turnover. Lines and shading denote the mean and standard deviation, respectively. (a) Matched simulations for different combinations of cost and turnover values ($(n_0, f_R) = (50\%, 1\%)$). For the ABM the mean and standard deviation of $n = 250$ independent replicates are shown. (b) Comparison of the TTP under AT in the ODE model and the ABM. The ABM predicts, in general, later progression but this difference becomes smaller in the presence of cost and turnover ($n = 1000$ replicates of the ABM). (c) Comparison of the relative time gained by AT ($(TTP_{\text{AT}} - TTP_{\text{CT}}) / TTP_{\text{CT}}$) predicted by the two models. The ABM predicts a significantly smaller benefit for AT from cost and turnover than the ODE model. (d) The two models predict different steady states of the resistant population (“effective carrying capacities”) in the presence of turnover. Shown are simulations of resistant cells grown in isolation (no drug) under different conditions ($n = 250$ replicates). The dashed line gives the steady state expected in the ODE model (Equation (3.3)). (e) Effective carrying capacity as a function of cost and turnover in the ABM model (points; colour reflects turnover) and ODE model (dashed line; Equation (3.3)). Values for the ABM were obtained by taking the final population size after 3650 days as in (d) ($n = 1000$ replicates per condition). The solid line shows the approximation for $K_{\text{Eff}}^{\text{ABM}}$ derived by Kimmel et al (Equation (3.4); in prep).

predict an extension of TTP for both CT and AT, but little extra benefit for AT. Quantifying these effects in more detail we find that both models predict relatively similar extensions to TTP_{AT} (the 0% turnover line (pink circles) runs almost parallel to the 45 degree line in Figure 3.11b), and agree closely on the predicted relative gain of AT (Figure 3.11c).

In contrast, there are significant discrepancies in the two models' predictions on the impact of turnover (Figure 3.11a). Firstly, while higher turnover results in later progression in both models, the impact is predicted to be larger in the ODE model. This can be seen in the examples in Figure 3.11a, where addition of both turnover and cost yields indefinite control in the ODE model but not in the ABM. Quantification of these observations for a wider range of parameters confirms that turnover extends TTP_{AT} more in the ODE model than in the ABM (curves bending to the right in Figure 3.11b). Moreover, this shows that turnover has a significantly greater predicted benefit for AT in the ODE model than in the ABM, and that this discrepancy increases if, additionally, a resistance cost is present (Figure 3.11c).

3.4.2.2 Weaker relationships between cost, turnover and effective carrying capacity explain the discrepancies

Overall, these observations confirm our hypothesis that the benefit of turnover for AT will be smaller in a spatial model than in the ODE model. While turnover reduces the net-growth rate it also makes it easier for resistant cells to escape trapping in the spatial model as it opens up space into which resistant cells can divide (see also the discussion in Section 3.3.5). But why do also resistance costs have so little impact (Figure 3.11c)? When we compare the ODE model and ABM simulations for the case $(c_r, d_T) = (30\%, 30\%)$ in Figure 3.11a we see that the two models predict different steady states. One of the insights we gained in Chapter 2 was that cost and turnover reduce the “effective carrying capacity” of the resistant population (Section 2.3.2.5). This was important because it meant the resistant population was effectively closer to carrying capacity than would be suggested by the available space/resources, which was why cost and turnover increased the benefit of AT. Thus, the simulations in Figure 3.11a indicate that the reason why the benefit of cost and turnover is smaller in the ABM is because they have less impact on the effective carrying capacity.

In order to test this hypothesis, we carried out simulations of just the resistant population for different cost and turnover values, and compared the steady state seen

in the ABM, $K_{\text{Eff}}^{\text{ABM}}$, with that in the ODE model, $K_{\text{Eff}}^{\text{ODE}}$. The latter is given by:

$$K_{\text{Eff}}^{\text{ODE}} = \left(1 - \frac{d_T}{r_R}\right) l^2. \quad (3.3)$$

This analysis shows that in the absence of turnover both the ABM and the ODE model predict that the population will saturate at the environmental carrying capacity, $K = l \times l$, also when a cost is present (Figure 3.11d). However, if turnover is introduced then $K_{\text{Eff}}^{\text{ABM}}$ is significantly larger than $K_{\text{Eff}}^{\text{ODE}}$, and this discrepancy increases further when we add a cost (Figure 3.11d). To quantify this, in Figure 3.11e we show the resistant population's steady state as a function of resistance cost for different values of turnover. This demonstrates that while both $K_{\text{Eff}}^{\text{ABM}}$ and $K_{\text{Eff}}^{\text{ODE}}$ have a similar relationship with cost, the impact of turnover on $K_{\text{Eff}}^{\text{ABM}}$ is much larger than that on $K_{\text{Eff}}^{\text{ODE}}$ (Figure 3.11e).

Working on a different project, two colleagues at the Moffitt Cancer Center (Dr Gregory Kimmel and Dr Jeffrey West) have derived expressions for the steady state in the ABM. Using a master equation approach they have shown that:

$$K_{\text{Eff}}^{\text{ABM}} \approx \left(1 - \frac{d_T}{r_R}\right)^{\frac{1}{a}} l^2, \quad (3.4)$$

where a denotes the number of neighbourhood sites a cell can potentially access for division ($a = 4$ in the case of the von Neumann neighbourhood used here). A publication with the details of their work is in preparation². Figure 3.11e shows that this expression approximates the observed relationship very well (solid lines). Based on this result we propose the following explanation for why the benefit of cost and turnover is reduced in the ABM: Because a cell has four potential sites into which it can divide, there is more opportunity for division in the ABM. As such, the impact of both cost and turnover is reduced relative to the ODE model because the higher cell production partially compensates for the cells lost due cost or turnover.

3.4.2.3 The ABM predicts faster cycling than the ODE model also in the presence of cost or turnover

A final observation we can make in the simulations in Figure 3.11a is that the differences in the cycle dynamics between the two models persist. Specifically, we can see that turnover changes the ratio of on- to off-treatment times in both models, but the frequency of cycling remains faster in the ABM. Again, this difference can be explained by the larger fraction of dividing sensitive cells in the ABM.

²Permission to discuss this result here was granted by Dr Kimmel.

3.4.2.4 Summary

To sum up, we have found stark quantitative differences in the predictions between the two models. Importantly, while both models agree that AT will provide benefits over CT, the ODE model predicts a larger relative benefit than the ABM. This suggests that the benefit of AT in practice may be smaller than expected from models based on logistic growth. Moreover, we have shown that these differences are due to the fact that the two models make different predictions about the fractions of dividing cells in the tumour, and that additionally in the ABM this fraction is dependent on the tumour's spatial architecture. In particular, we showed that the initial number of independent sites from which resistance emerges is a key determinant of the progression dynamics. This is an important observation as it suggests that resistance which is derived from mutations in a single cell will have a different response to AT than resistance which arises through phenotypic adaptation and so may occur in multiple locations at once. We will investigate this idea further in the following section.

3.5 Impact of the initial resistant cell distribution

Most previous work has focussed on analysing AT in a spherical model of tumour growth [115, 124]. Importantly, this assumes that resistance arises only in a single location within the tumour. However, for most drugs more than one resistance mechanism is known [46], and in light of the emerging role of phenotypic plasticity and environmental factors, there is increasing evidence that resistance may arise in multiple locations in parallel [47]. Evidence for such poly-clonal resistance has been observed, for example, in colorectal cancer [10], lung cancer [177], or melanoma [178]. We have already seen that the number of independent sites changes the tumour's growth and cycling kinetics, so that controlling multiple sites will likely represent a challenge for AT. In the following we will therefore perform a more systematic analysis of this problem, which will provide us with one of the key insights from this chapter: namely, that AT isn't just about competition of resistant with sensitive cells, but that competition of resistant cells with each other also plays an important role.

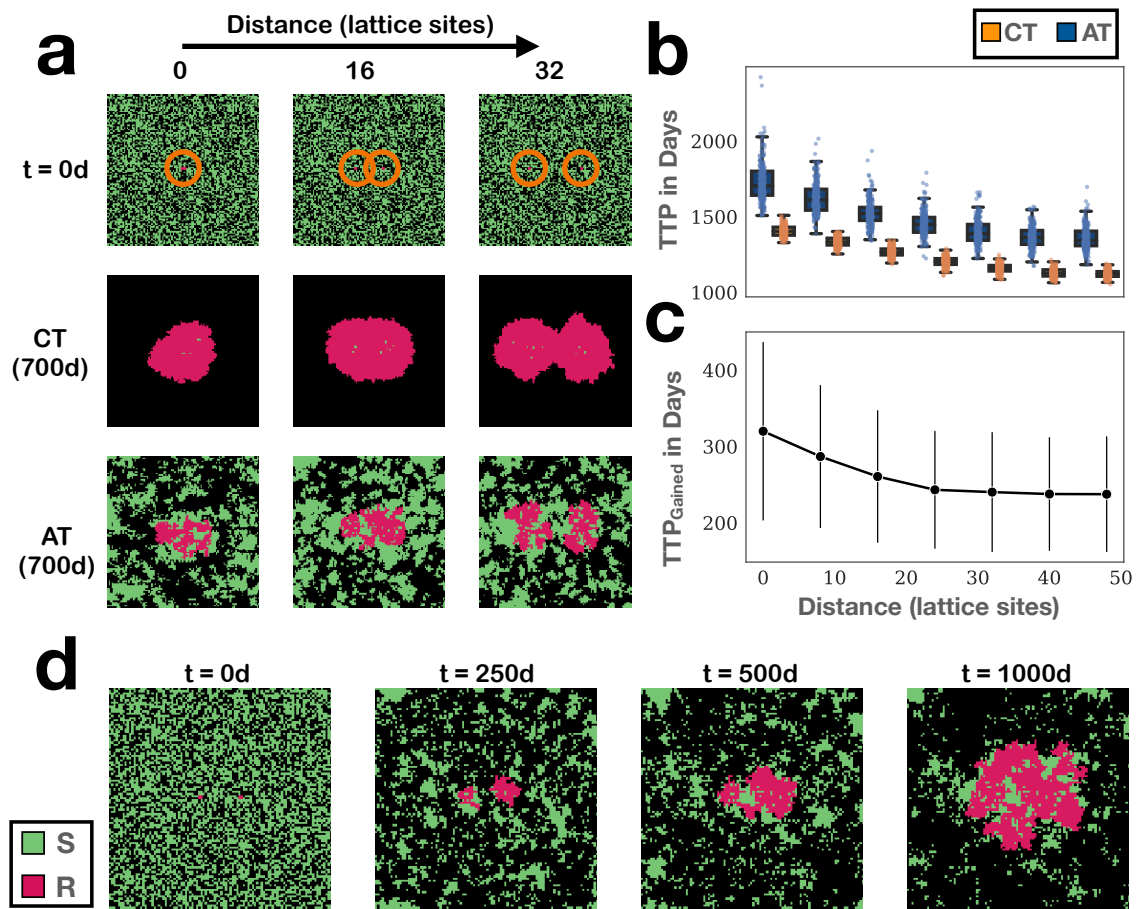


Figure 3.12: As a result of parallel mutation events or because of environmental factors mediating resistance, resistance can arise in multiple locations at once. This poses a number of challenges for resistance management. Simulations assume neither cost nor turnover ($(c_R, d_T) = (0\%, 0\%)$). **(a)** To investigate the impact of the initial spatial distribution of resistance on treatment response we seeded eight resistant cells either as a cluster in the centre, or as two nests at varying distances apart. Shown are representative snapshots from simulations for three different initial levels of separation. **(b)** TTP of CT and AT decreases as the separation between the nests increases ($n = 1000$ replicates per conditions). **(c)** Similarly, the benefit of AT decreases the more separated the nests ($n = 1000$ replicates per conditions). Error bars mark one standard deviation. **(d)** Representative snapshots from a simulation in which the nests are initially 16 lattice sites apart, which illustrate one of the reasons why control of multiple nests is more challenging. While the left nest can initially be controlled, the right nest escapes, and in turn triggers release of the left nest. This shows that the distribution of resistance, in particular the presence of resistance in multiple locations, represents a challenge for resistance management.

3.5.1 Mono-clonal resistance is more effectively controllable because it allows AT to leverage both inter- and intra-specific competition

We conducted a series of computational experiments in which we seeded the same number of resistant cells either as a single cluster or as a set of two clusters (Figure 3.12a; top row). To simulate the single cluster we placed eight resistant cells as a 2×4 rectangle in the centre of the domain. Subsequently, we compared this scenario to those in which we split this cluster into two nests of size 2×2 which we placed at varying distances apart from each other. Sensitive cells were seeded randomly in the domain to achieve a total initial density of $n_0 = 50\%$. An example of the initial conditions for three different levels of separation is shown in Figure 3.12a. We repeated this experiment with four instead of two nests, and with different values of n_0 , both of which gave qualitatively similar results. Thus, for brevity only the results for two nests and a value of $n_0 = 50\%$ will be discussed here.

In Figure 3.12a we illustrate the tumour growth patterns under CT and under AT for three different levels of separation between the nests. For simplicity we are assuming the presence of neither a cost nor turnover. Our results show very clearly that the speed of progression is determined, not only by the initial number of resistant cells, but also by their distribution. For example, even though all three simulations started from eight resistant cells, the further apart the two nests, the more resistant cells are present at 700 days. Computing TTP confirms this (Figure 3.12b).

Moreover, with increasing separation also the benefit of AT declines (Figure 3.12c). This is not unexpected, but has important implications. Based on the hypothesis that AT is driven by inter-specific competition between sensitive and resistant cells, we would have expected the opposite to hold true, as placing the nests apart maximises the opportunity for interaction between the cell populations. As such, these new findings indicate that not only inter- but also intra-specific competition plays a role in AT. This is because competition with sensitive cells is, in a sense, a double-edged sword. We can control the sensitive cells with treatment, and in the absence of treatment they may even have a competitive advantage over resistant cells. However, in the presence of treatment this advantage is lost, and by clearing sensitive cells we, in fact, open up space which allows resistant cells to expand. In contrast, when resistant cells grow adjacent to other resistant cells, any successful division is a zero-sum game, as it simply replaces a resistant cell which was previously in this position. As such, while there is more competition with sensitive cells when the resistant nests

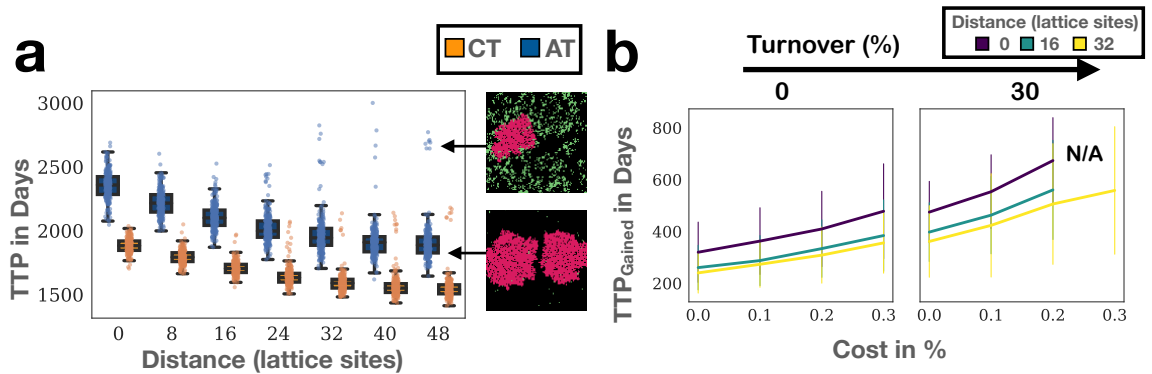


Figure 3.13: The spatial distribution of resistance modulates the impact of cost and turnover. Experimental setup as in Figure 3.13. (a) TTP as function of separation between nests in the presence of turnover ($d_T = 30\%$; compare to Figure 3.13b). Turnover can cause extinction of one of the two nests, which significantly increases TTP (1000 replicates per condition). Insets show illustrative snapshots at $t = 1500$ d. (b) Effect of the initial spatial distribution of resistance on the relationship between cost, turnover and gain of AT ($n = 1000$ replicates per condition). Error bars show one standard deviation.

are seeded apart, better control is achieved when they are clustered together because it allows AT to leverage both inter- as well as intra-specific competition.

Another reason for why AT is less effective when there are two separate nests, rather than one, is illustrated in Figure 3.12d. While the left of the two nests is initially constrained by sensitive cells ($t = 250$ d), the right nest is not and is able to expand ($t = 500$ d). This triggers more and more treatment which eventually results in the competitive release also of the left nest ($t = 1000$ d). We conclude that the spatial distribution of resistance plays a significant role in the response to AT, and its potential benefit.

3.5.2 The spatial distribution of resistance modulates the impact of cost and turnover

Next, we repeated the previous experiment in the presence of cell turnover. This shows a similar decrease in the mean TTP, and the benefit, of AT with increasing separation of the nests as seen in the absence of turnover (Figure 3.13a). That being said, there is a subtle, yet important, difference: turnover can cause the random extinction of one of the two nests, which significantly extends TTP and results in a high variability in outcomes (insets in Figure 3.13a). The impact of this is largest, the greater the separation between nests. Conversely, when we consider the relationship

between resistance costs and gain from AT, we find that this is adversely affected by the separation between the nests. The further the nests are apart, the smaller the benefit from resistance costs, even in the presence of turnover (Figure 3.13b).

Finally, we note that AT and CT again result in different growth patterns, similar to what we had reported in Section 3.3.6.4. While under CT the resistant population expands in circular fashion, under AT the obstruction from sensitive cells induces a branching structure (Figure 3.12a). That being said, as more and more treatment is administered and sensitive cells are cleared, this branching structure is lost, and this is also true if turnover is increased (not shown). Nevertheless, this supports the notion that AT may alter the tumour's spatial structure.

We conclude that resistance arising in multiple locations within a patient poses a significant challenge to resistance management. This raises an important, but so far overlooked, point: not only inter-specific competition between resistant and sensitive cell determines the outcome of AT, but also the intra-specific competition between resistant cells plays a critical role. The better control is achieved not by maximising contact between sensitive and resistant cells (large separation of resistant nests) but by blocking growth of the few resistant cells on the edge of a nest, and leveraging intra-specific competition to control the rest in its centre.

3.6 Further investigation of patient responses to intermittent androgen-deprivation treatment

Through the analyses in this chapter we have developed an in-depth understanding of intra-tumoral competition. This has revealed that the spatial structure of the tumour, in particular how resistant cells are distributed in space, is a key determinant of the benefit of AT. Moreover, we have found that this distribution modulates the impact of resistance costs and turnover and is, in turn, itself shaped by these factors. In light of our novel insights, the aim of this final section is to refine our understanding of the treatment dynamics of intermittently treated prostate cancer patients. In Chapter 2 (Section 2.5), we studied longitudinal PSA data from 67 patients from a Phase II study of intermittent androgen-deprivation therapy by Bruchovsky et al [109]. The trialled algorithm mimicked AT in that it administered treatment until the levels of PSA, a blood-based biomarker used for tracking tumour burden, were reduced below some threshold. Subsequently, treatment was withdrawn until PSA levels again exceeded some upper limit, when treatment was re-instated (recall Figure 2.13). Interestingly, we observed that some patients were cycling rapidly under treatment whereas others

were cycling more slowly. By fitting the ODE model from Chapter 2, we found that we were able to explain the cycling dynamics by using different combinations of the level of resistance cost and the rate of tumour turnover, suggestive of different underlying disease biology. In light of our new insights, we hypothesise that these differences may manifest themselves not only on the cell scale, but also as differences in the spatial structure of the tumours. In order to test this hypothesis, and gain further insight into how fast and slow cycling patients may differ, we fitted the ABM to these same data. In the following, we will first describe our fitting protocol, and subsequently discuss our results.

3.6.1 Model fitting

In order to examine whether the ABM can explain differences in the cycling speed of patients undergoing intermittent androgen deprivation therapy, we fit the ABM to the data by Bruchovsky et al [109] using a similar protocol to that employed in Chapter 2. The data are the same as those in Chapter 2, and we refer to Section 2.5.1 for further details. So as to avoid potentially confounding effects from a change in the number of lesions, patients who developed a metastasis were again excluded from this analysis. This yielded data from a total of 67 patients. The ABM was fitted to the normalised PSA measurements of each patient by minimising the RMSE between the data and the predictions. The predictions were based on the total cell number predicted by the ABM, $N(t)$, normalised relative to the initial cell count, N_0 . Given the stochastic nature of the ABM, each prediction was assembled from the mean of 25 independent stochastic replicates.

Initially we attempted to carry out optimisation using the same Levenberg-Marquardt algorithm from Chapter 2. However, we found that the stochastic nature of the optimisation problem caused problems in finding appropriate step sizes, and often resulted in premature termination (not shown). As a workaround, we employed the basin-hopping algorithm in Scipy [179] using default search parameters and a maximum of either 50 (when fitting 2 parameters) or 75 optimisation steps (when fitting 4 parameters). To escape potential local minima, optimisation was repeated 10 times for each patient from different, randomly chosen initial conditions, with only the best fit according to the AIC taken forward for analysis. Fitting was done using the `lmfit` package in Python [162].

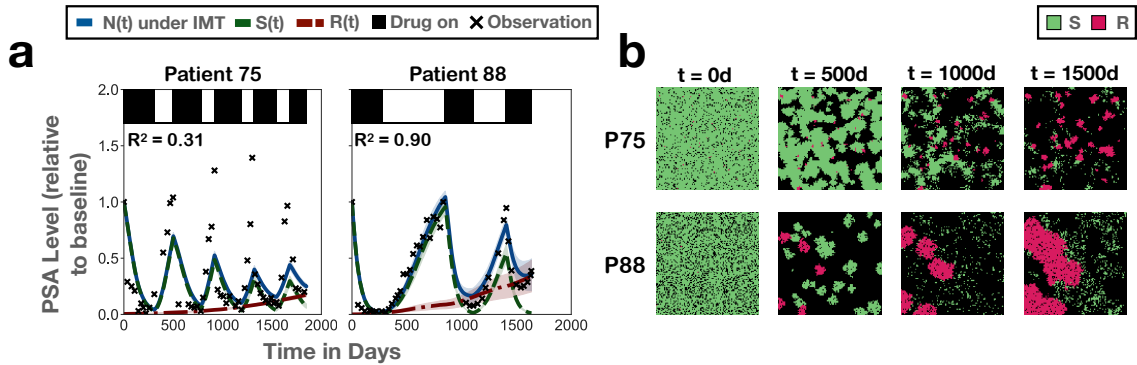


Figure 3.14: Representative fits of the ABM to the longitudinal PSA data of two prostate cancer patients undergoing intermittent androgen deprivation therapy from the trial by Bruchovsky et al [109]. (a) Tumour burden dynamics over time, as measured by PSA, and ABM fits. Solid line and shading mark the mean and standard deviation of 250 replicates. (b) Snapshots from the simulations in (a), showing distinct patterns of resistance growth in the two patients.

3.6.2 The ABM recapitulates fast and slow cycling patients

Our first question was whether the ABM could fit the dynamics observed in the patients. To test this, we fitted the model allowing n_0 , f_R , c_R , and d_T to be patient-specific parameters, as we had done for the ODE in Chapter 2. All other parameters were set to the values shown in Table 3.1. In Figure 3.14a we show example fits for two patients, one (Patient 75) who is cycling quickly, and one (Patient 88) who is cycling slowly. We find that in both cases the ABM can capture the cycling dynamics of the patients, although it provides a more accurate description of the slower cyler. In Patient 75, the model struggles to recapitulate the heights of the peaks in burden at the end of off-treatment intervals, an observation that also holds true for the other fast cyclers in the cohort. An overview of all fits is shown in Figure 3.15.

Having established that the model can fit the data, we next considered the spatial dynamics predicted to take place within each of the two patients. We find that fast and slow cyclers are associated with distinct spatial dynamics in our simulations (Figure 3.14b). In fast cycling patients, we predict that the tumour has a “carpet-like” structure with many small, independent, sensitive and resistant nests (Patient 75). In contrast, in slowly cycling patients growth is driven by only a handful of large, “patch-like”, colonies (Patient 88). This supports our hypothesis that not only variations at the cell-scale (e.g. resistance costs, turnover), but also at the tissue-scale, may contribute to the different behaviours of these tumours.




Figure 3.15: Overview of the ABM fits (fitting n_0 , f_R , c_R , and d_T) for all 67 patients, arranged by their r^2 value (showing the mean and standard deviation of 25 replicates per patient). Title colour indicates whether a patient relapses (orange) or not (green). Patients who were excluded from further analysis due to poor model fits are marked with a grey background.

3.6.3 Inter-patient variability is driven by differences in the resistance cost and turnover

In order to understand what was driving these differences, we explored two possible explanations: i) The differences are due to the initial conditions (n_0 and f_R), or ii) the behaviour is driven by differences in the cell kinetic parameters (c_R and d_T). To test these hypotheses, we fitted the ABM while keeping either the initial conditions (n_0 and f_R) or the cell kinetic parameters (c_R and d_T) fixed across the cohort. In each case we set the two parameters which we fixed to the mean of the values we had obtained when fitting all four parameters ($n_0 = 0.59$, $f_R = 0.04$, $c_R = 0.78$, and $d_T = 0.14$). Note that when computing these mean values we excluded 11 patients, for whom we had obtained only poor representations of the observed treatment dynamics (e.g. Patient 105), as the parameters associated with these fits are unlikely to be biologically meaningful (marked by grey boxes in Figure 3.15).

We find that allowing just cost and turnover to be patient-specific can explain the data almost as well as the full four parameter model (paired Student t-test: $t_{67,67} = 0.11$, $p = 0.91$). In contrast, the model assuming that inter-patient variability is caused by different initial conditions fits poorly (paired Student t-test: $t_{67,67} = 3.66$, $p < 0.01$; 95% CI for mean difference in r^2 : $[0.24, 0.84]$). For an overview of the fits, and a detailed patient-by-patient comparison we refer to Figures B.4 & B.5 in the Appendix. Overall, we conclude that according to the ABM resistance costs and turnover are the key parameters explaining the differences in the dynamics between patients.

3.6.4 Cycling speeds may reflect intra-tumoral organisation

In the next step we asked how cost and turnover relate to the differences in cycling speeds between patients. In Chapter 2, we found that fast cyclers were associated with large values of cost and small values of turnover, and vice versa for slow cyclers. Repeating this analysis with the ABM corroborates this result (Figure 3.16a). Firstly, we observe a negative correlation between the fitted values for the resistance cost and turnover (Pearson's correlation coefficient: $r = -0.76$, $p < 0.01$). Secondly, we find a positive correlation between turnover and the mean time taken per cycle (Pearson's correlation coefficient: $r = 0.74$, $p < 0.01$). Note that also in this case, 11 patients with unrepresentative fits were excluded from the analysis (see Figure B.4). Although weaker, a similar correlation is also seen when fitting all four parameters (not shown).

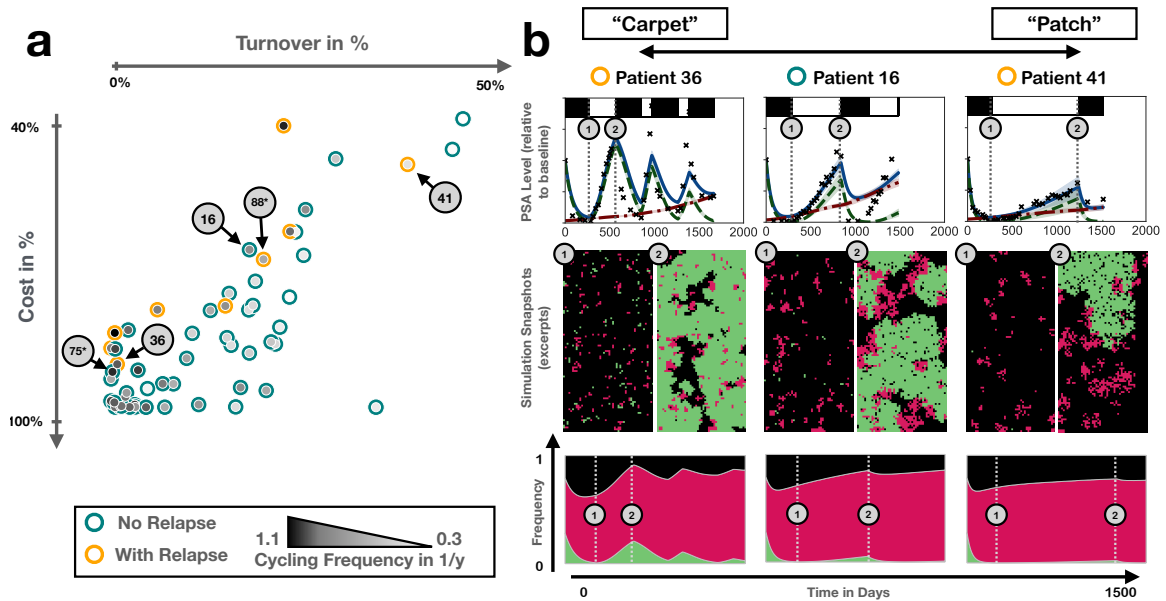


Figure 3.16: Analysis of the cycling dynamics in the full patient cohort. This suggests that different cycling speeds are linked to differences in the cost and turnover levels as well as in the spatial structure. Shown are the results for the 2-parameter model in which only cost and turnover were fitted. (a) Negative correlation between the estimated values of cost and turnover. Note that while Patients 75 and 88 are marked for reference, their fits here are not those shown in Figure 3.14a, which were fitted with all four parameters being allowed to vary. (b) Treatment trajectories, simulation snapshots and neighbourhood composition, illustrating the dynamics for patients in different areas of the parameter space in (a), ranging from what we term “carpet”-like appearance (many small, resistant colonies) to “patch”-like appearance (few, but large, resistant colonies). Treatment trajectories and neighbourhood analyses are based on 250 independent replicates. Neighbourhood plots were generated using EvoFreq [180].

Moreover, as we hypothesised, we observe a difference in the spatial structure of the tumour between slow and fast cycling patients (Figures 3.16a & b). Similar to our results from Figure 3.14, we find that fast cyclers are characterised by a diffuse structure with many small resistant nests, whereas in slow cyclers resistant colonies (and also sensitive colonies) are more defined and larger. The reason for this is that when turnover is low and cost is high (fast cyclers), most resistant cells present at the start of treatment will survive, but only expand very slowly, which yields the carpet-like appearance of these tumours. In contrast, when turnover is high and cost low (slow cyclers), many initially present resistant colonies will go extinct, but those that do survive will be able to expand more rapidly. Thus, the patch-like appearance of these tumours.

This observation has two implications. Firstly, it suggests that fast and slow cyclers may be characterised by different disease pathologies, with fast cyclers representing a more advanced stage of the disease with more invasive growth and/or metastasis, whereas slow cyclers represent an earlier, more confined stage. Secondly, it indicates that there may be differences in intra-tumoral competition and, thus, the benefit a patient can be expected to derive from an intermittent or adaptive approach. In order to illustrate this second point, we quantified the average neighbourhood composition of a resistant cell for different patients, i.e. what proportion of a cell's four neighbours is made up of which cell type. This shows that in slowly cycling patients almost all competition is intra-specific, whereas in fast cyclers, competition with sensitive cells plays more of a role (Figure 3.16b).

3.6.5 A combination of relatively smaller cost or turnover correlate with progression

To conclude, we revisit a final observation we made in Chapter 2, namely that patients who progressed on the trial were characterised not by a lack of cost or turnover but by a smaller combination of the two. The ABM analysis corroborates this context-dependence of the resistance costs. Progressors (yellow points) cluster along the upper boundary of the line of fits (Figure 3.16a). Accordingly, we detect no statistically significant difference in the turnover estimate (Mann-Whitney test, $U_{47,9} = 210$, $p = 0.5$), but a significant difference in the sum of the two ($c_R + d_T$; Mann-Whitney test, $U_{47,9} = 95$, $p < 0.01$). That being said, we do observe a statistically significant difference in the estimated cost values (Mann-Whitney test, $U_{47,9} = 45$, $p < 0.01$), which indicates that the resistance cost plays more of a role in the spatial ABM than it did in the non-spatial ODE model.

3.7 Discussion

In this chapter we extended our analysis of AT to consider the impact of space and of stochasticity. To do so, we developed a simple, 2-D ABM, based on the same set of assumptions as the ODE model in Chapter 2. The ABM is assumed to represent a small region of tumour tissue as a 2-D lattice, in which each tumour cell occupies one site, and is classified as either drug-sensitive or resistant. Using this model, we compared the 50% AT schedule from the prostate cancer clinical trial [97] with CT. We find that high initial tumour cell density, and low initial resistance fractions, maximise the benefit of AT, assuming a cure is not possible. This result corroborates

our findings from the previous chapter and is consistent with previous non-spatial [59, 1, 115, 122, 125, 126] and spatial theoretical studies [115, 124], and experimental evidence in cancer [115] and bacteria [107].

A question which was of particular interest to us was whether, and how, the impact of resistance costs and cell turnover is modified by space. We find that our conclusion that a resistance cost is not required for AT holds true also for the ABM. In combination with the work by Hansen et al [1] and Viossat and Noble [126], these results corroborate the hypothesis that AT can work even if there is uncertainty about the presence of resistance costs in patients. This is important given that resistance costs depend on the genetic [114, 150] and environmental context [181], and are, thus, likely to vary between, and possibly also within, a patient.

Similarly, we find that our key result from Chapter 2 that turnover may help resistance control by amplifying the impact of competition carries over to the ABM. In fact, in general, we conclude that there is good qualitative agreement between the ABM and the ODE model. Both agree that AT is, on average, beneficial when compared to CT, under a wide range of parameters. Theoretical studies to-date have employed either ODE models or ABMs, and even though reading of the literature indicates qualitative agreement, to the best of our knowledge this is the first time a direct comparison has been carried out.

That being said, when we perform a more detailed, quantitative, comparison of the two models, we observe significant differences. In particular, while AT is superior in both models, its relative benefit compared to CT is smaller in the ABM, and there is less gain from the presence of resistance costs and turnover. This is because the two models make different assumptions about spatial competition. In the ODE model, cells are assumed to be perfectly mixed, so that all cells experience the same competitive growth inhibition, which is simply a linear function of the total cell density. In contrast, the lack of migration in the ABM results in spatial segregation of different colonies, so that the competitive inhibition experienced by a cell depends on the cell's local neighbourhood, and varies across the tumour. Consequently, sensitive cells in the ODE model will always be able to competitively suppress resistant cells, whereas in the ABM this is only possible if the sensitive colony grows in close vicinity to the resistant colony. This indicates that the tumour's spatial organisation is likely to be an important factor in AT. Moreover, it shows that a detailed understanding of intra-tumoral competition is required, in order to determine whether or not a patient will receive a clinically meaningful gain from AT. This point is also supported by the recent work by Viossat and Noble [126], and Farrokhian et al [116], who found that

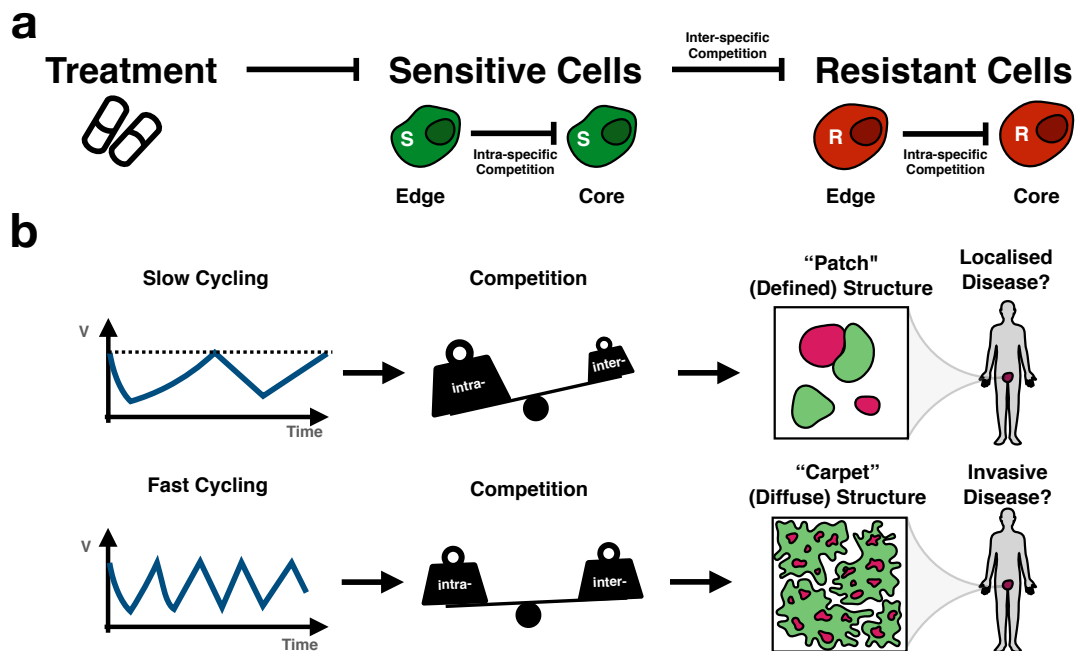


Figure 3.17: The two key results from this chapter. (a) The current paradigm foresees that AT keeps resistant cells in check by inter-specific competition from sensitive cells, which can be controlled via treatment [98, 104]. Here we show that intra-specific competition within each population is a further important factor. (b) Our analysis suggests that the more diffusely growing a tumour, the greater the fraction of inter-specific competition. In addition, this may be reflected in faster cycling frequency under intermittent treatment. Going forward, we propose that incorporating knowledge of the tumour’s spatial architecture and resulting competition structure will help to design more effective AT strategies.

while different ODE models of AT agree qualitatively, there are significant differences in their quantitative predictions depending on how competition is modelled.

In order to gain a better understanding of intra-tumoral competition during treatment, we leveraged the individual-based nature of our model to explicitly measure, for the first time, spatial competition between cells during therapy. This allowed us to visualise and quantify how treatment breaks during AT increase the competitive inhibition of resistant cells. Furthermore, we capitalised on this to explore how different model parameters affect competition, and to build on some of the ideas developed in the previous chapter. For example, we showed how reducing the initial cell density diminishes competition during AT, whereas a higher initial resistance fraction results in a similar level of inhibition per AT cycle, but fewer cycles.

Moreover, this analysis revealed that intra-specific competition of resistant cells with each other is an important, but so far overlooked, factor in AT. As a resistant cell

expands, most of its daughters will end up trapped in the central core of the resistant colony, able to divide only upon the death of another resistant cell (or upon migration out of the centre). As such, AT is most effective when resistance is clustered in a single location, and surrounded by sensitive cells, because it can leverage inter-specific competition at the edge of the resistant colony to maximise intra-specific competition between resistant cells at the core (Figure 3.17a).

An important implication of this observation is that it matters how resistance is distributed across the tumour. If resistance arises in a single location it can be controlled more effectively with AT than if resistance is present at multiple sites, either within the same lesion, or at different metastatic sites within the body. As such, we extend previous results by Bacevic et al [115] who showed in their computational model that resistant cells arising near the surface of a spheroid are harder to control than resistant cells arising near its core, and we characterise in more detail the non-trivial relationship between a tumour’s spatial organisation and AT.

How may we infer the spatial distribution of resistance? Tissue biopsies would provide the most direct and detailed measurements, but are invasive and often impractical. As an alternative, we show that it may be possible to use mathematical modelling to gather spatial insights from the patient’s longitudinal response dynamics (Figure 3.17b). We fit our ABM to PSA data from prostate cancer patients undergoing intermittent androgen deprivation therapy. We find that the speed at which patients cycle between treatment on- and off-periods correlates with distinct spatial organisations of the tumours in our simulations. While experimental validation of this hypothesis will be a question for future research, we do note that Bruchovsky et al [109] reported a “suggestive trend that a Gleason score <6 may be associated with a slightly longer time off treatment in the initial 2 cycles”. Indeed, lower Gleason scores indicate a more defined tissue architecture, which would be consistent with the more patch-like structure our model predicts for slowly cycling patients. As such, this hypothesis warrants further investigation.

In addition, it may also be possible to gather some information about the spatial distribution of resistance from the characteristics of the resistant population. In particular, if resistance is driven by a single clone, then it will likely be initially confined to a single, or at most a small number of, sites within the tumour. In contrast, if resistance is driven by multiple clones, as has, for example been observed in colorectal cancer [10], then it is likely to exist in multiple locations simultaneously. As such, liquid biopsies, which are showing promise at detecting and characterising the

clonality of emerging drug resistance, may allow inference of not only the abundance, but also the distribution of, resistance [47, 10].

But, so what? While we cannot easily alter it, understanding the spatial distribution of resistance in a patient may be relevant in the design of AT schedules. Gallaher et al [124] found in their off-lattice ABM that the rate of cell migration (and therefore of spatial mixing) determined whether a modulation-based adaptive algorithm (treatment is modulated in small increments, rather than withdrawn), or a vacation-based algorithm (treatment is either on or off) was more effective. In particular, spatially confined tumours favoured modulation, whereas in invasive tumours more benefit was derived from the vacation-based strategy [124]. As such, an important next step in extending this work would be investigation of how treatment should be adapted in light of different tumour architectures.

A final observation we make is that there can be significant variation in the benefit of AT between stochastic replicates despite identical model parameters. Variance depends on the number of resistant cells initially present in the simulations, and their distribution, which further highlights the importance of the spatial distribution of drug resistance within the tumour. Moreover, even though resistance costs and turnover increase the average benefit of AT they also increase variability in outcomes. As a result, the greater the average benefit of AT the more likely we are to see a patient fair much better or worse than expected. While the cell numbers in our simulations are unrealistically small, and as such our simulations are not suited to make quantitative statements about the magnitude of this problem, Hansen and Read [128] have recently raised similar concerns. We, thus, advocate further study of the impact of inter-patient variability in AT in order to inform future clinical trial design.

In aiming to keep our model tractable we have again made a number of simplifying assumptions. We assume no movement and no pushing of cells, which has been shown by Gallaher et al [124] and Bacevic et al [115] to reduce the benefit of AT, as it allows resistant cells to squeeze through surrounding sensitive cells. This may be another explanation as to why Bacevic et al [115] were unable to control resistance in 2-D cell culture, even though our work here suggests that control should have been possible. Moreover, for computational reasons, we restricted our analysis to a 2-D setting, which is arguably more representative of *in vitro* cell culture than a 3-D human tumour. We hypothesise that the extra dimension will hinder tumour control as it will allow resistant cells to more easily find space into which they can divide. That being said, we have herein also neglected the role of non-tumour tissue, which

acts as an additional competitor for space and resources, and may help to control resistant subpopulations.

A final, important caveat is that we have not explicitly modelled the mechanism by which resistance arises. Depending on whether it arises through mutation, phenotypic switching, or is environmentally induced, this may drive different initial distributions of resistant cells, and will also result in different dynamics during treatment due to *de novo* resistance acquisition (see also [1, 124, 126, 182]). Furthermore, our model cannot explain how some of the initial tumour compositions we have analysed would have arisen prior to treatment. For example, if resistance costs and turnover are assumed to be high, then resistance will disappear in our model, if the tumour is left untreated (not shown). This disconnect of the model of tumour evolution and treatment should be addressed in future research.

To sum up, in this chapter we have shown that the tumour's spatial architecture will likely be an important factor in AT, as it determines the balance of inter- and intra-specific competition. In addition, stochastic events may cause significant intra-patient heterogeneity, even when all tumour characteristics are the same. This strengthens the argument for patient-specific AT protocols. In the next chapter we will elaborate on this point by exploring whether we can use mathematical models to directly guide the treatment decision-making, in order to deliver a treatment schedule that is better tailored to each patient and their disease.

Chapter 4

Using modelling to guide adaptive therapy decision-making: two case studies in ovarian cancer

4.1 Introduction

In current clinical practice, cancer therapy is administered according to pre-determined treatment schedules. These are typically designed to maximise cell kill by administering as high a dose as frequently as toxicity permits [183]. Changes to these schedules are made only in the case of intolerable toxicity, or if the tumour progresses. In the previous two chapters we have shown how AT, which proposes to adjust treatment based on the tumour's drug response, can help to extend TTP. But, this requires a treatment algorithm to decide when, and how, to modulate treatment. So far we have investigated the algorithm used by Zhang et al [97] in their first-in-human clinical trial of AT in prostate cancer. While a mathematical model informed the trial, the final algorithm was a compromise of what the model suggested and what was clinically feasible (Bob Gatenby; personal communication). In fact, both the ODE model in Chapter 2, and the ABM in Chapter 3, predict that smaller burden reductions would have yielded longer tumour control, a finding also made by others [123, 126, 182]. To follow up on this observation, we will here consider the question of whether we can improve resistance management by integrating mathematical models *directly* into the decision-making process. We will use extensions of the ODE model from Chapter 2 to develop a framework for model-guided AT, and test it *in vitro* in the treatment of ovarian cancer spheroids with PARP inhibitors.

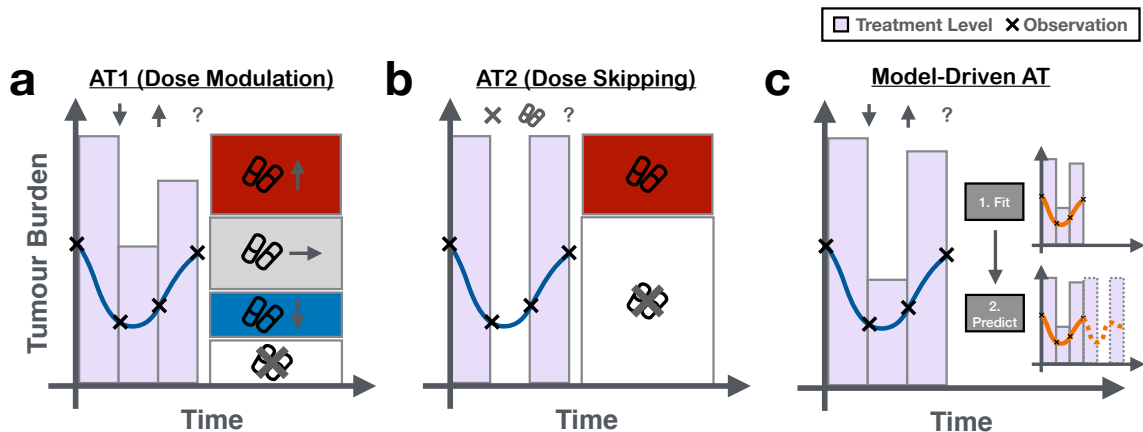


Figure 4.1: Overview of the different AT dosing algorithms proposed to-date. (a) AT1 modulates the treatment level, increasing it when the tumour grows (red), decreasing it when it shrinks (blue), and maintaining it the same when tumour size changes are small (grey). (b) AT2 takes a “bang-bang” approach, applying treatment when the tumour has grown (red), and withdrawing it when its expansion is under control (white). (c) The model-guided approach by Smalley et al [114] bases AT decisions on predictions of a mathematical model, which is continuously re-calibrated on the fly.

4.1.1 A review of AT dosing algorithms

While we have so far focussed on the 50% rule because of its clinical relevance (Equation (2.4)), this is not the only adaptive algorithm which has been proposed in the literature. Specifically, there are three others (Figure 4.1): i) dose-modulation (AT1), ii) dose-skipping (AT2), and iii) a model-driven algorithm by Smalley et al [114]. Before discussing how we may use modelling to improve AT, we will review each of these strategies, and discuss what is known about their relative merits. Subsequently, we will introduce the experimental system in which AT1 and AT2 will serve as benchmarks, and upon which we will seek to improve by extending the model-driven approach by Smalley et al [114].

Throughout this chapter we will assume that treatment is administered at k time points t_n , for $n = 0 \dots k$, where we will refer to t_n as a *treatment day*. We will define $V(t_n)$ as the tumour volume on day t_n , and denote the dose administered at time point t_n by, $D(t_n)$, or short, D_n . Administration could be as frequent as every day, but is more likely several days apart. The length of time between treatment days will be referred to as the *treatment interval*.

4.1.1.1 Dose Modulation (AT1)

AT1 was pioneered in the first paper on AT [60], and was subsequently refined in an *in vivo* study on breast cancer by Enriquez-Navas [113]. It modulates the dose administered at the current time point, $D(t_n)$, based on the change in tumour volume compared to the last treatment day, $V(t_n) - V(t_{n-1})$, and may be summarised as follows (Figure 4.1a):

$$D(t_n) = \begin{cases} 0, & \text{if } V(t_n) < V_{\min} \\ (1 + \alpha)D(t_{n-1}), & \text{if } V(t_n) > (1 + \beta)V(t_{n-1}) \\ (1 - \alpha)D(t_{n-1}), & \text{if } V(t_n) < (1 - \beta)V(t_{n-1}) \\ D(t_{n-1}), & \text{otherwise} \end{cases} \quad (4.1)$$

where α (> 0) describes the relative change in dose, and β (> 0) is the relative tumour size change which triggers treatment modulation. V_{\min} denotes a minimum tumour volume below which treatment is withdrawn completely. Therapy is started at some initial dose $D(t_0) = D_0$, and doses may not be escalated beyond the MTD, D_{Max} .

While D_0 and D_{Max} are determined by the experimental model system, it is, as of yet, not clear how the values of α and β should be chosen. In their study on ovarian cancer, Gatenby et al [60] used a decision threshold of $\beta = 10\%$, and modulated dose in absolute increments of 10 mg/kg. The subsequent study in breast cancer allowed for greater changes in tumour size ($\beta = 20\%$), and also made larger changes in the drug dose ($\alpha = 50\%$).

4.1.1.2 Dose Skipping (AT2)

The dose skipping strategy was developed by the same authors and is similar in spirit, except that treatment is withdrawn rather than modulated (Figure 4.1b). In addition, decisions are made based on volume changes over two treatment intervals, rather than just one:

$$D(t_n) = \begin{cases} D^*, & \text{if } V(t_n) \geq (1 + \beta)V(t_{n-2}) \\ 0, & \text{otherwise} \end{cases} \quad (4.2)$$

where again β denotes the relative change in tumour volume that triggers adaptation, and D^* is the administered drug dose. Therapy is started at $D(t_0) = D^*$.

This algorithm has only been experimentally tested once before, namely in the study by Enriquez-Navas et al [113], which used a value of $\beta = 25\%$. That being said it shares the vacation-based structure of the 50% rule from the clinical trial [97], which is further evidence that such a ‘‘bang-bang’’ approach can be effective in delaying progression.

4.1.1.3 Modulation vs skipping?

Which of these two approaches performs better? So far there have been two direct comparisons, one *in vivo* and one *in silico*, which indicate that the answer may depend on the tumour biology. Enriquez-Navas et al [113] benchmarked both AT1 and AT2 in the treatment of breast cancer with the chemotherapy paclitaxel in mice. They repeated the experiment with two human cell lines, one of which produced a fast growing and highly invasive cancer (MDA-MB-231/luc), and one of which was less aggressive (MCF7). The authors found that in both cases AT1 performed better than AT2, which was explained by the fact that the breaks during AT2 allowed for uncontrolled tumour growth, which could not be undone by subsequent treatment. They also found that AT1 resulted in less necrosis and better blood flow within the tumour [113].

In order to gain a deeper understanding of how different tumour characteristics may impact the choice of strategy, Gallaher et al [124] developed an off-lattice, ABM. Drug resistance of a cell was modelled as a continuum and assumed to be inversely correlated with the cell's proliferation rate (for more details on the model, see Chapter 3). The authors found that dose modulation did best when the mean level of resistance across the tumour was low and cells could not migrate nor mutate to a more drug-resistant phenotype. In contrast, in more resistant and invasive tumours, dose skipping was more effective. Indeed, also in the experiments by Enriquez-Navas et al [113], the difference between AT1 and AT2 was smaller in the invasive MDA cells than in the less aggressive MCF7 cells.

To sum up, the optimal strategy likely depends on the characteristics of the tumour. As such it is not surprising that Enriquez-Navas et al [113] observed significant heterogeneity in treatment responses, with some tumours being kept in check with ease, whereas others were never really brought under control, despite being treated according to the same strategy [113]. Together, these observations suggest that for AT to be most successful we should adapt not only treatment in a tumour-specific manner, but also the decision-making algorithm.

4.1.2 Use of mathematical models in treatment decision-making

4.1.2.1 Smalley et al: Model-guided AT

A first demonstration of how we may personalise AT using mathematical modelling was presented by Smalley et al [114] in an *in vivo* model of BRAF inhibitor treatment in melanoma (Figure 4.1c). The authors used a 2-population Lotka-Volterra ODE

model, similar to that studied in Chapter 2, to guide treatment individually for each mouse. A key ingredient of the model was the assumption that treatment induces the switching of cells from sensitive to resistant, and that in the absence of treatment this process is reversed. This was motivated by the observation of such phenotypic switching in *in vitro* experiments [114]. Each treatment day the model was fitted to the tumour's growth data, and used to predict tumour growth over the subsequent treatment interval. If this exceeded a given threshold, treatment would be applied. To capture uncertainty in the model parameters, the predictions were based on an ensemble of 200 fits, obtained via a genetic algorithm [114].

Smalley et al [114] compared their approach to a continuous and an intermittent schedule (2-weeks on followed by 1-week off). The individualised approach outperformed both, yielding the smallest tumour volume after 40 days of treatment [114]. In addition, there were noticeable differences in the treatment schedules between the mice, with those mice whose tumours were growing quickly receiving more treatment than those mice whose tumour had been effectively stabilised. Overall, this shows that modelling can be used to directly guide AT, and that it can yield better outcomes than fixed treatment schedules. Yet, it remains unknown how this approach compares against the simpler, more easily implemented, AT1 and AT2 algorithms.

4.1.2.2 Other model-guided treatment approaches

While the research by Smalley et al [114] is so far the first and only study to employ a model to directly drive decision-making in the context of AT, a number of other studies have explored similar concepts and are relevant to the work in this chapter. As we discussed in Section 1.2.7, there is a long tradition in the pharmacometrics community to use mathematical models to dynamically guide treatment schedules. This dates back to pioneering work by Lewis Sheiner [89] who proposed a Bayesian framework to estimate and to regularly update parameters of PK models from a patient's blood measurements, and to use the models to guide drug dosing for maximal safety and efficacy. In oncology, such model-guided TDM approaches have found particular applications in methotrexate treatment and in paediatric patients [90, 91, 92, 93, 94], but have been explored also in other settings [96, 87]. We note that while initially these models were often relatively simple (regression models or single compartments (e.g. [96, 90, 92])), increasingly more complex models of drug metabolism and side effects have been developed (e.g. [93, 94, 87]). The recent work by Meille et al [87] demonstrates how, through careful calibration and validation, comprehensive models can help to tackle more complex situations, such as combination treatment.

In parallel, a second community that has investigated dynamic treatment scheduling comprises researchers who have applied control theory to cancer therapy. When treatment is changed in response to its effect on the tumour, or the body, this can be translated into a “closed-loop” control problem in control theory. If a model is used to drive the decisions, as in the work of Smalley et al [114], this is known as *Model Predictive Control* (MPC). These approaches differ to those from the pharmacometrics community in that they have typically focussed on models with more details of the tumour biology and fewer considerations given to PK. In addition, very rarely have MPC approaches been tested in real life (wet-lab or clinic). To the best of our knowledge, some of the first contributions in this area came from Florian et al [184] and Chareyron et al [185], who performed MPC on a 2-state cell cycle ODE model of tumour growth, and an ODE model of tumour-immune interactions, respectively. Both concluded that MPC approaches are superior *in silico* to pre-determined schedules in addressing patient heterogeneity and uncertainty about model parameters, a result which has since been corroborated in several other studies [186, 187, 188].

The MPC work which is most pertinent to the current study can be found in the papers by Beckman et al [189] and Yeang and Beckman [190]. Beckman et al [189] performed MPC on a stochastic model of tumour evolution with two (or three) types of drugs and associated resistant states, in which cells become increasingly resistant through spontaneous mutation. They optimised treatment over intervals of 45 days allowing the model-based controller to choose between mono-therapy or combination therapy (the latter at reduced doses of the individual drugs). An interesting feature is that the authors compared different strategies. For example, they compared a strategy that always seeks to minimise predicted tumour size, to one in which the preferred treatment reduces the predicted level of multi-drug resistance, unless the tumour burden exceeds an intolerable level (named Strategy S2.2.). Subsequently, Beckman et al [189] carried out a large scale parameter sweep and found that MPC with Strategy S2.2. doubled median survival and increased cure rates from 1% to 15-20%. In a follow up study, Yeang and Beckman [190] investigated whether extending the optimisation horizon from 45 days to several months or years could help to further improve treatment. They found that this does not improve survival over their single-step approach but could further increase cure rates [190].

Together, these studies demonstrate that MPC in cancer treatment is in principle feasible and can yield superior tumour control compared to standard of care approaches. However, many studies assumed perfect knowledge of the underlying model, and to some extent model parameters, which is unrealistic in practice. This

may explain why in a majority of cases, experimental testing of the proposed frameworks is lacking. Moreover, relatively few studies have considered MPC in the context of drug resistance, and even though the work of Beckman and colleagues [189, 190] is comprehensive, they did not include cell-cell competition in their model.

4.1.3 Summary

To sum up, we draw three main conclusions from this review:

1. **The choice of AT algorithm matters:** The works of Gallaher et al [124] and Enriquez-Navas et al [113] demonstrate that it is not sufficient to simply maintain tumour burden, but that how this is done (skipping vs modulation) affects outcome.
2. **The superior adaptive algorithm depends on the underlying tumour biology:** This is shown by Gallaher et al [124] and suggests that to be most successful, AT should be carefully tailored to the particular clinical setting.
3. **Personalising schedules using modelling is *per se* feasible and can outperform standard-of-care regimes:** This is demonstrated, for example, by Smalley et al [114] *in vivo*, and Meille and colleagues [87, 88] in the clinic.

4.1.4 Outline

While the concept of using mathematical modelling to directly guide AT decision-making is appealing, and the *in vivo* results by Smalley et al [114] are encouraging, a number of open questions remain:

1. How does a model-driven approach compare to the simpler AT1 and AT2 algorithms? Is the additional complexity justified?
2. How can we address uncertainty about the underlying biology? Can we integrate predictions from multiple plausible models into decision-making?

In the following, we will present theoretical and experimental work aimed at addressing these questions. This project was a close collaboration with Professor Mehdi Damaghi from the Department for Cancer Physiology at the Moffitt Cancer Center. All wet-lab work was performed by him and his assistant, Samantha Byrne, whereas data analysis (unless otherwise stated) and modelling were carried out by myself. We will begin by introducing the experimental spheroid model system and the *in*

in vitro data which motivated this project. Subsequently, we will present a first study in which we use the ODE model by Smalley et al [114] to perform open-loop optimisation of the treatment schedule and test this schedule against AT1 and AT2 in the wet lab. As this performs poorly and we see significant heterogeneity between spheroids in their treatment response, we next develop an MPC approach to allow for individualised scheduling. To account for uncertainty in the underlying biology we propose a MPC framework which makes its decisions based not on a single, but on an ensemble of different ODE models, and we test this framework *in vitro*. Although the complexity of this project means that a number of unanswered questions remain, we gain useful insights into the practicalities of using mathematical models to drive treatment decisions in the real world.

4.2 Open-loop optimisation of PARPi treatment in ovarian cancer

4.2.1 Background

Ovarian cancer is one of the deadliest cancers in women, with a 5-year survival rate of 30% [18]. The primary reasons for the poor prognosis are the difficulty in detecting the cancer early and its high rate of recurrence [16]. Upon diagnosis, the current standard-of-care foresees cyto-reductive surgery to remove as much primary and metastatic tumour mass as possible, followed by platinum-based chemotherapy [16]. While ovarian cancer shows initially some of the highest platinum response rates seen in solid tumours, with up to 80% of patients responding, most patients will see their disease return within 3 years [16, 17]. Despite large-scale research efforts, the survival rate has changed little over the past 30 years [16, 191].

However, the recent introduction of PARP inhibitors (PARPis) is promising to revolutionise the field [192, 193, 191]. These agents target the PARP (poly-adenosine ribose polymerase) family of proteins, which are key orchestrators of single strand DNA repair [193, 191]. This results in the rapid accumulation of double strand DNA damage in proliferating cells, and eventually causes cell death [194, 195, 191]. Ovarian cancers are particularly sensitive to PARPis because their DNA repair is frequently deficient as a result of mutations in the Breast Cancer Associated (BRCA) genes (BRCA1 and BRCA2; [191]). As our approach in this chapter is relatively generic, we will defer a more detailed discussion of the molecular biology to the next chapter.

There are currently four FDA-approved PARPis on the market, all of which are given orally, and are primarily used as maintenance therapy after chemotherapy [191]. This means that treatment is administered after chemotherapy has been completed, in an attempt to eradicate the disease or, at least, push back progression. In the following, we will focus on the best characterised PARPi, Olaparib, which has been shown to increase TTP by between 3.6 months [196] to more than 30 months [197], depending on the BRCA status of the disease, and depending on prior treatment history. While this is a significant achievement, most patients still see their tumour recurring within 5 years (see also Section 5.1.1.2 for more details; [198, 197, 199]). In addition, around 40% of patients require dose adjustments due to serious Grade 3 or 4 adverse events, such as anaemia [17], and the cost for one month of treatment in the USA is about \$14,000 [200]. Thus, there is a need to develop strategies to further extend TTP under Olaparib, and to investigate whether, and how, cumulative drug use may be reduced.

4.2.2 A pilot experiment comparing continuous and adaptive Olaparib treatment *in vitro*

Given the promising AT results from our work and from the literature, we wondered whether we could extend TTP by scheduling Olaparib adaptively. To test this, Professor Damaghi carried out an experiment in which he compared continuous Olaparib treatment (CT) of *in vitro* ovarian cancer spheroids at MTD with three different schedules: two adaptive and one intermittent. The intermittent schedule (IMT) alternated between one week on MTD and one week off treatment. The adaptive schedules consisted of the dose-modulation schedule (AT1) and the dose-skipping schedule (AT2) from Enriquez-Navas et al [113] (see also Section 4.1.1), with starting doses $D(0) = D_{\text{Max}}$ and $D(0) = 1/2D_{\text{Max}}$. Treatment was administered as part of the culture medium in which the cells were immersed, at an MTD of $D_{\text{Max}} = 10\mu\text{M}$. The culture medium was replaced twice a week, allowing treatment adaptation every 2-3 days. Eight spheroids were treated according to each schedule (16 for control), where the decisions for the adaptive approaches were based on the average growth observed across the cohort. Growth was monitored for a total of 80 days.

4.2.2.1 Experimental methods

Professor Damaghi carried out this comparison on two cell lines, which are commonly used as ovarian cancer models: UWB^{WT} and UWB^{BRCA} . UWB^{WT} possesses a

mutation in BRCA1 making it *homologous repair deficient* (HRD) [201], and thus representative of the type of patient in whom PARPis would be strongly recommended in the clinic [191]. In contrast, UWB^{+BRCA} is a derivative of UWB^{WT} in which BRCA1 activity has been restored via transfection with a plasmid containing an intact copy of the *brca1* gene [201]. Both UWB^{WT} and UWB^{+BRCA} are quite slow growing, with 2-D culture doubling times of around 50 hours [201].

The spheroids were grown by seeding 5000 cells per well in a non-adherent 96 well U-bottom plate (Corning), covering them with Matrigel (ThermoFisher), and immersing them in culture medium. As culture medium, a 1:1 mixture of Roswell Park Memorial Institute Medium (RPMI) (ThermoFisher) and Mammary Epithelial Cell Growth Medium (PromoCell) was used, which was supplemented with 10% Fetal Bovine Serum and 1% penicillin/streptomycin. Cells were cultured for ten days in drug-free medium to allow aggregation into spheroids before treatment was started. Throughout the experiment cells were kept at 37°C and 5% CO₂.

Tumour growth was measured every 6 hours with an IncuCyte ZOOM time-lapse microscopy system (Essen BioScience). Phase-contrast white light, as well as green and red fluorescent channels, were recorded. Alive cells were found to be red auto-fluorescing (Figure 4.2a), possibly because of indigestion of the phenol in the medium. This information, together with the phase-contrast information, were used to segment the images to obtain the spheroid size (measured as the size of the fluorescing area in the images). Image processing was carried out by Mr Matthew Wicker who was a summer intern at the Anderson lab at the time, using openCV [176]. His algorithm first pre-processed each colour channel by using binary and Otsu-thresholding, followed by a blurring and dilation step to enhance the regions in which cells were present. Next, it segmented these regions using the `findContours()` method in openCV. Finally, it filtered the identified regions and removed those which were too small (less than 35 px wide or high), or which were close to the edge of the image, and as such were unlikely to be part of the spheroid. All subsequent data analysis (e.g. comparison of treatment arms) was performed by myself.

4.2.2.2 Results

In Figure 4.2 we show the results of this initial experiment. When they were grown in the absence of treatment, the spheroids expanded until they reached carrying capacity, at which point most of the cell mass was fluorescing and there was only a small corona of debris from dead cells around the outside (Figure 4.2a). In contrast, when treated with continuous Olaparib, growth was slowed down and there were many more dead

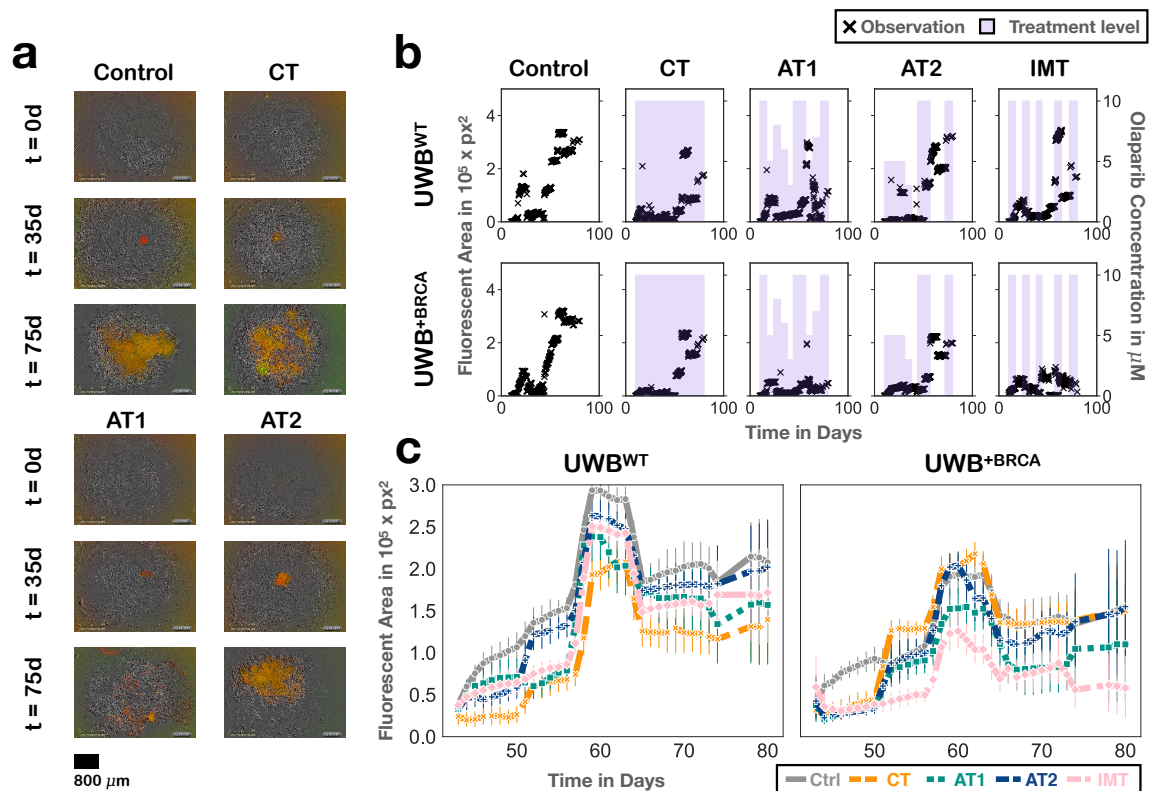


Figure 4.2: Pilot experiment comparing continuous Olaparib treatment to intermittent and adaptive administration. (a) Example microscopy images of UWB+BRCA tumour spheroids from different treatment conditions at different time points. Pictures show overlaid phase, red, and green channels. (b) Representative growth curves of individual spheroids under different treatment conditions. For UWB+BRCA, these correspond to the images in (a). Pink shading indicates Olaparib concentration in the culture medium. (c) Comparison of average tumour size across the different treatment strategies. While for UWB^{WT} continuous treatment tends to do best, for UWB+BRCA there appears to be an opportunity to improve control via alternative scheduling. Markers indicate daily averages, and error bars denote 95% CIs.

cells (Figure 4.2a). The same held true for the adaptive strategies and for intermittent treatment (the latter is not shown). Furthermore, both AT1 and AT2 spheroids saw adaptations of their treatment over time (Figure 4.2b).

In most spheroids we saw an initial expansion followed by a rapid regression over the first month of the experiment (Figure 4.2b). As this regression occurred in both treated and untreated cases, we reasoned that this must be part of the growth dynamics of these spheroids, rather than being an effect of treatment. As such, we decided to focus on the time period after this initial “bump” to compare the performance of the different treatment schedules (day 42 of the experiment). We find

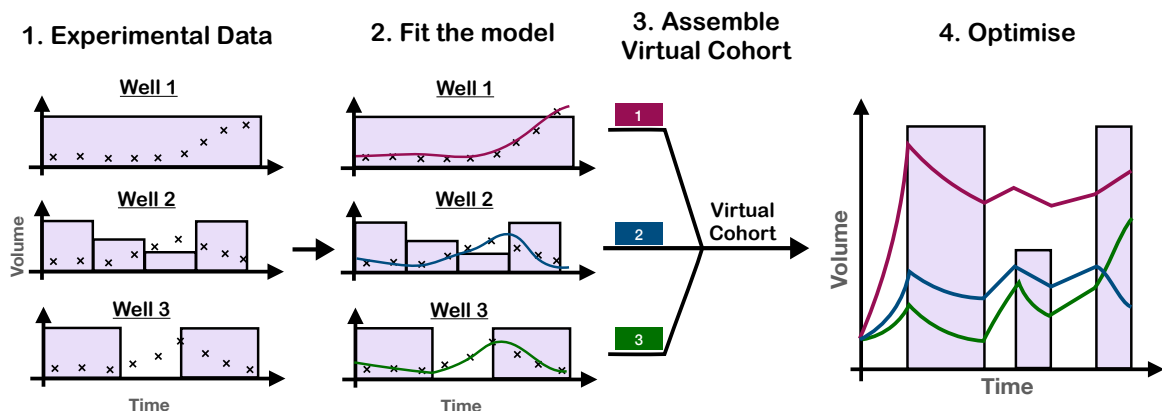


Figure 4.3: The virtual cohort approach to treatment optimisation. Using the experimental data from the pilot experiment, we calibrated a mathematical model to each individual spheroid (Steps 1 & 2). Next, from the best fit for each spheroid we assembled a “virtual cohort” of spheroids (Step 3). Finally, we carried out *in silico* treatment optimisation on this cohort, seeking a schedule that performed well across the different “virtual spheroids” (Step 4).

that in the UWB^{WT} spheroids, CT tended to keep the tumours smallest throughout the experiment, followed by AT1 and IMT which performed similarly, whilst AT2 performed poorest (Figure 4.2c). That being said, none of the differences between treated and un-treated tumours at 80 days was statistically significant (Student t-tests, all p-values > 0.05). On the other hand, in UWB^{+BRCA} spheroids, CT and AT2 only yielded very short-lived tumour size regression, whereas, AT1 and, in particular IMT, were able to better control tumour growth (Student t-test of the difference between Ctrl and IMT at day 80: $t_{16,8} = 2.69$; $p = 0.013$ Figure 4.2c).

Finally, we noticed that in both cell lines there was significant heterogeneity in treatment response. Despite identical experimental conditions and treatment schedules, some spheroids responded much better to treatment than others (Figure C.1). Together, these observations indicated that by expanding on the success of the IMT and AT1 schedules, and by better addressing the heterogeneity in drug responses, we may be able to improve TTP.

4.2.3 A virtual cohort approach to improve scheduling

Given these results, we asked whether we could use a mathematical model to derive an even better treatment schedule than IMT or AT1. To test this, we carried out model-based schedule optimisation and organised for Professor Damaghi to test our improved schedule in the next round of his *in vitro* experiments.

We identified two key challenges which our approach had to overcome. Firstly, while we had data available taken every 6 hours from a range of treatment conditions, all they provided us with were a measure of spheroid size. In addition, the data were quite noisy (see Figure 4.2b). We did consider re-analysing the images to characterise the tumours in more detail, by, for example, measuring the number of the dead cells over time, but we decided against this as we only had 6 weeks available before Professor Damaghi had scheduled the next set of experiments. This meant that even with a simple model we would likely face uncertainty in certain model parameters.

The second issue we needed to address was the heterogeneity in treatment response. This implied that even if we were to identify all parameters of our model, different spheroids would likely require different parameter combinations, and that, ideally, any schedule we found should perform well over a range of these combinations.

To address these issues we decided to employ a virtual cohort approach [202, 203]. This meant that rather than choosing a single set of model parameters for which to carry out treatment optimisation, we evaluated the performance of each candidate schedule on a set (“cohort”) of plausible parameter combinations *simultaneously* (Figure 4.3a). To assemble our virtual cohort we leveraged the structure of the pilot experiment. Specifically, we fitted the model to each of the spheroids in the pilot experiment individually to obtain a “virtual representation” of that spheroid, defined by the inferred parameters. In the next step we then carried out numerical optimisation to find the schedule that best controlled tumour growth *across* the cohort. We repeated this independently for the UWB^{WT} and UWB^{+BRCA} cells and subsequently handed our optimised schedules to Professor Damaghi for *in vitro* testing.

4.2.3.1 The mathematical model

Model of the intra-tumoural, eco-evolutionary dynamics: To the best of our knowledge, no mathematical model of the dynamics under PARPi treatment exists (see also discussion in Section 5.3). We therefore decided to adapt the model used by Smalley et al [114] as this had been successful in improving drug scheduling in their *in vivo* experiments whilst being sufficiently general to be a plausible model also for our current setting. The model is very similar to that which we analysed in Chapter 2, and assumes that we can divide the tumour into drug-sensitive cells ($S(t)$) and fully drug-resistant cells ($R(t)$), where competition between the two is described by

Table 4.1: Summary of mathematical model variables and parameters for the logistic model of spheroid growth (Equations (4.3), (4.4) and (4.7)). Also shown are the bounds imposed during model fitting and treatment optimisation.

Parameter	Description	Lower bound	Upper bound
$S(t)$	Density of sensitive cells (in cells/well)	0	K
$R(t)$	Density of resistant cells (in cells/well)	0	K
r_S	Sensitive cell growth rate (in d^{-1})	1×10^{-7}	5×10^{-1}
r_R	Resistant cell growth rate (in d^{-1})	1×10^{-7}	5×10^{-1}
K	Spheroid carrying capacity (in cells/well)	1×10^3	5×10^4
d_D	Drug-induced kill rate of sensitive cells (in $\text{d}^{-1} \mu\text{M}^{-1}$)	1×10^{-7}	5×10^{-1}
D_{Max}	Maximum tolerated dose (in μM ; not varied)	10	10
α	Drug-induced resistance acquisition rate (in $\text{d}^{-1} \mu\text{M}^{-1}$)	1×10^{-7}	5×10^{-1}
β	Re-sensitisation rate of resistant cells (in d^{-1})	1×10^{-7}	5×10^{-1}
S_0	Initial density of sensitive cells (in cells/well)	0	7500
R_0	Initial density of resistant cells (in cells/well)	0	7500
D_0	Initial drug concentration (in μM)	0	10

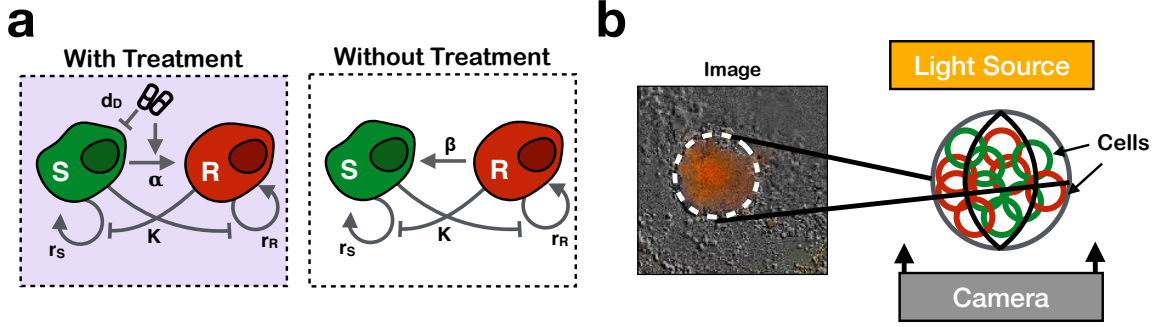


Figure 4.4: The mathematical model used to describe spheroid growth and drug response (Equations (4.3), (4.4) and (4.7)). **(a)** Tumour cells are divided into drug-sensitive and resistant sub-populations, which compete with each other according to Lotka-Volterra dynamics. Treatment kills sensitive cells and can induce resistance in the survivors. If treatment is withdrawn, then some resistant cells may lose resistance and re-sensitise. For a description of all model parameters, see Table 4.1. **(b)** To translate cell numbers to the fluorescent area seen on the images we assume that cells aggregate into a perfect sphere. On the images we then see a projection of this sphere onto a 2-D plane, so that the observed fluorescent area is approximately that of a plane section through the sphere.

a Lotka-Volterra model. The full model reads as follows (see also Figure 4.4a):

$$\frac{dS}{dt} = \underbrace{r_S \left(1 - \frac{S+R}{K}\right) S}_{\text{Growth}} - \underbrace{d_D D(t) S}_{\text{Drug Kill}} - \underbrace{\frac{\alpha}{D_{\text{Max}}} D(t) S}_{\text{Acquisition of Resistance}} + \underbrace{\beta \left(1 - \frac{D(t)}{D_{\text{Max}}}\right) R}_{\text{Re-Sensitisation}}, \quad (4.3)$$

$$\frac{dR}{dt} = \underbrace{r_R \left(1 - \frac{S+R}{K}\right) R}_{\text{Growth}} + \underbrace{\frac{\alpha}{D_{\text{Max}}} D(t) S}_{\text{Acquisition of Resistance}} - \underbrace{\beta \left(1 - \frac{D(t)}{D_{\text{Max}}}\right) R}_{\text{Re-Sensitisation}}, \quad (4.4)$$

where $D(t)$ describes the drug concentration in the medium, and the definitions of all other parameters are given in Table 4.1. The initial conditions are given by

$$S(0) = S_0, \quad R(0) = R_0 \quad \text{and} \quad D(0) = D_0.$$

As before this model assumes that cells divide at rates r_S and r_R , respectively, which may differ from each other to account for a cost of resistance. Furthermore, for simplicity, we assumed that both populations share the same carrying capacity K , and that both affect each other's growth equally (Lotka-Volterra competition coefficients set equal to 1).

There are two key differences in this model compared to that in Chapter 2. Firstly, this model assumes that treatment kills cells independently of their division status,

as proposed by Skipper et al [30]. While this is a strong assumption to make, it yields the simplest and possibly most common treatment law in the literature (see e.g. [28, 122, 146]), and it was successful in the work by Smalley et al [114]. Secondly, it assumes that cells can switch between drug-sensitive and resistant phenotypes and vice versa. With this assumption we can represent a wide range of plausible biological scenarios. For example, a current key discussion is about whether resistance is pre-existing or arises during treatment [46]. We did not know which was the case in the spheroids, but using this formulation we could explore both possibilities by varying the values of α , β , and R_0 .

The drug concentration, $D(t)$, is assumed to be that in the growth medium. As the medium is replaced frequently, and there is no enzymatic system which may clear drug, it is assumed $D(t)$ remains constant between medium changes.

Measurement model: While Equations (4.3) and (4.4) describe the total number of cells in the well, the microscopy images provide us not with a cell number, but with a measure of the size of the tumour spheroid. To translate the tumour cell numbers predicted by our model to the fluorescent area we would see on the images, we developed the following measurement model (Figure 4.4b). We assumed that all cells in the well aggregate into the spheroid, so that the volume of the spheroid at time t is given by:

$$V(t) = V_{\text{Cell}}(S(t) + R(t)),$$

where V_{Cell} is the volume of an individual cell, and we are assuming that both sensitive and resistant cells have the same size. Next, we assumed that the shape of the spheroid was approximately that of a perfect sphere, so that:

$$V_{\text{Cell}}(S(t) + R(t)) = \frac{4}{3}\pi r^3(t), \quad (4.5)$$

where $r(t)$ is the radius of the spheroid. The IncuCyte machine takes a picture of the spheroid from below, which means that what we see on the image is a projection of the sphere onto a 2-D plane. Thus, we assume that the fluorescent area which we observe on the images, $A_{\text{F}}(t)$, is approximately the area of a plane section through the sphere at its equator, i.e.:

$$A_{\text{F}}(t) = \pi r^2 \quad (4.6)$$

Solving Equation (4.5) for $r(t)$ and substituting the result into Equation (4.6) we obtain the following expression for $A_{\text{F}}(t)$:

$$A_{\text{F}}(t) = \left(\frac{3\sqrt{\pi}}{4} V_{\text{Cell}}(S(t) + R(t)) \right)^{2/3}. \quad (4.7)$$

Based on Professor Damaghi’s reports that UWB cells were approximately spherical with a diameter of $20\mu\text{m}$, we used the following estimate of cell volume in our calculations:

$$V_{\text{Cell}} = \frac{4}{3}\pi 10^3 \mu\text{m}^3. \quad (4.8)$$

Finally, to convert from the observed area in μm^2 to its size in pixels on the image, we obtained from the IncuCyte Zoom manual [204] that: $1\text{px} = 10\mu\text{m}^2$.

In order to allow plotting of the total predicted tumour size as well as the proportions of sensitive and resistant cells on one axis, we introduce $A_{\text{F}}^{\text{S}}(t) = \frac{S(t)}{S(t)+R(t)}A_{\text{F}}(t)$ and $A_{\text{F}}^{\text{R}}(t) = \frac{R(t)}{S(t)+R(t)}A_{\text{F}}(t)$ to denote the fluorescent area accounted for by sensitive cells and resistant cells, respectively.

The model was solved using the Runge-Kutta 4(5) scheme implemented in Scipy [179] with default solver options.

4.2.3.2 Model calibration

In the next step we fitted the model to the data from the pilot experiment by minimising the RMSE between the observed and the predicted fluorescent areas. We fitted a total of eight parameters, $\vec{\theta} = (r_{\text{S}}, r_{\text{R}}, K, d_{\text{D}}, \alpha, \beta, S_0, R_0)$. The reason why we included the initial conditions for $S(t)$ and $R(t)$ (S_0 and R_0 , respectively) as parameters was to account for the imprecision in seeding precisely 5000 cells. The bounds imposed on each parameter are shown in Table 4.1. These were derived from a combination of discussions with Professor Damaghi, and a review of the literature. Optimisation was carried out using the `basinhopping()` routine in Scipy [179], with default settings and a maximum number of 50 optimisation steps. For each spheroid we performed 25 independent fits from different initial conditions, and the best fit (lowest RMSE) was taken forward for further analysis. An overview of the best fit for each spheroid is shown in Figure C.1 in the appendix.

As we discussed in Section 4.2.2.2, most spheroids underwent an initial expansion and contraction which could not be explained as being the result of cell kill caused by treatment. In preliminary experiments we found that this can cause problems for the fitting routine as the model struggles to reconcile these rapid initial dynamics with the slower subsequent growth (not shown). We, therefore, decided to calibrate the model based only on the data from day 35 onwards, in line with our earlier analysis focussing on this second part of the experiment. In order to reduce the noise in the data we averaged it over 2 day intervals.

4.2.3.3 Schedule optimisation

Having fitted the model, we next assembled our virtual spheroid cohort and used it to predict an improved treatment schedule. For each of the two cell lines, the cohort consisted of the best fits for each of the 32 spheroids which had been subjected to treatment. That is, if $\hat{\theta}_j$ denotes the best fit parameter combination for spheroid in well j , then the cohort for UWB^{WT} cells is given by $\mathcal{C} = \left(\hat{\theta}_j\right)_{j \in \mathcal{S}^{\text{WT}}}$, where \mathcal{S}^{WT} is the set of treated UWB^{WT} spheroids, and similarly for the UWB^{+BRCA} cohort. Fits from untreated tumours were excluded because no treatment information had gone into calibrating their treatment parameters.

We then carried out optimisation to find the treatment schedule, \vec{x} , which minimised the following two-part, objective function:

$$\mathcal{J}(\vec{x}; \mathcal{C}) = \sum_{\hat{\theta}_j \in \mathcal{C}} \frac{1}{T_{\text{End}}} \underbrace{\sum_{t=0}^{T_{\text{End}}} A_{\text{F}}(t; \hat{\theta}_j)}_{\text{Mean Spheroid Size}} + \underbrace{w A_{\text{F}}(T_{\text{End}}; \hat{\theta}_j)}_{\text{Final Spheroid Size}}, \quad (4.9)$$

where \mathcal{C} is the cohort for which we were carrying out optimisation, and T_{End} is the length of the experiment in days. The reason why we chose to minimise both the time-averaged tumour size as well as the final size was that if we only targeted the latter, then the resulting schedules would simply delay treatment until the end of the experiment (not shown). This is because the tumours then finish at their nadir, as resistance has not had time to emerge, but it leaves the spheroids to grow in an uncontrolled fashion for most of the experiment. The multiplier $w \in \mathbb{R}^+$ sets the relative weight that is given to minimising the average and final tumour size. We tested different values of w and finally settled on a value of $w = 1.25$, which we found to provide a good trade-off between global and final size reduction.

The plan for the experiments was as follows: after letting the spheroids aggregate for a week, they would be treated every Tuesday and Friday for a total of $T_{\text{End}} = 73d$, starting on a Friday. We therefore explored candidate schedules that consisted of a sequence of 19 treatment intervals, $\vec{x} = (\{t_n^{\text{Start}}, t_n^{\text{End}}, D_n\})_{n=0}^{19}$, where t_n^{Start} and t_n^{End} are, respectively, the start and end time of each interval. D_n is the drug concentration in the medium, which is administered on day t_n^{Start} and is assumed to remain constant until t_n^{End} . Since the timing of the treatment is fixed by the Tuesday/Friday schedule of the experiment, this yields a discrete optimisation problem consisting of the identification of the optimal sequence of the 19 doses to administer. To reduce the size of the search space and to ensure that the prescribed doses were practical to

administer, we restricted the possible values for D_n to be one of $D_n \in \{0, 5, 10\} \mu\text{M}$. Despite these constraints this still left over 1 million candidate schedules to explore.

To solve the resulting discrete optimisation problem we employed a genetic algorithm (GA). These algorithms gradually evolve a population of candidate solutions through a process mimicking Darwinian evolution. At each generation, the best performing (“fittest”) candidates within the population are selected, and then recombined and mutated to produce the offspring population for the next generation. We implemented the GA using the DEAP package in Python [205], using as simple a structure as possible. We carried out selection using so-called “tournaments” in which five schedules are randomly picked at a time and the fittest is allowed to contribute one offspring to the next generation. Tournaments are sampled with replacement so that a schedule can potentially contribute multiple offspring. Subsequently, variation is introduced into the offspring. We implemented two types of variation: i) “mutation” in which one of the doses in the schedule is changed randomly, and ii) “cross-over” in which two schedules randomly exchange their recommendations for one of the treatment intervals. During the variation step, cross-over is applied first, followed by mutation, where, each time, a solution has a 50% chance of being altered. We carried out optimisation using a population size of 500 schedules, initiated randomly, and evolved it for a total of 100 generations. We also tested alternative optimisation parameters to those above, but found none which improved on the results (not shown). The final schedule for each cell line is given in Table 4.2.

4.2.3.4 Experimental methods

The experimental protocol was the same as in Section 4.2.2.2, except that this time the cells had been fluorescently labelled using red fluorescent protein (RFP). In addition, the initial treatment-free aggregation period was shortened from 10 days to 7 days. As before, images were analysed by Mr Wicker using a re-calibrated version of the algorithm from Section 4.2.2.2. All post-hoc analysis was carried out by myself.

4.2.4 Results

4.2.4.1 The model recapitulates tumour growth and yields plausible treatment schedules

To begin with, we attempted to fit the model to all spheroids using a single, shared parameter set. This failed to provide a good fit (not shown). To address this issue, we fitted the model to each spheroid individually. In Figure 4.5a we show example fits for

Table 4.2: The improved treatment schedules predicted by the mathematical model. Columns labelled UWB^{WT} and UWB^{+BRCA} give the recommended dose.

Interval	Date	Week Day	UWB^{WT} (in μM)	UWB^{+BRCA} (in μM)
0	27.07.2018	Friday	10	10
1	31.07.2018	Tuesday	10	5
2	03.08.2018	Friday	10	10
3	07.08.2018	Tuesday	0	0
4	10.08.2018	Friday	10	10
5	14.08.2018	Tuesday	0	5
6	17.08.2018	Friday	10	0
7	21.08.2018	Tuesday	5	10
8	24.08.2018	Friday	10	0
9	28.08.2018	Tuesday	0	10
10	31.08.2018	Friday	10	0
11	04.09.2018	Tuesday	0	10
12	07.09.2018	Friday	10	5
13	11.09.2018	Tuesday	0	10
14	14.09.2018	Friday	10	10
15	18.09.2018	Tuesday	10	0
16	21.09.2018	Friday	5	10
17	25.09.2018	Tuesday	10	10
18	28.09.2018	Friday	10	10
19	02.10.2018	Tuesday	10	10

spheroids from each of the different treatment branches for the UWB^{+BRCA} cell line. We find that despite its generic nature, the model can now describe the dynamics for most spheroids well (median $r^2 = 0.58$; see also Figure C.1 for an overview of all fits). At the same time, we also find evidence for issues with parameter identifiability. For example, for the spheroid in well A1, the model predicts that it consists entirely of resistant cells (Figure C.1). This is unlikely to be true, but because no treatment was administered in this well there is no information for the model on the basis of which to differentiate between drug-sensitive and resistant cells. As a result, we decided to exclude fits from untreated tumours when predicting an improved schedule.

In order to model the heterogeneity in treatment response and address the uncertainty in the model parameters, we assembled a virtual cohort, composed of the fits of each treated spheroid. Based on this cohort we derived the schedules shown in Table 4.2, which are intended to minimise tumour growth across the cohort. We find that, even though the two schedules are designed for different cell lines, they share a similar structure: To begin with, there is a treatment intense, lead-in phase, which is followed by a period of intermittent treatment in which doses alternate between 0 and D_{Max} . Towards the end of the experiment, both regimens then recommend CT at MTD. The intermittent-like structure of the proposed schedule is consistent with the success of IMT in the UWB^{+BRCA} cells. Similarly, the fact that the UWB^{WT} schedule recommends a slightly higher cumulative dose than the UWB^{+BRCA} schedule ($140\mu\text{M}$ vs $135\mu\text{M}$, respectively) matches the better performance of high dose CT in the UWB^{WT} cells. That being said, the overall similarity between the two schedules for the two cell lines is unexpected, given the differences between the cell lines in the *in vitro* experiments. Furthermore, we note that the intermediate dose of $5\mu\text{M}$ is recommended only rarely and without a clear pattern, even though AT1 performed well in both experiments. This suggests that the model predicts that a dose skipping approach is superior to dose modulation in this setting.

4.2.4.2 The virtual cohort approach captures both quantitative and qualitative differences in treatment response between spheroids

To better understand why the schedules in Table 4.2 were predicted to improve outcome, we studied the associated treatment responses in more detail. We will discuss here the results for UWB^{+BRCA}, but similar observations hold true also for UWB^{WT}. We find that the model predicted a range of different outcomes across the spheroids in the virtual cohort (Figure 4.5b). In addition, whilst in some cases the optimised

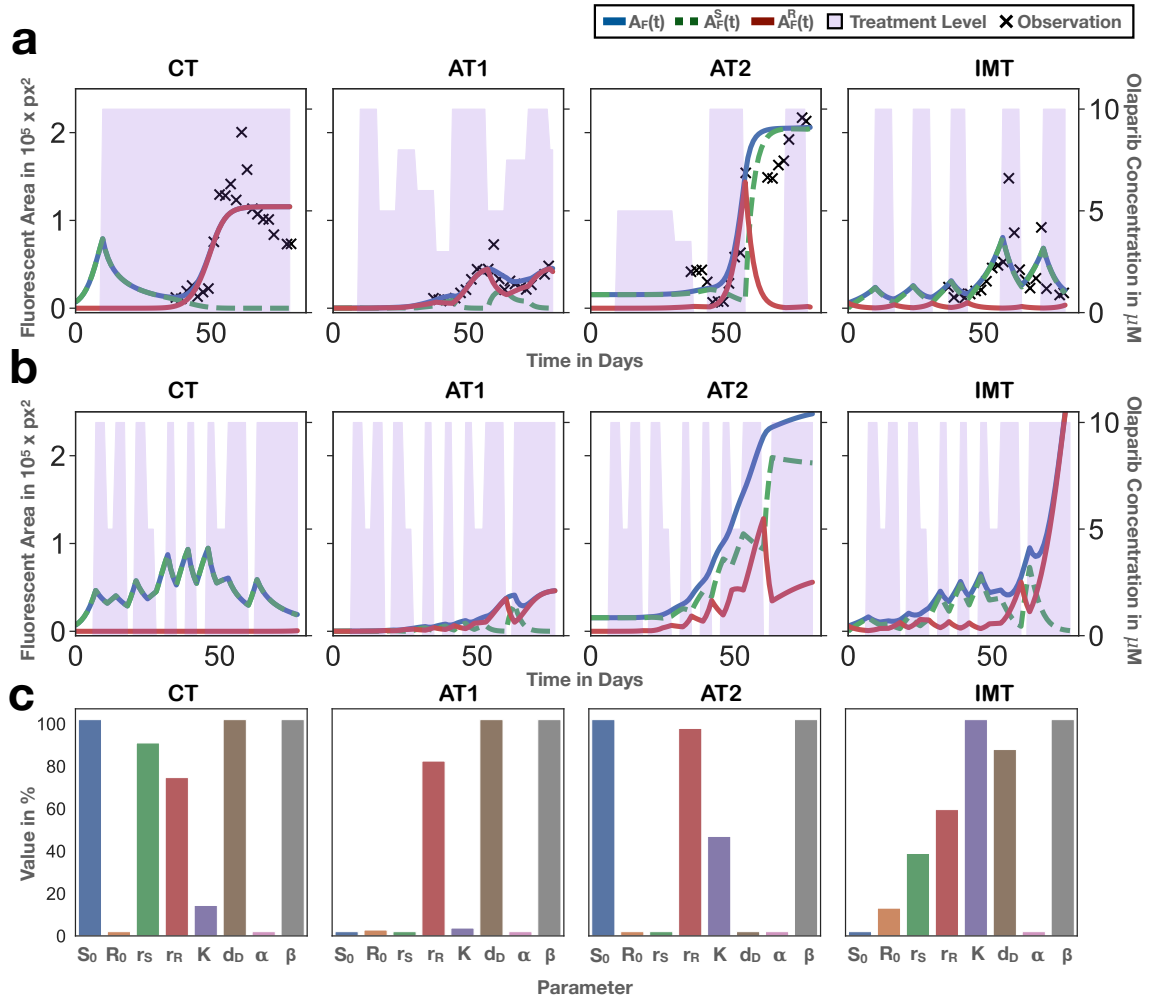


Figure 4.5: Analysis of the model-derived schedule for the $\text{UWB}^{+\text{BRCA}}$ cells. (a) Representative model fits to spheroids from the pilot experiment from each of the different treatment branches. Shown is the best fit in each case. (b) Prediction of the tumour's trajectory under the model-derived treatment schedule, based on the fits shown in (a). (c) Model parameters associated with the fits in (a) and (b). To allow plotting on one axis, each parameter is scaled relative to its range (0%: lower bound; 100%: upper bound). For the definition and bounds of each parameter, see Table 4.1.

schedule was predicted to perform significantly better than the schedule that was actually administered to this spheroid in the experiment (Figures 4.5a & b; left panel), in other cases the opposite was true (Figures 4.5a & b; right panel). While we did not have data to validate the proposed intra-tumoral dynamics, this showed that, at least phenomenologically, our cohort approach had been successful in capturing possible heterogeneity in treatment responses. Moreover, the modelled heterogeneity was both quantitative as well as qualitative: not only were there quantitative differ-

ences in how much an individual spheroid was predicted to benefit from the optimised schedule but, in fact, some spheroids were predicted to require a different schedule to others.

4.2.4.3 Differences in model predictions are driven by differences in the treatment schedule of the training data

Next, we investigated what was underlying the heterogeneity in the virtual cohort. We found that how a spheroid was predicted to respond to the optimised schedule depended strongly on the calibration data underlying that fit (Figure 4.5a & b). When trained on data from a spheroid treated with CT, the model inferred dominance of sensitive over resistant cells in the absence of treatment ($S_0 > R_0$ and $r_S > r_R$; Figure 4.5, left column). Thus, for CT it predicted competitive release, whereas for the optimised schedule it predicted that drug sensitivity in this spheroid would be maintained by a combination of competitive suppression, and re-sensitisation during the breaks (Figures 4.5a & 4.5b; left panel). In contrast, if the model was calibrated with data from a spheroid that underwent IMT, the model predicted an initial dominance of resistant cells ($R_0 > S_0$ and $r_R > r_S$; Figure 4.5c). Because of the long breaks and the high re-sensitisation rate, β , the resistant population was predicted to be held at bay during IMT. However, when treated with the optimised schedule, the breaks were shorter and treatment was escalated towards the end, so that the model predicted that the “optimised” schedule would actually perform poorer on this spheroid than IMT. Two further scenarios were observed in fits derived from spheroids treated with either AT1 or AT2.

These results provide two insights. Firstly, there is a strong influence of the training data, in particular the treatment schedule, on the subsequent treatment predictions. This is likely because it determines which parameters can be identified, and indicates that careful thought is required when selecting the data for calibrating a treatment model. Secondly, this shows that the final schedule obtained by our optimisation is not optimal in that it yields the best result for every spheroid in the cohort. Instead it is a compromise between the requirements set by different possible scenarios. Depending on whether these scenarios actually occur in the experiment, and at what frequency, we may see the schedule perform better or worse than expected.

4.2.4.4 The model-predicted schedule fails to improve outcome *in vitro*

In the final step, we asked Professor Damaghi to test our model-derived schedules *in vitro*. Because of insufficient resources, the experiment had to be limited to only

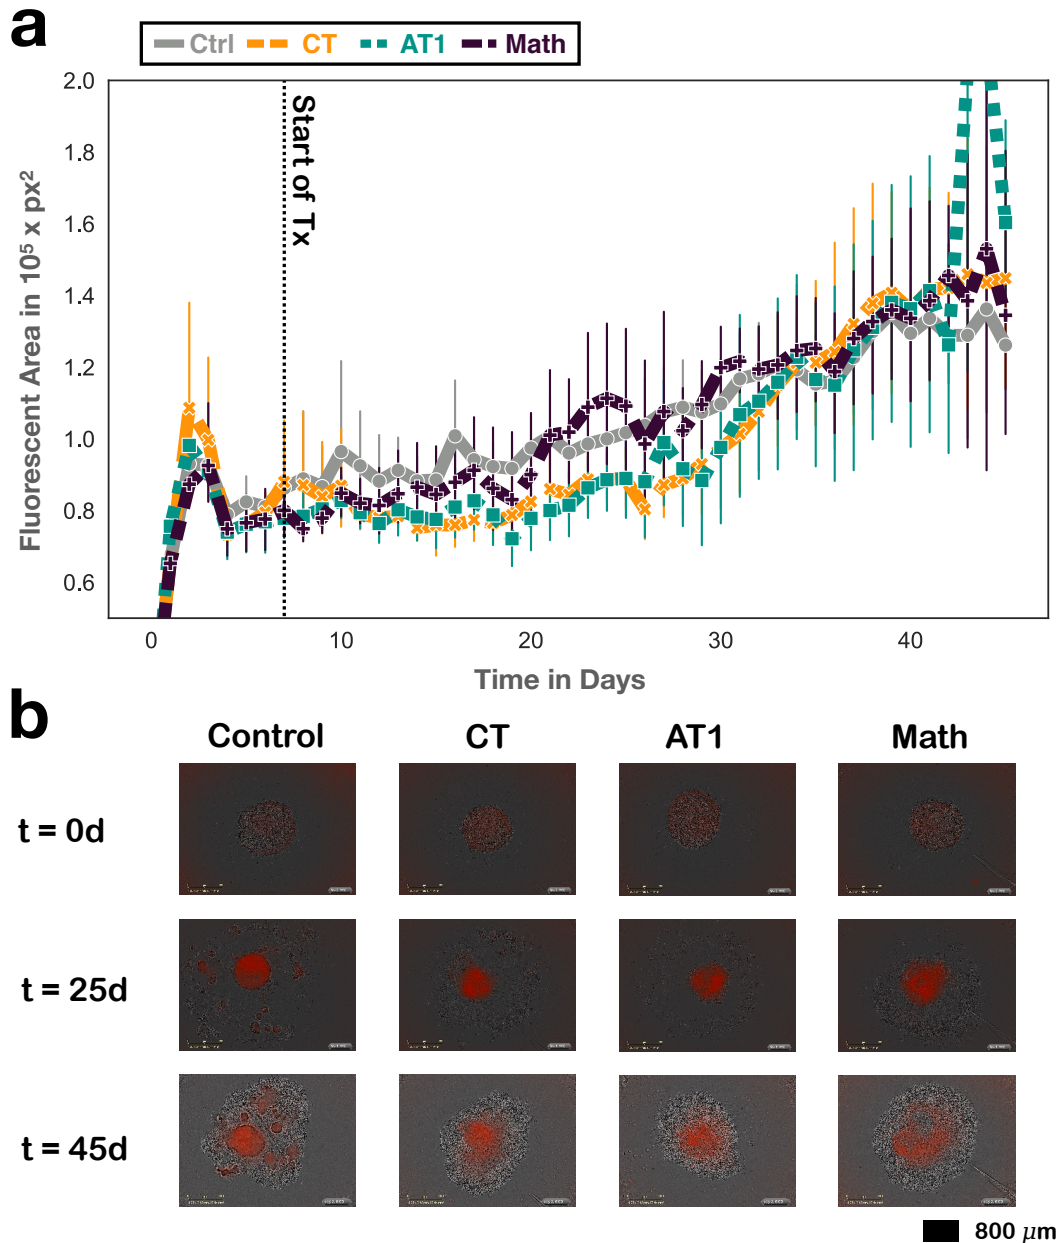


Figure 4.6: Test of the model-predicted schedule (Math) on $\text{UWB}^{+\text{BRCA}}$ cells *in vitro* (schedule as shown in Table 4.2). **(a)** Comparison of spheroid size over time under different treatment schedules. The CT and AT1 schedules can achieve a temporary slow down of tumour growth, whereas the spheroids under the Math schedule grow in a manner almost indistinguishable from untreated tumours. Markers denote daily mean sizes and error bars indicate 95% CIs. Sample size: Ctrl (16), CT (16), AT1 (8), and Math (8). **(b)** Representative time-lapse microscopy images for each treatment branch. Shown are phase-contrast images with overlaid red channel.

one of the two cell lines. We decided to test the $\text{UWB}^{+\text{BRCA}}$ schedule because of the promising results of IMT in these cells during the pilot experiment. Aside from the model-derived schedule (Math; Table 4.2, right column), Professor Damaghi also trialled the CT and the AT1 adaptive schedule. Figure 4.6 shows the results of this experiment. We find that the model-predicted schedule (Math) performs more poorly than either CT or AT1. Whilst both CT and AT1 were able to induce temporary reductions in the tumour growth rate, the Math treated spheroids grew almost identically to the untreated tumours. That being said, the treatment effect was relatively small and after about 40 days the different treatment branches were all indistinguishable from each other, and from untreated tumours. In addition, whenever the medium was changed, the spheroids were moved slightly, which could make them look bigger or smaller, or bring them out of the plane of focus of the camera (Figure C.2). Given that the resulting measurement error was on the order of magnitude of the differences between treatment branches, it was difficult to draw strong conclusions about the performance of each schedule. As a result, Professor Damaghi decided to terminate the experiment early.

4.2.4.5 Heterogeneity in response to Math schedule *in vitro* agrees with that predicted by the model

In order to understand why the Math schedule performed poorly, we studied the treatment trajectories of individual spheroids in the experimental cohort. We find that, as expected, there was heterogeneity in the responses, with some spheroids being controlled by treatment, whereas others escaped (Figure 4.7). Moreover, there was qualitative agreement between the different types of responses observed *in vitro* and those predicted by the model (compare Figure 4.7a to Figure 4.5b). While this observation does not validate the mechanism by which the model predicts how heterogeneity arises, it does show that the heterogeneity captured within the virtual cohort is realistic.

4.2.5 Discussion

The aim of this section was to test whether we could use a simple mathematical model of the interactions between sensitive and resistant cells to develop a more effective PARPi treatment schedule in ovarian cancer. We adapted the mathematical model by Smalley et al [114], which is similar to that which we had studied in Chapter 2, and calibrated it using data from an *in vitro* spheroid model, provided to us from a collaborator. We found that despite its simplicity, the model was able to represent the

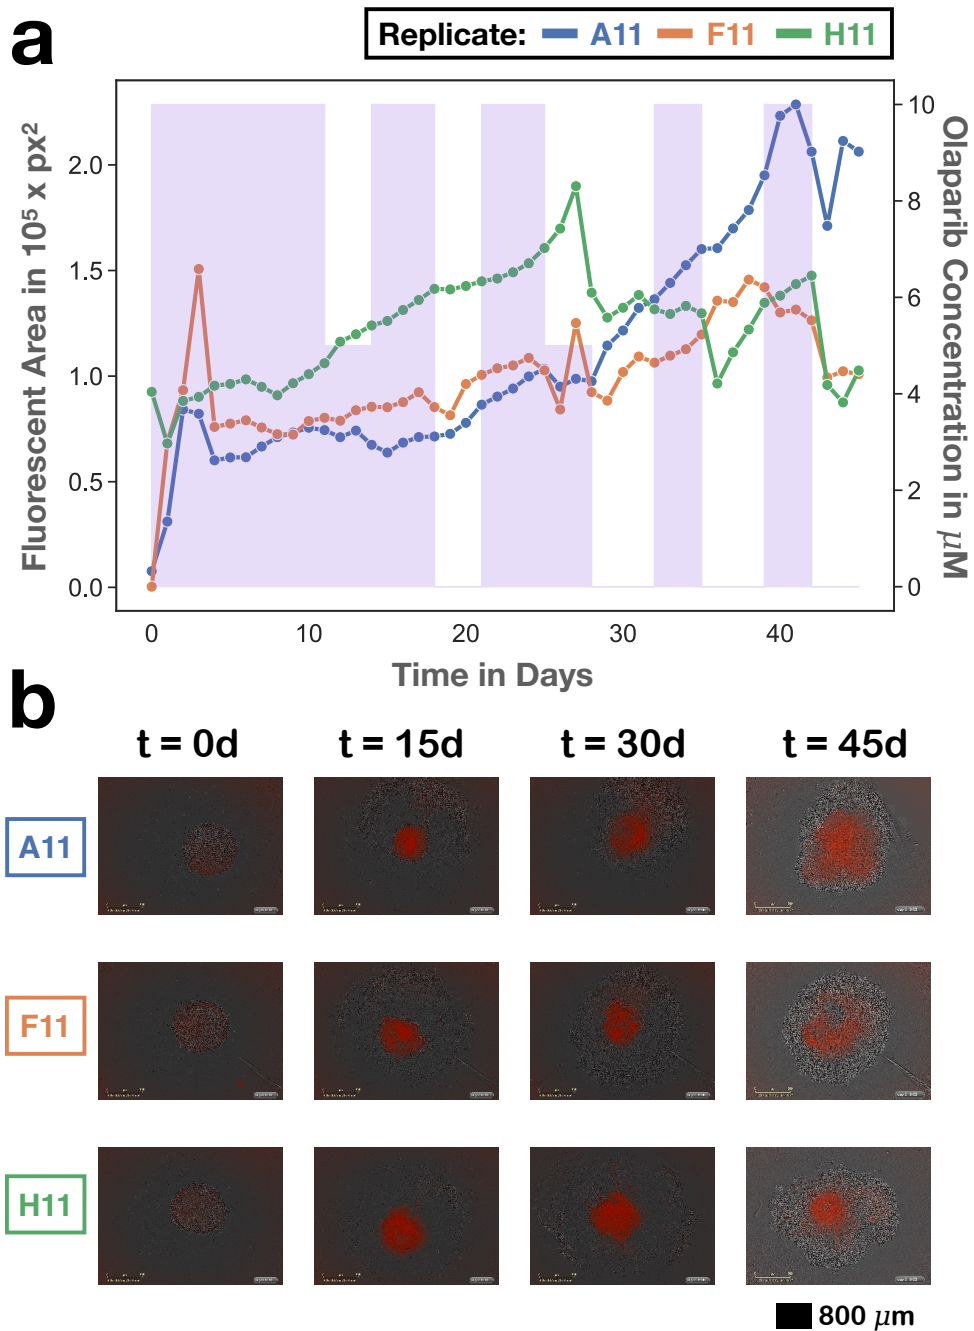


Figure 4.7: Heterogeneity in the *in vitro* response to the model-predicted schedule. (a) Spheroid size over time for three spheroids in the Math cohort (schedule as in Table 4.2; UWB^{+BRCA} cells). While the spheroids in wells F11 and H11 are controlled by treatment, the spheroid in well A11 escapes. The trajectories qualitatively agree with the possible outcomes predicted by the model (Figure 4.5b). (b) Representative time-lapse microscopy images for each treatment branch. Shown are phase-contrast images with overlaid red channel.

spheroid growth dynamics, and we were able to derive a plausible treatment schedule, which was subsequently tested back in the wet lab.

One of the key challenges we faced in our approach was calibrating the model. Even though the data were collected frequently, they only provided us with a measure of the total tumour size. As a result it was difficult to identify certain parameters with confidence (e.g. switching rates), and there was a strong influence of the calibration treatment schedule on the values of the inferred parameters. The latter can be explained by the fact that a schedule with both on- as well as off-treatment periods provides more information on both sensitive and resistant cell growth, than one in which treatment is mostly on or off. This indicates that it is likely we were overfitting the data. One approach to overcome this problem could have been, for example, to more carefully integrate the data from different treatment branches via a hierarchical or mixed effect model [25, 206, 207]. This would have also allowed us to make use of the untreated control data, which we have neglected in our current approach.

To account for the uncertainty in the model parameters, we optimised the treatment schedule not for a single, but for a cohort, of different model fits. Virtual cohort approaches are increasingly being used to compare different treatment protocols in the pharmacometrics community and in industry (e.g. [87, 208, 203, 81, 209]). Again, mixed effect models provide a very intuitive and powerful approach for capturing inter-subject variability. To do so, distributions are estimated for the relevant parameters which are then sampled to simulate the range of plausible dynamics. However, this requires detailed understanding of the correlation between model parameters, in order to ensure that the sampled parameter combinations still produce realistic behaviour [209]. Because of the time constraints (we only had 6 weeks available), we instead decided on a simpler approach in which we assembled our cohorts directly from the best fit for each of the treated spheroids. This was meant to ensure that each “virtual spheroid” was reflective of responses which had actually been observed *in vitro*. That being said, if this work was to be repeated it would be better to use a mixed effect framework. Software packages, such as Monolix©(Lixoft) or NONMEM©(ICON), provide powerful tools for the development of mixed effect models, and facilitate access also to those not familiar with the technique.

Nevertheless, our results show that the virtual cohort we assembled in this fashion does successfully capture heterogeneity in treatment response in which some spheroids will be controlled by the model-predicted schedule, and others will not. Moreover, the predicted trajectories do phenomenologically mimic those observed *in vitro*. While

the mechanisms by which the model predicts heterogeneity to arise remain to be validated, this does suggest the approach has potential.

However, just because we can account for heterogeneity does not mean we can manage it. Our aim was to find a schedule that works robustly for different assumptions on the spheroid's eco-evolutionary dynamics. To do so, we minimised tumour growth across the cohort. We identified two challenges: Firstly, careful choice of the objective function is required. Secondly, when optimising across several different model parameterisations simultaneously, the final schedule was a compromise between the requirements set by the different scenarios. As a result, it was predicted to improve outcome only for certain conditions, whereas in others it made little difference or was even predicted to worsen outcome. With more knowledge of the underlying biology, we could potentially weigh the different parameterisations to ensure outcome is improved in those scenarios which are most common in reality. That being said, the problem that treatment remains optimised only for a subgroup of cases persists.

To sum up, in this section we performed open-loop control on a simple model of sensitive-resistant cell competition in order to try to derive a better treatment schedule for PARPi treatment of ovarian cancer spheroids. We found that even though the model described overall tumour growth well, the predicted schedule did not improve on CT or an adaptive AT1 schedule. A key issue was heterogeneity in the treatment response. This is what we will seek to address in the next section.

4.3 Individualised, closed-loop optimisation of PARP inhibitor treatment in ovarian cancer

4.3.1 Introduction

AT is a multi-faceted concept, which incorporates both the notion that treatment should be adapted based on the tumour's response, as well as the idea that by de-escalating treatment we can leverage intra-tumoral competition for resistance control. In the previous section, we attempted to do the latter without the former. We aimed to leverage competition by incorporating it into the model with which we determined the treatment - but once the schedule was set we did not modify it. A key issue with this approach was that *in silico*, as well as *in vitro*, the schedule worked well for some spheroids whereas for others it did not. This suggests that treatment should be more carefully tailored to the competition and response dynamics of an individual spheroid.

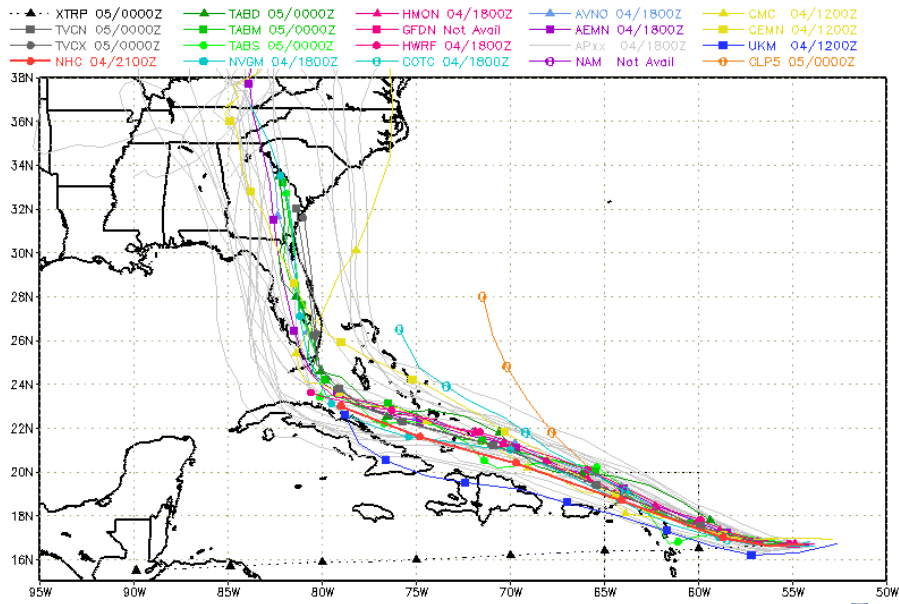


Figure 4.8: Spaghetti plot showing different model predictions for the trajectory of Hurricane Irma in September 2017. Each line represents a different model, or model parametrisation. Examining different models allows us to gage the uncertainty in the predictions, and perform more robust emergency preparations. Reproduced with permission from the South Florida Water Management District webpage (www.sfwmd.gov).

Our work in the previous section also revealed that it was difficult to determine with certainty, from the data at hand, what mechanisms were driving the treatment response. To address this issue we had based our predictions, not on a single model parameter set, but on a cohort of plausible values. However, given we were unable to fit all the data with one parameter set, this suggests that the model we used provided an incomplete description of the underlying dynamics. Traditionally, in mathematical oncology, at this point we would carefully review the biology and seek to develop a more accurate model. While we subsequently did so (see Chapter 5), we first wondered whether we could get around this problem by working, not with a single, but with an ensemble of *different models*.

Independent of how carefully we refine models, they will always remain simplifications, and as such falsifications, of the reality [210]. As mathematical models make their way into clinical decision-making, even with carefully chosen models and model systems, there will always be uncertainty about whether the model captures all the relevant dynamics within a particular patient. A cautionary example of this can be seen in weather forecasting: Despite the huge amount of effort that has gone into model development, and the vast amounts of data that are collected for calibration

every day, uncertainty in the predictions remains. Nevertheless, these models are used to guide important decisions, such as which regions to evacuate when a hurricane approaches. To take into account the uncertainty about the precise underlying physics, or model parametrisation, the predictions of different models are considered. These are typically visualised as “spaghetti plots” as shown in Figure 4.8, where each line represents a plausible path the hurricane may take. The predictions are then compared with the observed trajectory of the storm, in order to determine which model (or models) capture the dynamics most accurately. Based on this assessment, guidance is then issued, and decisions are made.

Given that mathematical models of human cancers are less advanced than weather models, and given that available data for calibration are a lot sparser, we hypothesise that uncertainty about whether a model accurately describes a patient’s disease will be an important challenge to overcome in integrating models into clinical decision-making. In light of this, we decided to explore whether we could adapt the idea of using a model ensemble, in order to address the uncertainty about the biology in the current experiment. We developed a simple proof-of-concept algorithm which we tested in the next round of Professor Damaghi’s experiments. In the following, we will first describe our framework, and subsequently discuss its performance relative to CT, AT1 and AT2 *in vitro*. To the best of our knowledge, this is the first time an ensemble has been used in cancer treatment modelling, and while our results are preliminary, they do demonstrate some of the potential benefits of such an approach. Moreover, our work identifies important questions which require future investigation.

4.3.2 The framework

4.3.2.1 Overview

The aim of our approach was to deliver individualised treatment recommendations, which were updated each treatment day, based on the most recent data. Moreover, to account for the uncertainty about the underlying biology we wanted to base our decisions not on a single model, but on an ensemble of different models. In this fashion we could explore different plausible scenarios about the biology, and consider the treatment recommendation in each case.

Importantly, what was of interest to us was the degree to which the models agreed on which treatment decision was the best to take. If all models agree on a schedule as being the best way to treat the spheroid, then we can adopt this schedule with confidence: independent of the exact spheroid biology, as represented by the

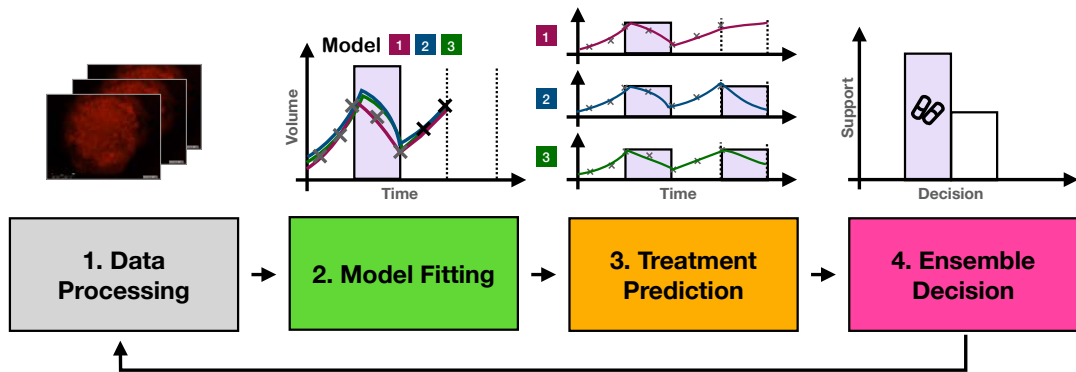


Figure 4.9: A closed-loop, ensemble-based method for treatment decision making. Decisions are made not based on a single, but on an ensemble of models, representing different assumptions about the underlying biology. At each decision point, each of the models in the ensemble is calibrated with the most up-to-date data, and is used to make a prediction about the optimal treatment schedule. Subsequently, these predictions are compared with each other, to identify the schedule with the most support, which is then administered.

different models, there is no alternative decision which will yield a better outcome. Importantly, this does not imply that the models necessarily agree on the quantitative benefit of the predicted schedule, or even that the proposed schedule achieves a particularly good tumour size control. But it does mean that within the options available to us, there is no better alternative. Conversely, if the recommendations vary between models, or even between fits of the same model, then this indicates that the models need to be scrutinised carefully, to choose the one in which we have the most confidence.

With this in mind we developed a proof-of-principle framework based on an ensemble of six ODE models. Broadly, our approach involved four steps, which were carried out individually for each spheroid every 3-4 days in order to determine which dose to administer next (Figure 4.9):

1. **Data Processing:** To begin with, we processed the most recent imaging data to obtain up-to-date information on the spheroid’s treatment response.
2. **Model Fitting:** In the next step, we fitted each of the models in our ensemble to the available data.
3. **Treatment Prediction:** Subsequently, we carried out *in silico* optimisation on each of the models in order to obtain their treatment recommendation.

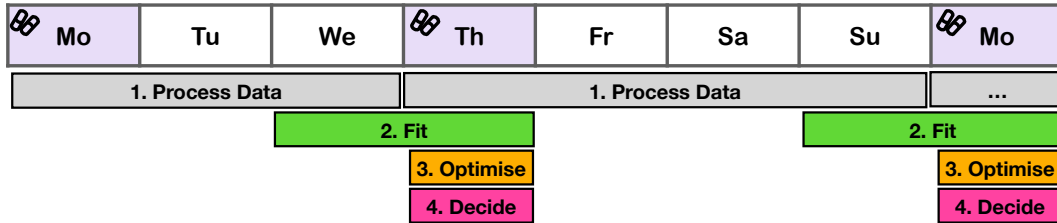


Figure 4.10: The typical weekly schedule during the course of the experiment. Treatment decisions were made on Mondays and Thursdays, when the growth medium was changed. Data export and transfer took 2-3 days and so was started as soon as the medium change had been carried out. Once processing of the most up-to-date data was completed, the fitting and optimisation pipelines were started. When these steps had been completed (12-24 hours), the treatment decision for the next time interval was made, the treatment administered, and the process repeated.

4. **Ensemble Decision-Making:** Finally, we integrated the recommendations of the different models to decide which treatment dose to administer.

In the following, we will first introduce the model ensemble, and subsequently discuss each of the steps in detail. The typical weekly schedule of the experiment is illustrated in Figure 4.10, indicating the timing of the different steps. One added complexity in carrying out this experiment was that the laboratory work was carried out by Professor Damaghi in Tampa, whereas I was based in Oxford at the time. By making use of the high performance computing facilities at the Mathematical Institute in Oxford, I was able to carry out all computations within a day. However, transferring the images, so that I could process them, represented a significant time bottleneck (Figure 4.10).

4.3.2.2 The model ensemble

While, in principle, there are many different models one could explore, our main hypothesis was that it would be intra-tumoral competition which would allow us to improve TTP. Thus, we assumed in all models that the spheroid was divided into drug-sensitive and resistant cells, and we focussed on two aspects which Chapters 2 & 3 had identified as being crucial in choosing a treatment schedule which aims to exploit competition: i) the strength of competition, and ii) the rate at which cells switch between sensitive and resistant states.

We explored a total of six different models (Table 4.3). These represented a combination of three different models of competition, and two assumptions about

Table 4.3: The ensemble of models used to guide treatment decisions. Model: Name of model; GL: growth law; Swg: Switching; #Ps: Number of parameters; $\hat{D}(t) = D(t)/D_{\text{Max}}$. For a description of each model parameter, see Table 4.1.

Model	GL	Swg?	#Ps	Equations
E-NS	Exp.	No	4	$\frac{dS}{dt} = r_S S - d_D \hat{D}(t) S$ $\frac{dR}{dt} = r_R R$
E-S	Exp.	Yes	6	$\frac{dS}{dt} = r_S S - d_D \hat{D}(t) S - \alpha \hat{D}(t) S + \beta (1 - \hat{D}(t)) R$ $\frac{dR}{dt} = r_R R + \alpha \hat{D}(t) S - \beta (1 - \hat{D}(t)) R$
L-NS	Log.	No	5	$\frac{dS}{dt} = r_S \left(1 - \frac{S+R}{K}\right) S - d_D \hat{D}(t) S$ $\frac{dR}{dt} = r_R \left(1 - \frac{S+R}{K}\right) R$
L-S	Log.	Yes	7	$\frac{dS}{dt} = r_S \left(1 - \frac{S+R}{K}\right) S - d_D \hat{D}(t) S - \alpha \hat{D}(t) S + \beta (1 - \hat{D}(t)) R$ $\frac{dR}{dt} = r_R \left(1 - \frac{S+R}{K}\right) R + \alpha \hat{D}(t) S - \beta (1 - \hat{D}(t)) R$
S-NS	Sph.	No	4	$\frac{dS}{dt} = r_S \frac{S}{\sqrt[3]{S+R}} - d_D \hat{D}(t) S$ $\frac{dR}{dt} = r_R \frac{R}{\sqrt[3]{S+R}}$
S-S	Sph.	Yes	6	$\frac{dS}{dt} = r_S \frac{S}{\sqrt[3]{S+R}} - d_D \hat{D}(t) S - \alpha \hat{D}(t) S + \beta (1 - \hat{D}(t)) R$ $\frac{dR}{dt} = r_R \frac{R}{\sqrt[3]{S+R}} + \alpha \hat{D}(t) S - \beta (1 - \hat{D}(t)) R$

switching (either it was possible, or not). The reason we chose to vary the strength of competition was that in Chapter 3 we found that the Lotka-Volterra model can provide a poor representation of the competition within a tumour, resulting in the over- or under-estimation of the benefit of treatment breaks. In addition, several prior studies have shown that the logistic growth model provides a poor fit to spheroid growth data (e.g. [148, 41]). Therefore, we explored the following three models:

1. **Exponential Growth:** Both populations grow exponentially and independently from each other. This model was meant to serve as a “worst case” scenario, both in terms of the speed of tumour growth, as well as in terms of the gain from competition (more specifically, the lack thereof).
2. **Logistic Growth:** Both populations grow logistically and interact in a Lotka-Volterra fashion, which assumes growth inhibition from competition is linearly proportional to overall cell density. This was chosen because it is the most common model to have been used for AT, and because it was successful in the *in vivo* experiment by Smalley et al [114].
3. **Spherical Growth:** The cells are assumed to grow as a sphere in which only the cells on the surface of the sphere can divide. As such, competitive inhibition is proportional to both the spheroid’s surface area as well as the relative abundance of each population (see below for details).

While the first two options are classical choices from the literature, the spherical growth model was specifically designed by us to provide a more accurate representation of the spheroid’s growth dynamics. It is inspired by the von Bertalanffy growth law [38], which has been shown to fit tumour growth well [148, 41]:

$$\frac{dV(t)}{dt} = rV^{2/3},$$

where $V(t)$ is the tumour volume, and r its growth rate. The assumption underlying this model is that the tumour grows as a perfect sphere, and that only the cells on its surface are able to divide. Thus, the number of dividing cells is taken to be proportional to the sphere’s surface, $V^{2/3}$ [38].

To incorporate the interaction between two cell populations, we extended this model as follows. Assuming the spheroid is composed of sensitive and resistant sub-populations, its volume is given by $V(t) = V_{\text{Cell}}(S(t) + R(t))$. Thus, its surface area and so total number of dividing cells is proportional to $(S(t)+R(t))^{2/3}$. Assuming that

sensitive and resistant cells are homogeneously distributed throughout the spheroid, the fraction of sensitive cells on the surface is proportional to:

$$\frac{S(t)}{S(t) + R(t)}(S(t) + R(t))^{2/3},$$

and vice versa for resistant cells. Assuming no switching between cell states, this yields the following set of equations for $S(t)$ and $R(t)$:

$$\begin{aligned}\frac{dS}{dt} &= r_S \frac{S(t)}{S(t) + R(t)}(S(t) + R(t))^{2/3} - d_D \hat{D}(t)S \\ \frac{dR}{dt} &= r_R \frac{R(t)}{S(t) + R(t)}(S(t) + R(t))^{2/3}\end{aligned}$$

where $\hat{D}(t) = D(t)/D_{\text{Max}}$, r_S and r_R are the cell growth rates, and d_D is the drug kill rate. Simplifying the first term in each equation, and optionally adding switching, then gives the equations in Table 4.3.

We note that both the exponential and the spherical model allow for unbounded growth of the spheroid. While this is unrealistic in the long run, it is plausible for the shorter time periods over which we will seek to perform optimisation. Conversely, the logistic model makes the more realistic assumption of a maximum size K . However, inferring this size before it is reached is difficult, and the logistic growth model has been shown to be prone to underpredicting growth as a result [41, 211]. We reasoned that by including all three options (fast unbounded, slow unbounded, and saturating growth), the models might balance out each others' weaknesses, and we would be able to select the most appropriate model and treatment strategy.

The other aspect we decided to vary between our models was the rate of drug-induced switching between sensitive and resistant cells. The reason why this distinction is important is our result from Chapter 2, that if there is no switching from one state to the other, then it is always beneficial for resistance control to maintain sensitive cells at the largest possible population size. In contrast, as Hansen et al [1, 126] and Viossat and Noble [128] have shown, if a drug induces resistance in sensitive cells, then it can be beneficial to reduce the size of the sensitive population (see also Section C in the Appendix for an illustration of this). Thus, depending on the rate of switching, we may want to take opposite treatment approaches.

While in principle the non-switching models in Table 4.3 are special cases of the switching models (with switching rates set to 0), we decided to include them explicitly, as we were uncertain whether we would be able to infer switching, or lack thereof, from the data (see also the identifiability analysis by Greene et al [146]). By considering

both options explicitly, we reasoned that we would be able to gain greater confidence in our decisions.

Finally, we note that for all models we carried over the assumptions from the previous section that the treatment level, $D(t)$, was constant between treatment days and changed in a step-wise fashion upon medium replacement.

As before, we implemented all models in Python and solved them using the `solve_ivp()` method in Scipy [179]. During the fitting, problems with stiffness could sometimes arise, so we used the implicit, backwards differentiation method provided in Scipy (BDF) to solve the ODEs.

4.3.2.3 Experimental setup

The experimental protocol remained similar to that used before. Cells were grown as spheroids, with medium changes twice a week (Mondays and Thursdays), and time-lapse imaging every 6 hours. However, because of their slow growth and weak treatment response, Professor Damaghi decided to exchange the cell lines for faster growing alternatives and increase the MTD to $D_{\text{Max}} = 100\mu\text{M}$. UWB^{WT} was substituted with A2780s, which also carries a mutation in BRCA, albeit BRCA2 [201]. $\text{UWB}^{+\text{BRCA}}$ was exchanged with TOV112d, which does not have a BRCA mutation [201]. Their doubling times are about half those of the UWB cell lines (18h and 23h, respectively), and both produce a durable response to Olaparib when treated at MTD (see also Figure 4.18). Due to the fast growth of the cells, no mammary epithelial growth supplement was used, and only two days were required for the cells to assemble into spheroids. The cells were transfected with RFP to allow for fluorescent imaging. Otherwise the same reagents and protocols were used as in Sections 4.2.2.2 & 4.2.3.4.

In addition to the model-driven AT strategy, Professor Damaghi also tested the CT (at MTD), AT1, and AT2 schedules. The sample size for each cell line was 12 spheroids as untreated control, and 8 spheroids for each treatment branch, except for AT1 where only 4 spheroids were treated. The experiment took 50 days.

4.3.2.4 The lead-in phase

At the beginning of treatment no data were available with which to calibrate the models, so, initially, a different algorithm was required to decide which dose to administer. For this “lead-in phase” we decided to treat the spheroids for 1 week with CT, followed by 2 weeks with the AT2 schedule. The 1 week of CT was recommended to us by Professor Damaghi in order to establish initial control over the spheroids.

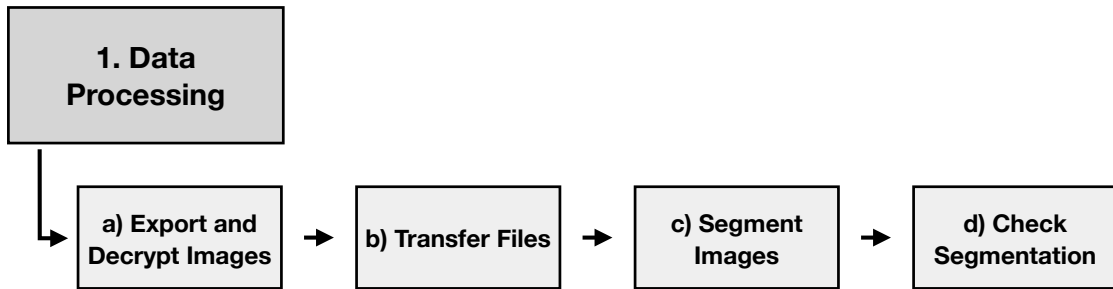


Figure 4.11: Overview of the steps involved in data processing.

The 2 weeks of AT2 were chosen to provide us with both on- as well as off-treatment data. After this 3 week lead-in period, all decisions were driven by the model.

4.3.2.5 Data processing

In the first step of our pipeline, we obtained and processed the most recent imaging data. While conceptually simple, this was the most time consuming step of all (Figure 4.10). Firstly, the images had to be exported from the IncuCyte machine via its proprietary software suite (IncuCyte Zoom 2). Next, the images had to be decrypted in order to allow them to be transferred to a non-Moffitt computer. File encryption had been introduced just a few weeks before the start of our experiments to improve cyber security at Moffitt. As the system was still in its early stages, and it had not been foreseen that large amounts of raw data would be exported, this step was quite slow: exporting 2-3 days worth of images took itself 2-3 days. All these steps were carried out by Professor Damaghi, or his assistant, Ms Byrne.

Subsequently, we transferred the images to my computer in Oxford, where I segmented them using the Python script written by Mr Wicker. During the pilot experiments in Section 4.2 we had sometimes encountered problems with imaging or segmentation, which resulted in distortion of the tumour’s growth trajectory (Figure C.2). As the data in the current experiment would directly drive treatment decisions, it was vital that they were accurately reflecting the growth dynamics of the spheroid. To ensure the integrity of the data, I therefore manually checked each image and its associated segmentation, and corrected distortions as necessary. Similar to the previous experiments, we found that after a change in medium there was usually a discontinuity in the reported spheroid size (Figure 4.12a). One reason for this was that after a change of medium, the fluorescence often temporarily dropped, but then rapidly recovered (Figure 4.12b). Even though this resulted in discontinuities in the data we decided not to alter or neglect these jumps, as they were part of the

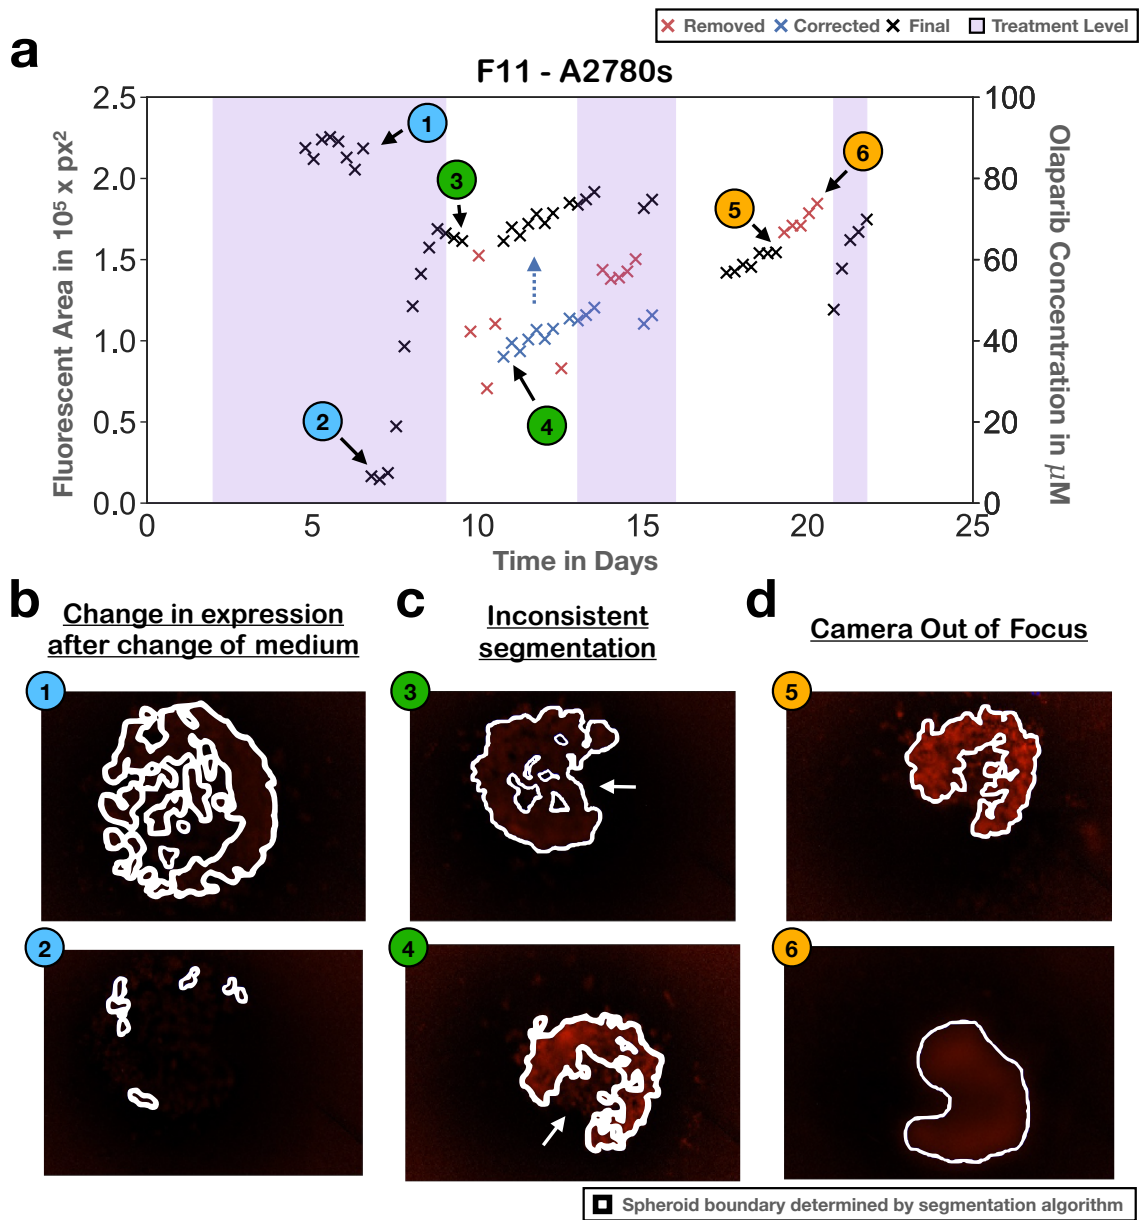


Figure 4.12: Problems encountered during image processing. (a) Tumour growth trajectory as measured by fluorescent area, computed from the images. Black crosses indicate data used for model fitting, red crosses denote values excluded, and blue crosses mark raw values which were corrected (to black). (b) Temporary drop in fluorescence after a growth medium change. (c) Improved image quality after a medium change decreases the reported tumour size (for inferred size see (a)). This is because the outline of the spheroid is now segmented more accurately (arrows). (d) A spheroid moving out of the plane of focus after a medium change results in a blurry image and an inflated reported size.

growth dynamics of the spheroids, and occurred at slightly different speeds for each spheroid (see also Figure C.3). More problematic discontinuities were caused when the spheroid was accidentally moved when aspirating the medium, so that it was now imaged in a different orientation or from a different plane. This could result in the analysis pipeline falsely reporting decreases or increases in tumour size, where there were none. For example, when after a change of medium the image quality increased, the segmentation algorithm could falsely indicate a decrease in tumour size, as it was now able to detect the outline of the spheroid in finer detail (Figure 4.12c). Conversely, if the image was blurred, then this could inflate the reported spheroid size (Figure 4.12d). Depending on how many images were affected, we either excluded misleading (usually poor quality) images, or corrected the fluorescent area values so that they more accurately reflected the size on the image.

Unlike what may be suggested by Figure 4.12, which was chosen for representative purposes, there were relatively few problems overall (see Figure C.3 for an overview of all spheroids). Generally, there were more problems with the A2780s cell line than TOV112d, because of its more diffuse growth pattern. However, in both cases data acquisition became more reliable as the tumour size increased. Where we did have to make corrections, the modified values are typically consistent with the data observed at subsequent times, indicating that our alterations were appropriate (Figure C.3).

4.3.2.6 Model fitting

After having processed the data, we used them to calibrate each model in the ensemble. In trying to keep the complexity of the framework manageable, we decided to drop the measurement model from Section 4.2. Instead we assumed that the observed fluorescent area, $\hat{A}_F(t)$, was simply directly proportional to the total tumour cell density, $N(t) = S(t) + R(t)$. Thus, the predicted fluorescent area, $A_F(t)$, is given by:

$$A_F(t) = \theta(S(t) + R(t)),$$

where θ is a scaling factor, which we allowed to vary to account for differences in the spatial organisation and fluorescence between spheroids. We justify the choice of the linear model in three ways. Firstly, it is still able to fit the data. Secondly, a linear model is often and successfully used in the literature, even when the relationship between the measurement and the cell density is likely more complex (e.g. [148, 212, 114]). Finally, adding a non-linear measurement model may blur the differences between the different growth laws in the ensemble.

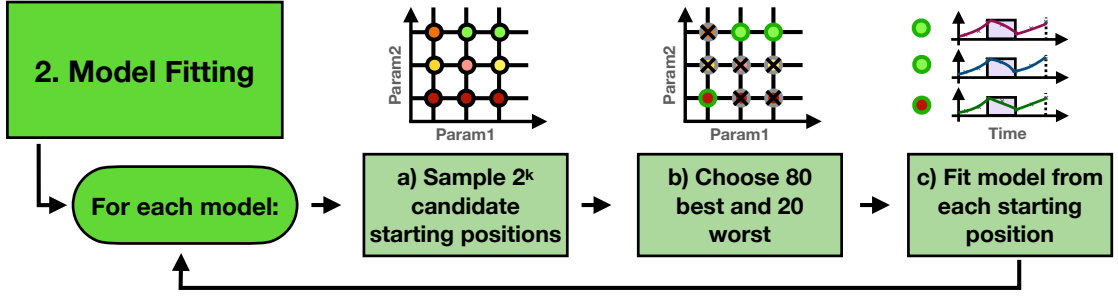


Figure 4.13: Overview of the steps involved in model fitting. Each of the six models was fitted 100 times from different initial parameter guesses, which were selected using an importance sampling based approach.

We fitted the model by minimising the RMSE between the observed fluorescent area, $\hat{A}_F(t)$, and the predicted area $A_F(t)$. To do so we employed the `least_squares()` routine in Scipy [179], which uses a trust region method for minimisation. In preliminary experiments we found this method to be faster performing than the basin-hopping routine we used in Section 4.2 (not shown). Default optimisation options were used for the `least_squares()` method.

For each model we fitted all parameters, in addition to the scaling factor, θ . Because we did not have data available from the first days of the experiment, we found that the fitting routine struggled to delineate the impact of θ from that of the initial conditions. We therefore constrained the initial conditions so that the initial cell number equalled that which was seeded, $S_0 + R_0 = 5000$, and only allowed S_0 to vary. The bounds imposed on each parameter were the same for all models, and were decided on in consultation with Professor Damaghi. A summary can be found in Table 4.4.

In order to account for the possibility that there may be multiple parameter combinations which yield equally good fits, and in order to capture some of the uncertainty in the parameter estimation, we fitted each of the six models 100 times from different initial starting positions. Subsequently all 600 fits were taken forward to the treatment prediction step.

Table 4.4: Bounds on the model parameters when fitting the models in the ensemble (Table 4.3). The same bounds were used for all models. *Note: R_0 was determined from S_0 as $R_0 = 5000 - S_0$.

Bound	r_S	r_R	K	d_D	α	β	S_0	R_0^*	θ
Lower	0	0	0	0	0	10^4	0	0	10
Upper	10^{-2}	10^{-2}	10^{-2}	10^{-2}	10^{-2}	10^7	5×10^3	5×10^3	200

The challenge in selecting the 100 guesses from which to start the optimisation, is to choose these in a fashion which provides as good a coverage of the parameter space as possible. In Section 4.2 we selected these randomly, but this carries the risk that some regions of parameter space with good fits may be completely missed, whereas other, poor fitting regions may be over-represented. An alternative to this approach is to take the guesses from an even-spaced grid in parameter space. But given the dimensionality of our problems (the minimum number of parameters to be fitted was five), this still typically produced more candidates than could be examined. As a workaround, we adopted an approach inspired by importance sampling (Figure 4.13; see Lambert [213] for an introduction to importance sampling). In the first step we probed the parameter space by evaluating the fit of the model from all grid points on an equi-spaced grid spanning the space. We constructed this grid by taking a sample at each first and second tertile of each parameter value ($1/3$ and $2/3$ the value along its range). This gave two possible values per parameter, and a total of 2^k candidate initial conditions, where k is the number of parameters to be fitted. Next, from these 2^k conditions we chose those 80 with the best fit and 20 with the worst fit to serve as initial conditions for the fitting. The motivation to include the 20 poor fits was so as to find possible alternative optima. If $2^k < 100$, then we used all points from the grid, and chose any extra required guesses by repeating the points from the grid in order of their goodness-of-fit. The parametrisation of this algorithm (number of samples per parameter, proportions of good and bad candidates) were chosen in preliminary experiments as a trade-off between balancing exploration of the parameter space and computational cost.

4.3.2.7 Treatment prediction

In the next step, we used the 600 fits to assess different treatment options for the upcoming treatment interval. In developing this part of the framework there were three key design choices to make. Firstly, we had to decide over what time frame to optimise treatment, the so-called *optimisation horizon*, \mathcal{H} . If this horizon is too short for resistance to arise, then the best choice will always be to treat as harshly as possible, as this will minimise tumour size in the short term. Conversely, the further we predict into the future, the better we can leverage the ability of the models to foresee the possible emergence of resistance, but also the greater the computational burden of the optimisation problem. After discussions with Professor Damaghi we settled for an optimisation horizon of $\mathcal{H} = 14\text{d}$ (four treatment intervals), as a biologically reasonable time frame which was still computationally feasible.

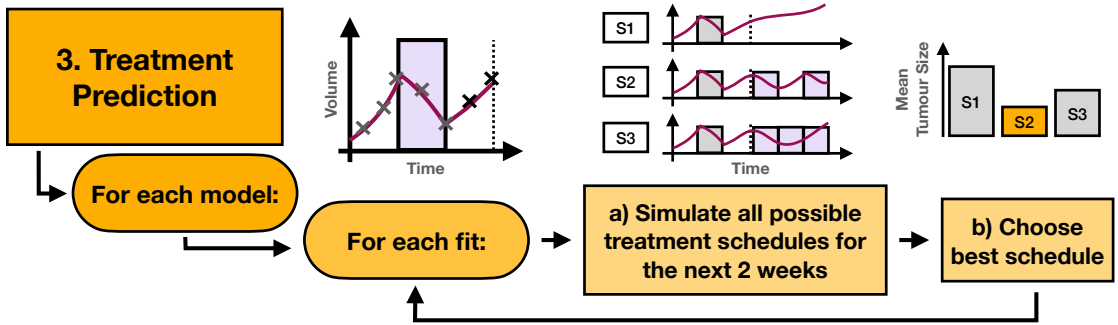


Figure 4.14: Overview of the steps involved in assessing possible treatment schedules for the upcoming treatment interval. For each model and fit (600 total) all possible vacation-based treatment schedules were simulated ($D(t) \in \{0, 100\} \mu\text{M}$), and evaluated according to time-averaged spheroid size (Equation (4.10)).

The second question to address was the choice of objective function. Section 4.2 showed that optimising the final spheroid size will select for schedules which intensify treatment towards the end of the optimisation horizon. This was true even when we added the time-averaged spheroid size as a regularisation term. The success of late intensification relies on the fact that the experiment finishes before the resistance it releases becomes a problem. However, this approach is problematic in a closed-loop setting, which does not finish at the end of the optimisation horizon. To avoid misguided treatment intensification, we therefore decided to optimise schedules not by the final, but by the time-averaged, size of the spheroid:

$$\mathcal{J}(\vec{x}) = \frac{1}{\mathcal{H}} \int_{t=t_n}^{t_n+\mathcal{H}} N(t; \vec{x}) dt \quad (4.10)$$

where \vec{x} is the treatment schedule, and t_n is the current treatment day.

Finally, we had to decide on the complexity of schedules to explore. Again, the timing of treatment was predetermined by the experimental protocol, so that each schedule was simply a vector of $\vec{x} \in (\{D_n\})_{n=0}^3$, where D_n is the treatment dose to give $n = 0, 1, 2, 3$ treatment days from today. We decided to constrain ourselves to vacation-based schedules, so that $D_n \in \{0, 100\} \mu\text{M}$, which we did for two reasons. Firstly, it significantly reduced the size of the search space to 16 candidates, so that we could avoid heuristic optimisation, and could instead examine all possible treatment schedules. Secondly, for simplicity, the models in the ensemble assumed a linear dose-response relationship. While this could capture switching between two drug doses, it was likely to be inaccurate for predicting the effect of intermediate doses, as dose-response curves are typically non-linear [214].

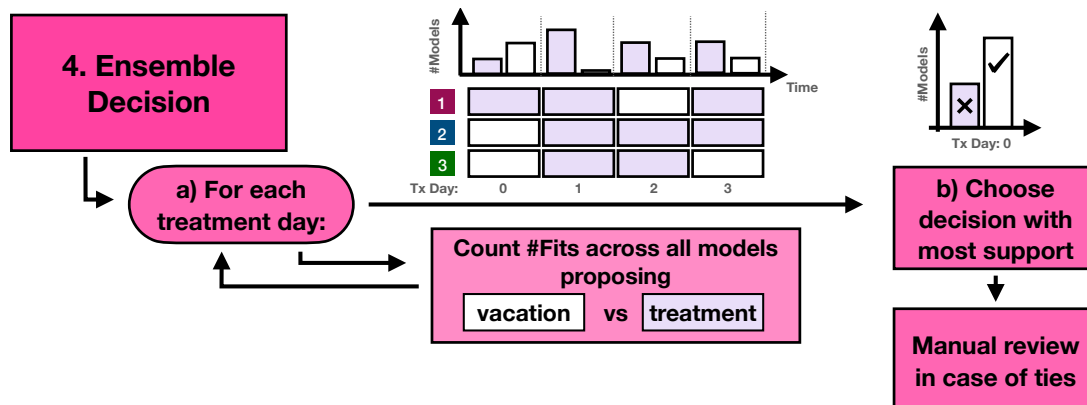


Figure 4.15: The algorithm used for assembling the decision from the ensemble. We chose a simple “majority rule” approach, in which the option which was recommended by most of the fits was the one selected. In the case of ties, we manually reviewed the evidence for each choice to make the decision.

To sum up, on each treatment day we leveraged the 600 model fits to simulate all 16 possible vacation-based schedules for the next two weeks (Figure 4.14). Based on these predictions, we ranked the schedules according to the predicted mean tumour size across the 14 day period, and selected the top schedule from each fit. In the final step of our pipeline we then integrated these 600 recommendations to decide what dose to administer.

4.3.2.8 Ensemble decision-making

The motivation to use an ensemble of different mathematical models was to gain an understanding of the robustness of the predicted treatment schedule. In the final step, we therefore compared the treatment recommendations made by each of the 600 fits (Figure 4.15). In the interests of simplicity, we employed a simple “majority rule” approach. To do so, we counted for each of the four treatment days in the two week optimisation window, how many of the fits recommended treatment that day ($D_n = 100\mu\text{M}$) compared to a treatment vacation ($D_n = 0\mu\text{M}$). We then based our treatment decision for that day on which of the two options was supported by a greater number of fits. In the case of ties, we manually reviewed the evidence for each choice. To do so, we followed the principle of AT and, unless the spheroid was growing, or there was a significant difference in the goodness-of-fit between the models advocating breaks or treatment, we selected breaks. This occurred a total of six times during the experiment: four times in spheroid D5 (TOV112d), and twice in spheroid G11 (A2780s; see also Figure 4.19).

4.3.2.9 Statistical analyses

In order to compare the performance of our framework to that of the other treatment strategies, we obtained and segmented the images using the image analysis pipeline. However, we did not have the resources to manually check these images in the same way in which we had checked those for the model-guided strategy. To avoid a bias in the analysis due to the curation, we therefore used the raw, uncurated segmentations when comparing the model-guided strategy to other schedules. All statistical analyses were carried out in Python 3.6, using Scipy 1.5.1 [179] and statsmodels 0.9.0 [215].

4.3.3 Results

4.3.3.1 An example of the framework in action

The aim of this work was to test, in an *in vitro* setting, whether we could use an ensemble of mathematical models to guide individualised, AT decision-making. In Figure 4.16 we illustrate the application of our framework to one spheroid. To begin with, the spheroid was subjected to a lead-in treatment in order to obtain calibration data for the models. During this time we see an expansion when treatment is withdrawn and a contraction when it is re-applied (Figures 4.16a & b). This indicates that the spheroid can be successfully re-challenged after a break, which is encouraging for AT.

After the lead-in period, we switched to model-guided AT to determine the next drug dose. Following data processing, we fitted each of the six models 100 times to the data, which yielded the fits shown in Figure 4.16c. There was close agreement between the different fits of the same model for all but the E-NS model, indicating that there was a single, well-defined optimum in each case. Moreover, all models concurred in their prediction that the spheroid was made up almost entirely of sensitive cells, although the E-NS model predicted more resistance than the other models (Figure 4.16c).

Aside from the inferred resistance fractions, there were also other differences between the models (Figure 4.16c). Firstly, the exponential and logistic models predicted a slightly larger initial spheroid size than the spherical models. One reason for this uncertainty in the initial conditions was the paucity of data from the first days of the experiment. This was because Professor Damaghi only started imaging the cells on Day 4 of the experiment, as the machine was not available earlier. In addition, the image quality was initially poor, so that reliable data for this spheroid were available

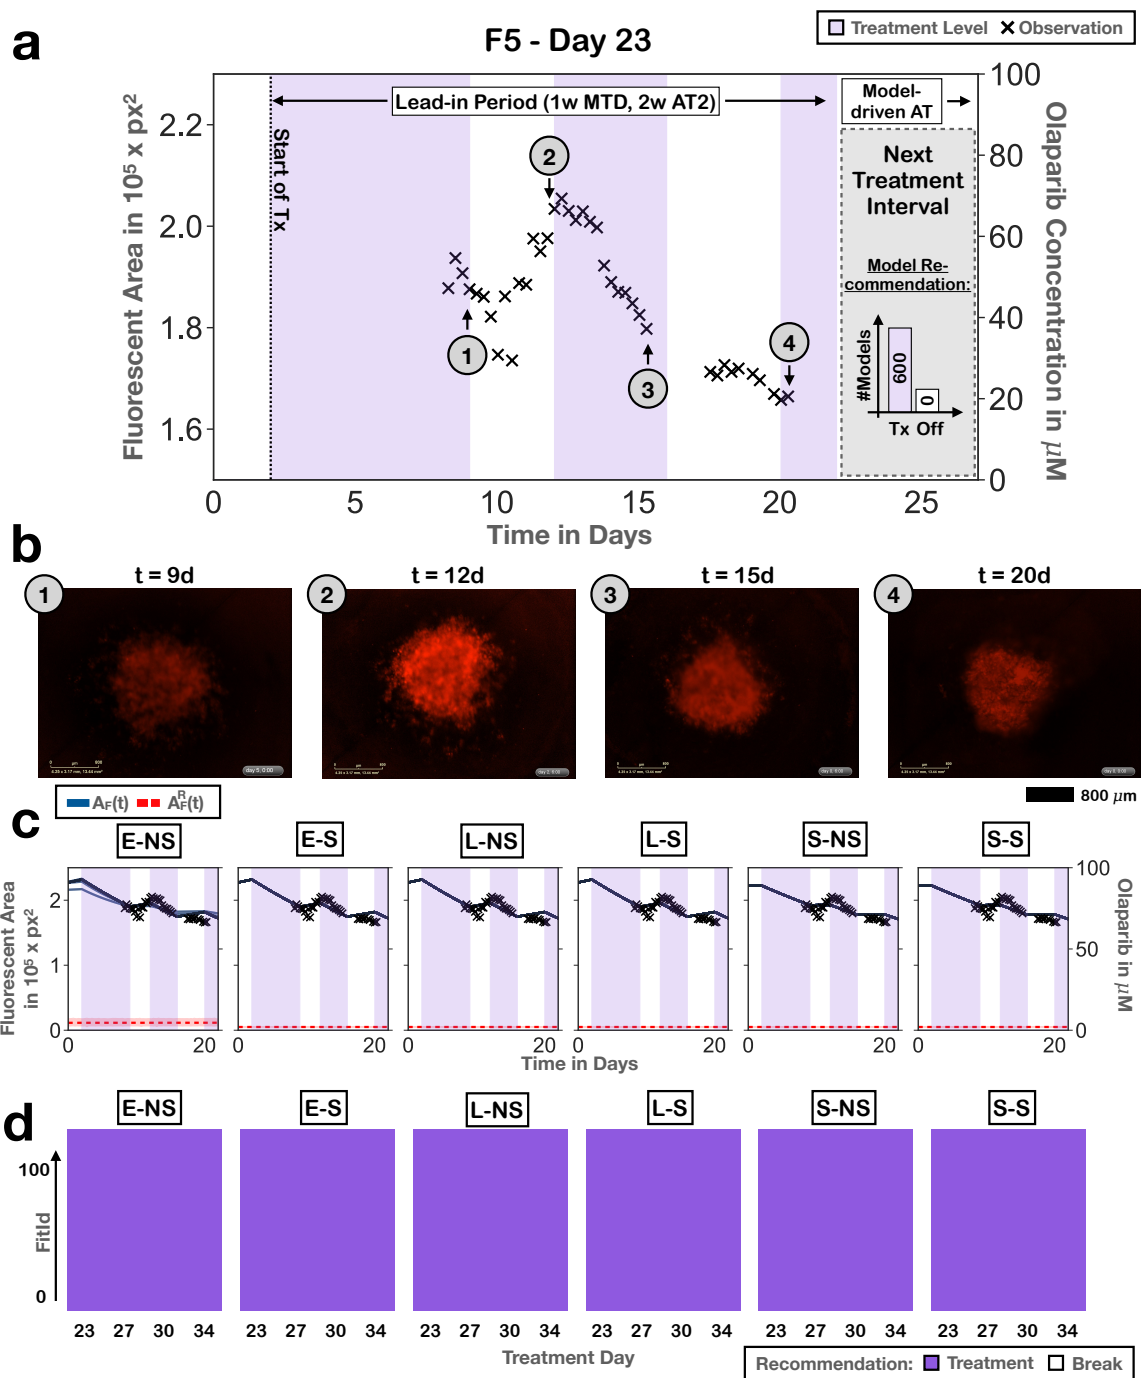


Figure 4.16 (*previous page*): Example of using the ensemble framework to guide individualised treatment decisions. **(a)** Tumour growth trajectory up to the current day (here Day 23). As it was the first time the framework was used, only data from the lead-in phase were available. The grey box depicts the upcoming treatment interval for which a decision was required, together with the framework’s recommendation to treat. **(b)** Representative images of the spheroid over time. **(c)** Fits of each of the different models to the data (solid, blue lines), together with inferred size of the resistant population (dashed, red lines). Each line represents one of the 100 fits, with the hue of the colour representing the fit’s index (bright: 0, dark: 100; difficult to see here because all fits concurred closely). Shading marks a 95% CI. Generally there was close agreement on spheroid growth dynamics and resistant fractions between the models and the runs of each model. **(d)** Optimal treatment schedule for the next 14 days according to each of the 100 fits for each of the six models. As all recommended treatment, it was decided to treat that day.

only from Day 8 onwards. Thus, aside from the knowledge that a total of 5000 cells were seeded, there was little to constrain the initial conditions with.

A second difference could be seen in the predicted rate of re-growth during treatment breaks (Figure 4.16c). As would be expected, the exponential models (E-NS and E-S) predicted the fastest, and the spherical model (S-NS and S-S) the slowest, re-growth. That being said, all models over-predicted the growth seen during the second treatment break, suggesting that none of them could fully explain the data.

Subsequently, we used each of the 600 fits to determine which combination of treatment and vacation was predicted to minimise the mean tumour size over the next two weeks. In this case, all fits agreed that the optimal strategy was to treat at MTD for all four upcoming treatment intervals (Figure 4.16d). This held true even for those E-NS fits which predicted a larger fraction of resistance. Given the strong concordance between the model predictions, we decided to treat this spheroid that day.

4.3.3.2 An example of disagreement in the model predictions

However, it was not always the case that the models agree closely. For example, in Figure 4.17 we show the model predictions for a different spheroid from the same cell line and from the same day. While this spheroid also responded to the re-challenge during the lead-in phase, its response was slightly weaker and it remained on treatment until the end of the lead-in period (Figures 4.17a & b). Interestingly, in this case all models indicated a large fraction of resistant cells within the spheroid, although the estimates for the size of the resistant population varied both within, and between,

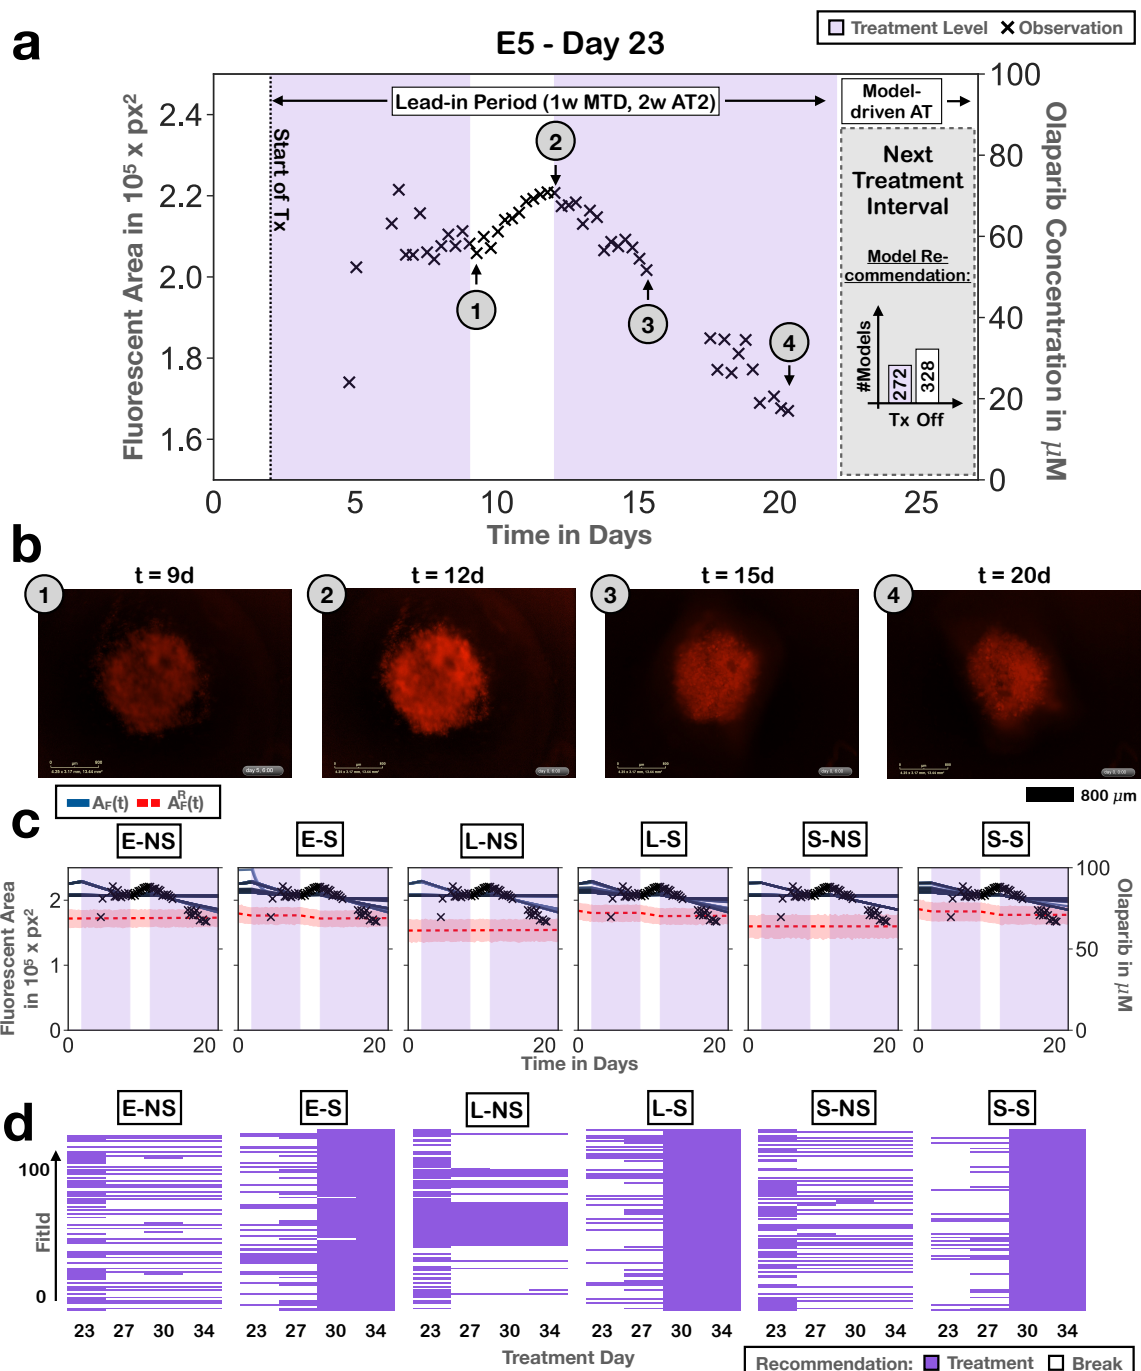


Figure 4.17 (*previous page*): Example of disagreement on the recommended treatment strategy between models. **(a)** Tumour growth trajectory up to the current treatment day (Day 23). The grey box depicts the upcoming treatment interval for which a decision was required, together with the framework’s recommendation. In this case, there was strong disagreement between whether one should treat or not. **(b)** Representative images of the spheroid over time. **(c)** Fits of each of the different models to the data (solid, blue lines), together with inferred size of the resistant population (dashed, red lines). Each line represents one of the 100 fits, with the hue of the colour representing the fit’s index (bright: 0, dark: 100). Shading marks a 95% CI. All models predicted a substantial resistance fraction, although the estimate for its size and dynamics varied between models. **(d)** Optimal treatment schedule for the next 14 days according to each of the 100 fits for each of the six models. There was strong disagreement, with non-switching models advocating treatment followed by breaks, whereas switching models recommended the opposite. Given the continued drug response of the spheroid in (a), it was decided to follow the majority recommendation and withhold treatment.

different models (Figure 4.17c). We observe that models with switching tended to infer a slightly higher resistance fraction than the corresponding non-switching model (e.g. S-NS vs S-S). In addition, the switching models predicted a re-sensitisation of the spheroid during the treatment break, whereas in the non-switching models the break was not predicted to have much of an effect. Accordingly, there were conflicting treatment recommendations, with switching models suggesting treatment breaks followed by a re-challenge, whereas non-switching models recommended the opposite approach (Figure 4.17 d). Nevertheless, given that a slight majority of models recommended a break (328 vs 272), and that between Days 15 and 20 the spheroid had continued to respond well to treatment, we decided to withhold treatment in the coming interval.

4.3.3.3 Model-guided AT can match, and even outperform, rule-of-thumb based AT algorithms

By applying our framework as illustrated in Figures 4.16 & 4.17, we guided treatment for 16 spheroids (8 per cell line) for a total of 27 days. In this process we made a total of 8 decisions per spheroid. To assess the performance of our individualised approach we compared it to CT, as well as AT1 and AT2 (Figures 4.18a & b). This showed that our framework outperformed AT2, producing both smaller mean tumour sizes and a smaller final tumour size (Table C.1). Also the comparison to AT1 was favourable, especially in the A2780s cell line (Figure 4.18a). That being said, the sample size

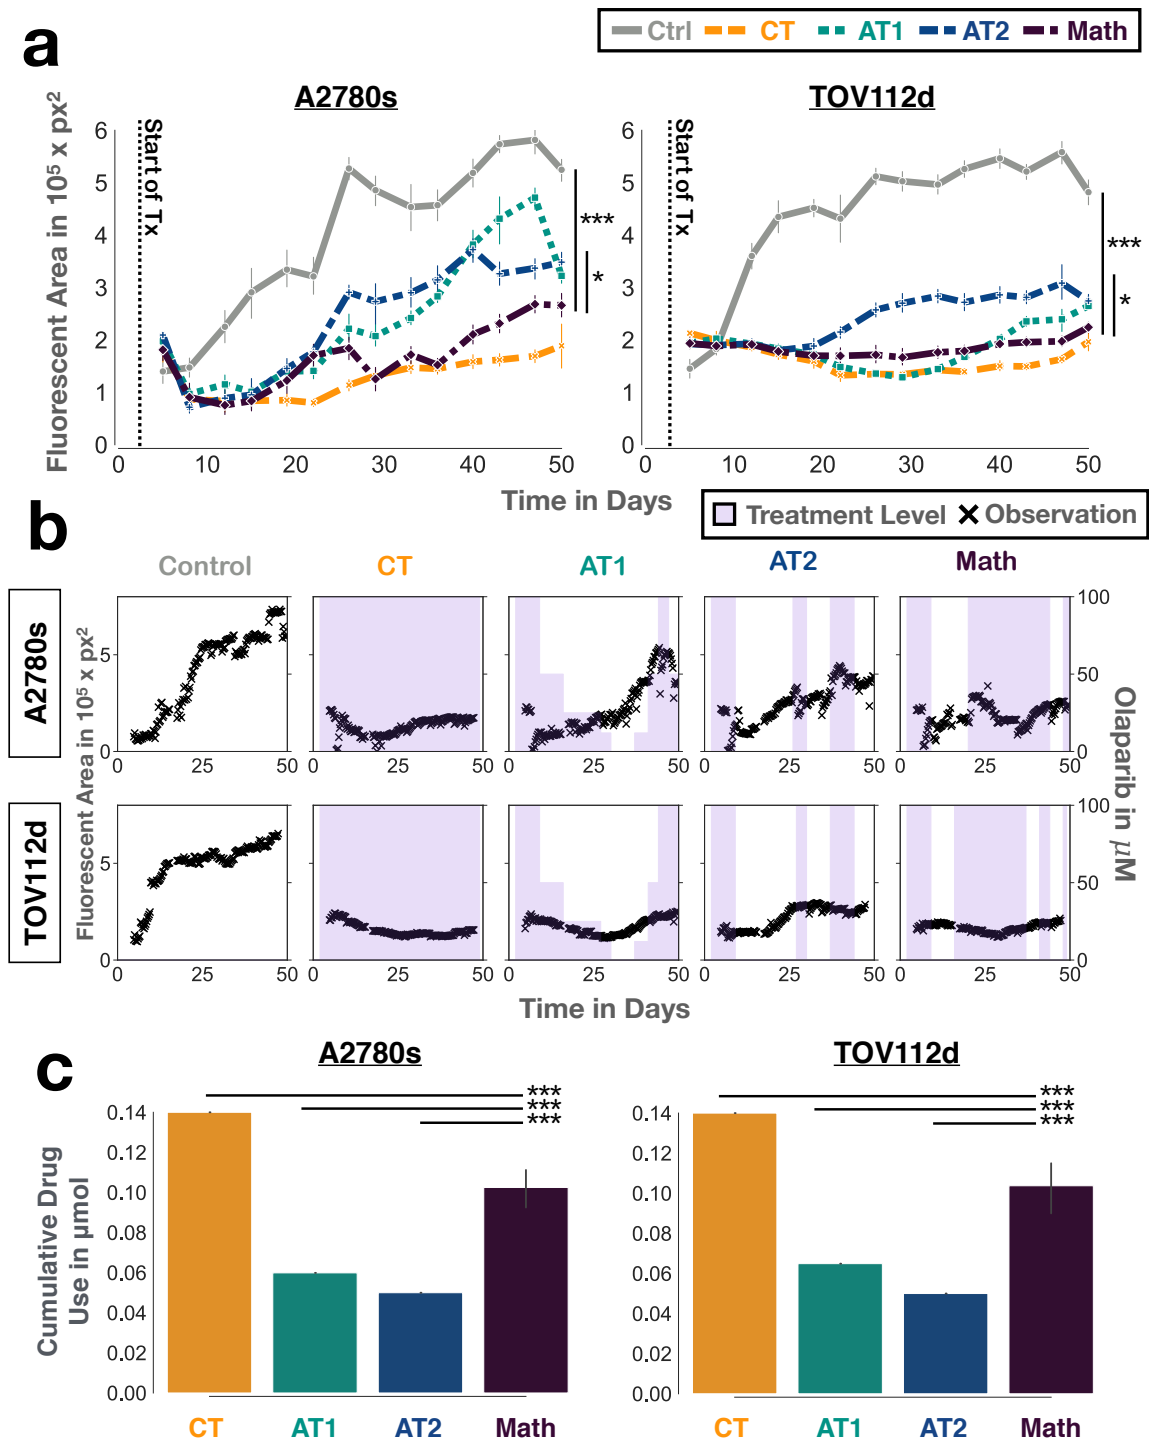


Figure 4.18: Comparison of model-guided AT with CT, AT1, and AT2 *in vitro*. The ensemble algorithm (Math) performed as well, or better, than AT1 and AT2, but poorer than CT (see also Table C.1). * and *** indicate statistical significance at the 95% and 99% confidence level, respectively. Error bars denote 95% CIs. (a) Spheroid growth over time under the different treatment schedules. Only the measurements at the end of each treatment interval are shown. (b) Example trajectories of individual spheroids from the experiment. (c) Cumulative drug use by the different treatment schedules.

of the AT1 branch was small (4 spheroids), and the differences were not statistically significant (Table C.1), so it was not possible to draw definite conclusions. Nevertheless, these results are promising, and indicate that it is feasible to use mathematical modelling to guide AT decision making, and that doing so may produce better results than simple “rule-of-thumb” algorithms such as AT1 or AT2 (see also Figure C.4 for an overview of all spheroids).

4.3.3.4 Not adaptive, but continuous, PARPi administration maximises tumour control

A somewhat unexpected result from our experiment was that the schedule that best controlled the spheroids overall was CT (Figures 4.18a & b). This was contrary to our hypothesis that CT would cause competitive release of resistant cells. While spheroids under CT began to regrow, they did so at a very slow rate. In line with this we note that the ranking of the schedules correlates with the cumulative amount of drug administered (Figure 4.18c). CT, which performed best, administered the most drug, followed by the model-guided schedule, AT1, and finally AT2.

4.3.3.5 Model-guided AT adaptively intensifies treatment

In the light of this observation, we note that the model-guided strategy administered the most drug out of all adaptive strategies (Figure 4.18c). Moreover, when we examine the schedules produced by our framework in more detail, we see that therapy typically begins with a long period of CT (Figures 4.19a & b). This suggests that the reason for the success of the model-driven strategy was not that it leveraged competition more effectively than AT1 or AT2. Rather, these results indicate that it was successful because it detected that treatment breaks did not have the desired effect, and that a treatment-intense approach would be more successful.

Nevertheless, while the model-guided AT administered more intensive treatment, it still did so in an adaptive fashion. Firstly, the more slowly growing TOV112d spheroids saw earlier treatment withdrawal than the faster growing A2780s spheroids (Figures 4.19a & b). In addition, there were differences between spheroids of the same cell lines. For example, A5 was treated continuously, whereas B5 was treated only sporadically (Figure 4.19b). We also note that even though the model-guided schedule administered more treatment than the other adaptive schedules, it still used significantly less drug than CT (Figure 4.18c).

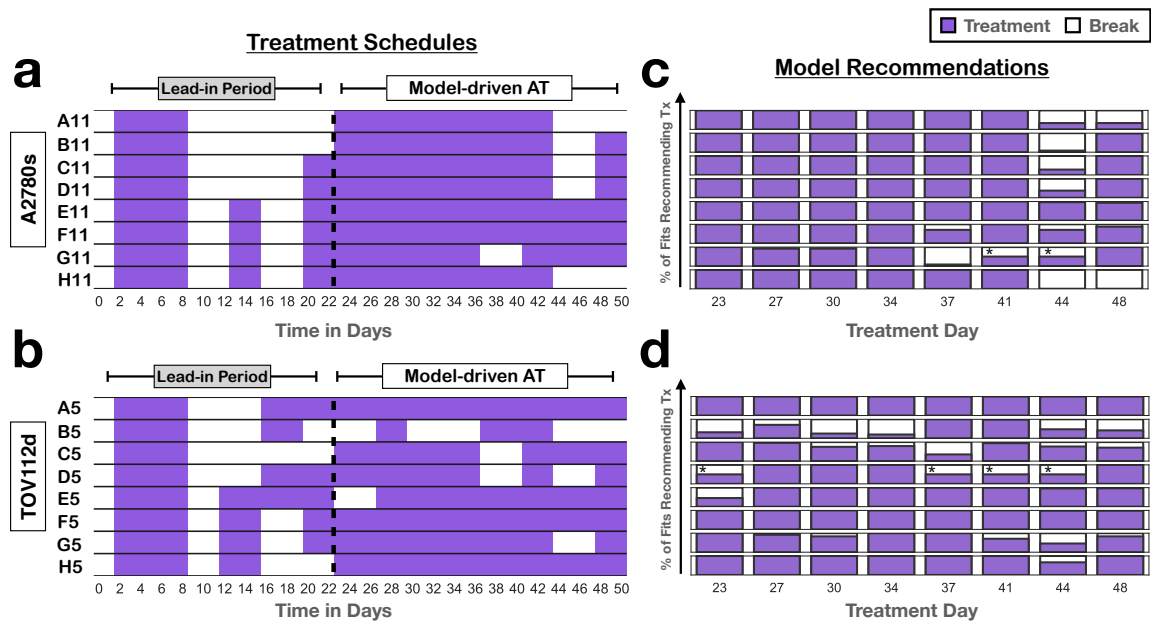


Figure 4.19: The treatment schedules produced by the model-guided AT framework. (a) & (b) Schedules for the A2780s and TOV112d spheroids, respectively. (c) & (d) Model recommendations for each spheroid for the two cell lines, respectively. * denotes ties. In line with the superior performance of CT in the experiment, the model-guided AT treated more intensely than the other adaptive strategies. That being said, the degree of intensification was specific to each spheroid. For a plot of the growth dynamics of each spheroid over time, see Figure C.3.

4.3.3.6 While in some spheroids there is strong concordance between models, in others there is continued disagreement

One aim of using an ensemble of different models was to try to quantify the confidence in treatment decisions. The more models agreed with a decision, the more confident we could be that it would be the most effective action to take at this point in time. In concordance with the observation that CT performed well, we find that the recommendation to treat was regularly supported by all models (Figures 4.19c & d). Conversely, the decision to withhold treatment was typically contested. Only twice did all models recommend a break (Spheroid H11). Moreover, we also here observed heterogeneity between the cell lines and individual spheroids. Whereas in the A2780s spheroids there was typically strong agreement between the models about how to treat at each step (e.g. A11 or H11), in the TOV112 cells there was more disagreement (e.g. B5 or D5). To sum up, this illustrates how examining not just a single model, but an ensemble of different models, may potentially provide warning signs, when a treatment approach is not working as planned.

4.3.4 Discussion

In this section we presented proof-of-principle data on using an ensemble of models to drive AT decision-making. Importantly, not only did we develop a theoretical pipeline, but we also tested it on a real-life *in vitro* system. Our framework fitted six different ODE models to each spheroid on every treatment day, and subsequently used each of the 600 fits to predict the best vacation-based treatment strategy for the upcoming two week time period. We find that our algorithm delivers individualised treatment schedules which are unique to each spheroid and which, in a direct comparison with the rule-of-thumb AT1 and AT2 algorithms, perform favourably.

To the best of our knowledge, this is the first time that an ensemble of different models has been used to guide cancer treatment decisions. We assembled our ensemble by varying two assumptions which Chapters 2 & 3 showed to be key for AT: i) the growth law which describes the strength of intra-tumoral competition, and ii) the rate of switching between sensitive and resistant cells. We find that the models could fit the data, and provided scrutiny of the treatment decisions. While in some cases all models agreed with a treatment decision, in other cases there was discord.

Interestingly, the decisions which tended to be contested were those on administering treatment breaks. While this was contrary to our hypothesis that treatment de-escalation would help with long-term tumour control, it does match our *in vitro* result that not AT, but CT, is the best performing strategy in this case. This indicates an unexpected, but important, potential benefit of model-driven AT: if designed correctly, then such a framework could act as an early-warning system of when *not* to proceed with AT. That being said, at this point we do not have sufficient evidence to conclude whether the algorithm tested here “detected” that CT was the more advantageous route, or whether its parametrisation biased it towards more treatment-intense schedules. In subsequent extensions, it would be important to test our framework on an experimental model system in which AT is known to be superior to CT.

However, before proceeding with more practical testing of the framework, this pilot experiment has shown that there are a number of issues which should be addressed first. While the models can fit some of the dynamics, they are unable to fully describe the spheroid’s treatment response, as shown in Figures 4.16c & 4.17c. Also predictions, such as those in Figure 4.17c, that a majority of the spheroid is resistant, seem unlikely to be true. Furthermore, a method should be developed to account for the discontinuities in the growth curve caused by medium replenishment, so that the model can distinguish these more reliably from treatment-induced changes

in spheroid size. Perhaps a Hidden Markov Model [216] might provide a way to more accurately connect predicted and observed spheroid sizes.

Moreover, a parameter sensitivity analysis should be carried out to better understand the confidence with which the different model parameters may be inferred. In particular, focus should be put on how to infer the initial resistance fraction and the switching rates, which are important factors in deciding a treatment strategy. Greene et al [146] have recently proposed an experimental protocol based on an identifiability analysis on a model similar to Equations (4.3)-(4.4). It would be useful to incorporate their work, also with a view towards identifying a schedule for the lead-in phase which provides as much information about the tumour as possible, whilst not foregoing the chance of a cure or of tumour control.

A further issue which should be addressed, is the timing of data acquisition. Because of the delay caused by the decrypting process, we did not know the current size of the spheroid, when making a treatment decision. Nor did we know its response to the most recent treatment decision. This meant that we could not incorporate the models' predicted power into the decision-making, as would be done in weather forecasting. There are two ways to address this issue: Firstly, the logistics should be streamlined to avoid costly data transfer. Secondly, if lags remain, the timing of the decision-making should be aligned with that of data acquisition. For example, we could have updated our decisions every two treatment intervals, so that at least the feedback on our first decision was available before making the next one.

Finally, an important question which warrants further investigation is how the predictions of the different models in the ensemble should best be integrated. As a proof-of-concept, we chose here a simple, majority-rule approach. However, our work shows a number of pitfalls of this strategy. Firstly, it does not take into account how well each of the models fit, or predict the data. Even poor fits contribute equally. Secondly, there is no memory of prior decisions in our framework which is problematic in pursuing more long-term strategies. For example, while the switching models in Figure 4.17 recommend a break for the upcoming two treatment intervals, they subsequently advise to re-apply treatment. How can we ensure that such longer-term plans are remembered in making subsequent decisions? Also, what should be the relationship between the time horizon over which predictions are made, and over which they are applied? Yeang and Beckmann [190] have shown that longer prediction horizons yield better outcomes in their MPC framework. But the longer the horizon, the greater also the computational burden of the optimisation problem [190].

Future research should investigate this problem, drawing from experience with closed-loop control systems in other areas of medicine, such as glucose control (the “virtual pancreas” [217, 218]) or anaesthesia [219].

4.4 Concluding remarks

This collaboration provided us with valuable insights not only into how to improve AT decision-making, but also into the challenges of translating treatment plans from the screen to the lab bench. To conclude, we want to reflect on this unique and formative experience.

We identify three main challenges. Firstly, we observed significant heterogeneity in treatment responses between individual spheroids. This means that the parameters in the model fits varied significantly, and what was predicted to be the best way to treat one spheroid was not necessarily the best for another. While the field has developed great expertise in optimising the schedule for an individual parameter set, little is known about how to choose treatment schedules in light of this heterogeneity. How robust are our schedules to such heterogeneity? Moreover, should we be optimising the mean outcome, or should we be trying to reduce the extremes? Virtual cohort approaches, which have become widely used in industrial PK/PD modelling, may provide a tool kit to address these questions [79, 208].

A second important question is how do we respond to unforeseen interruptions, and how do we proceed at the end of the optimisation period? For example, what do we do if there is a data lag, or loss, as we saw in the ensemble experiment? Or if treatment has to be given a day late or early, due to, for example, a holiday (or illness in case of a patient)? These are problems not typically encountered when optimising treatment *in silico*, but which play an important role as we move from the screen, to the bench, or even bedside. It may be possible to learn strategies from the control theory literature to address these problems, especially from closed-loop control.

Thirdly, while we have tight control over the mechanisms driving treatment response in our models, this is not the case in real biological systems. Even in a well controlled experimental system, we can never be sure that our model will capture all relevant biology. As such, it is important to investigate how we can address this uncertainty. Here we proposed an ensemble approach, inspired by ideas from weather forecasting. While our framework has a number of limitations, our results are promising, and we believe that they warrant further investigation. That being said, the *in vitro* experiments are expensive and take a long time to run, which makes

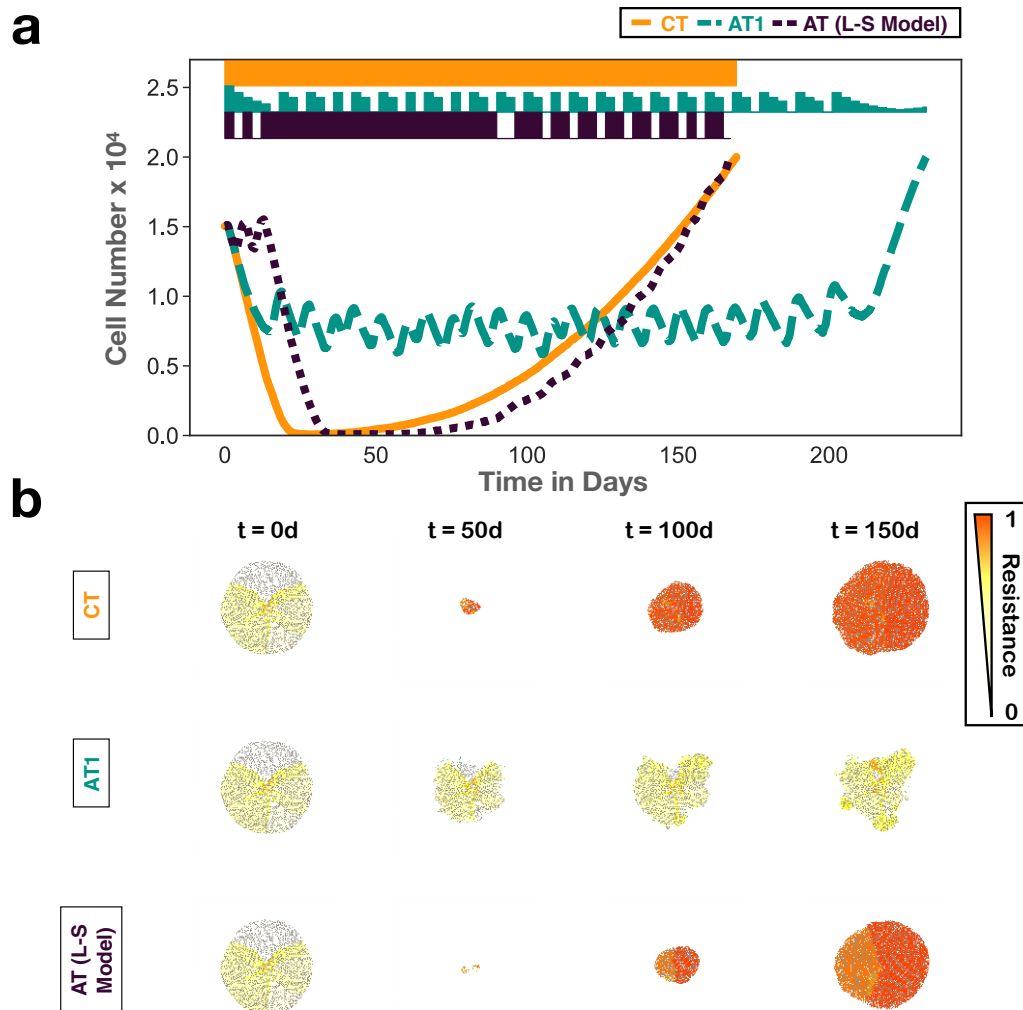


Figure 4.20: While essential, the *in vitro* experiments are time-consuming and costly. As a work-around, we propose to next fine-tune our framework on other *in silico* models before re-applying it *in vitro*. This will also allow us to explore how well the framework would work if there was a mis-match between the models in the ensemble and the tumour biology. As a proof-of-concept, we here present preliminary work applying the framework to treat the ABM by Gallaher et al [124]. (a) Growth trajectories of CT, AT1, and the model-driven strategy. Bars denote the treatment schedule. (b) Snapshots from the simulations in (a). Here we only used 10 fits of the L-S model to form the ensemble, which resulted in over-treatment of the tumour. Next, it would be interesting to test whether the ensemble could overcome this problem.

it challenging to develop our framework further. As a work-around, we have started experimenting with optimising it *in silico* first, by applying it to other tumour models. This allows us to run experiments in a number of days rather than weeks, and it allows us to investigate what is the impact of having models in the ensemble that do

not represent the true underlying biology. As an example, in Figure 4.20 we use the L-S model to guide treatment of the ABM by Gallaher et al [124]. The ABM assumes a spectrum of drug resistance, which is pre-existing, and which correlates inversely with the cells' division rates (see also Section 3.1.1). We find that the model-driven AT performs similarly to CT, whilst using less drug, but worse than AT1, as it administers treatment breaks too late. In the next step, it would be interesting to test whether using the full ensemble of models can help to overcome this problem.

Furthermore, trialling model-derived schedules in the wet lab also comes with experimental challenges. One needs to select a model system in which one can reliably distinguish between the perturbation caused by the administration of treatment, and the treatment effect. Moreover, for adaptive approaches, one requires timely and reliable data. These are both points which we underestimated, and which are a weakness of the work in this chapter. While the *in vitro* spheroid model provided a good trade-off between realism and costs, changing the growth medium caused large perturbations which confounded the true treatment effect. Animal experiments are a more robust alternative, but are more costly and allow for less insight into the intratumoral dynamics. Perhaps flow chamber systems [220], could provide work-arounds in the future.

Finally, there are the challenges of building inter-disciplinary collaborations. For example, while ideally we would have liked to work with an experimental model system which was well understood, and for which plenty of data were available on which to build a model, this was not what our collaborator was working on. Similarly, while we would have liked to collect more calibration data and explore our models in more detail before applying them, this would have been beyond the time scale of our collaborator's grant. We could have waited for a better opportunity to arise, but it is questionable whether one would have appeared. By seizing this opportunity, we learned valuable lessons, and established a connection which has allowed us to have a more active say in the subsequent follow-up work.

In the next chapter, we will present this follow-up work. While we have focussed so far on the merits of different AT strategies, in almost all experiments CT performed as well, if not better than, AT. This suggests that treatment de-escalation of PARPi has to be carefully planned. To conclude this thesis, we will investigate why this is the case, and develop a more specific mathematical model to aid with future exploration of PARPi scheduling.

Chapter 5

Integrated experimental and mathematical evaluation of the opportunities for adaptive PARPi treatment

5.1 Introduction

Ovarian cancer is a leading cause of cancer death in women [159]. For stage III or IV disease, which comprises the majority of patients (70%), 5-year survival rates remain below 30% [221]. PARPis have shown great promise, but recurrence remains a significant problem, especially in patients who have failed on first-line therapy [191, 222]. In the previous chapter, we investigated whether AT could help to delay progression compared to CT, which is the current standard-of-care. However, we found that this was not the case, and that CT provided the best tumour size control over the time frame of the *in vitro* experiments. We now investigate why this is the case. Thanks to the collaboration we had established with Professor Damaghi, I was able to receive training in 2-D cell culture techniques in his laboratory. Under his supervision, I carried out three sets of experiments aimed at better understanding PARPi treatment response. Our results show that cells recover quickly once treatment is withdrawn, so that treatment de-escalation has to be carefully planned. To help with this, we leverage our data to develop a mathematical model of PARPi response, and we present preliminary results on how one might address the issues that we encountered with the AT algorithms in the previous chapter. Overall, our work stresses the importance of carefully tailoring the AT algorithm to the clinical setting, and has potential implications for the management of PARPi-related toxicities.

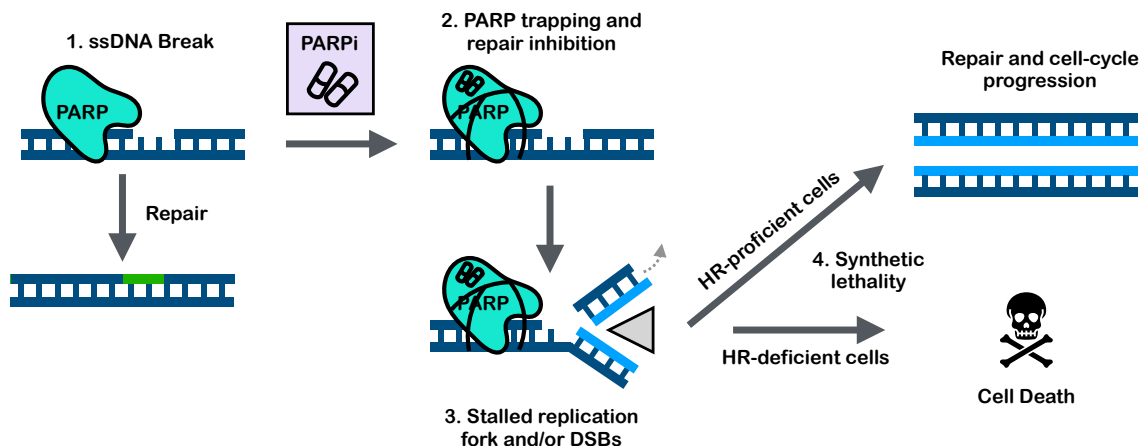


Figure 5.1: The mechanism of PARPi-mediated cytotoxicity. PARPs are important proteins in the repair of single strand DNA (ssDNA) breaks. But PARPis can trap PARPs on the DNA, which results in stalled replication forks and DSBs during DNA replication. Only cells with intact HR pathways can repair this damage, but tumour cells are typically HR-deficient, and are thus killed (see the main text for abbreviations). Figure re-drawn and modified with permission from Noordermeer and van Attikum [222].

5.1.1 Review of PARPi biology and clinical use

As we used relatively generic mathematical models in Chapter 4, we introduced little of the biology underlying PARPi treatment. However, given that even an ensemble of six different models was insufficient to adequately describe the dynamics, this indicates that a more in-depth understanding is required. In the following we will review the key principles of PARPi biology, resistance, and clinical use. Based on this, and the set of experimental results in Section 5.2 we will propose an explanation for the results from the previous chapter, and develop a more accurate mathematical model of PARPi treatment.

5.1.1.1 The biological mechanisms of PARPi action

Tags consisting of short sugar chains, called poly-Adenosine Di-Phosphate Ribose (PAR) chains, are important signals in managing the access to and the organisation of the genome [191]. The family of 17 proteins responsible for attachment of these tags to the genome are known as PAR polymerases (PARPs), and play a key role in DNA repair [223, 191, 222]. PARPs, in particular PARP1 and PARP2, help to detect single strand DNA (ssDNA) damage, and orchestrate the subsequent repair (Figure 5.1).

PARPis inhibit this process and cause cell death through a combination of two processes (Figure 5.1): i) PARP trapping, and ii) synthetic lethality (Figure 5.1). Firstly, PARPis obstruct the catalytic site which is required for releasing PARPs from the DNA [224]. Importantly, this not only prevents repair, but also traps PARPs on the DNA, so that they become an obstacle to further repair, regulation and replication efforts [191]. In particular, if the cell attempts to divide, the replication fork will stall in this position, which causes cell cycle arrest, and can potentially also lead to double strand breaks (DSBs; [223, 191, 222]).

The second important aspect of PARPi action is so-called *synthetic lethality* (Figure 5.1). The cell cycle arrest and DSBs caused by PARP trapping are not per se lethal, but can be repaired via the homologous repair (HR) pathway. However, HR is deficient in many tumours due to, for example, mutations in BRCA1 or BRCA2, which are key orchestrators of HR [195, 194]. As such, HR-deficient (HRD) tumour cells have to rely on more error prone backup mechanisms, and are likely to die from the damage caused by PARPis [223, 191, 222]. It is estimated that around 50% of ovarian cancer patients have mutations affecting HR [225].

5.1.1.2 PARPis in the treatment of ovarian cancer

There are currently three FDA-approved PARPis for ovarian cancer (Olaparib, Rucaparib, Niraparib), which differ mostly in which members of the PARP family they target, and how well they are able to trap PARPs on the DNA [223, 191]. We focussed our work on Olaparib, as it has been approved for the longest, and is the best studied among the four. Olaparib is taken orally for two years after chemotherapy, unless unacceptable toxicity occurs or the tumour progresses [226]. Because it is poorly soluble and is cleared rapidly, with a half life of only 5-7h [227], pills have to be taken twice a day. The MTD is 400mg per administration, although the recommended dose is 300mg due to toxicity [228, 226]. The therapeutic window of Olaparib appears to be quite wide, with reductions in efficacy reported only for doses below 200mg [227, 229, 228]. Most patients (98%) report adverse events, such as nausea, fatigue, and gastro-intestinal problems [17]. Around 40% of patients require dose adjustments due to more serious Grade 3 or 4 events such as anaemia [17].

The benefit of single agent Olaparib maintenance in patients with BRCA mutations has been demonstrated in two major Phase III trials. In newly diagnosed patients, SOLO1 (NCT01844986) showed that Olaparib maintenance for two years after the end of chemotherapy treatment increased median progression free survival (PFS) from 14 months under placebo to 56 months [17, 197]. Similarly, SOLO2




Figure 5.2: The four main mechanisms of PARPi resistance. These are not mutually exclusive, and each may be realised via several biological routes.

(NCT01874353) showed that in patients who were treated for recurrent disease, Olaparib maintenance was able to push back progression from a median PFS of 5.5 months to 19.1 months and, in addition, found a 12.9 month benefit in median overall survival [198, 199].

There are two key biomarkers of PARPi response in patients. Firstly, there is the presence of BRCA mutations or other HRDs. For example, while Lederman et al [196] found a significant PFS benefit of Olaparib for patients with a BRCA mutation, they did not detect a benefit for patients who lacked a BRCA mutation. A second important biomarker is the response to the prior round of platinum-based chemotherapy, as these share common resistance mechanisms [191]. As such, the currently recommended use for Olaparib maintenance is in patients with a BRCA mutation who have shown at least a partial response in the previous round of platinum-based chemotherapy [226].

5.1.1.3 PARPi resistance

Despite the success of PARPis in prolonging TTP, the majority of patients still see their disease returning within five years, either whilst under PARPi treatment, or after maintenance has stopped [197, 199]. Four main types of resistance mechanisms have been identified, which interfere with PARPi action at different points during its mode of action (Figure 5.2). For brevity, we will only outline these here, and we refer to the excellent reviews by Noordermeer and van Attikum [222] and by d'Andrea [223] for more biological detail.

Firstly, cells can up-regulate efflux pumps, such as the ABC transporters, which allows them to clear drug from the cell. Secondly, the efficacy of PARP trapping can become reduced, either through mutations in PARP1/2, or by changes in the pathways which are responsible for triggering PARP release (loss of PARG). Thirdly,




Figure 5.3: Representative microscopy images illustrating the phenotypes of the four human cell lines which we investigated. Cells were seeded sparsely (5-10% confluence), and left to grow, or treated at MTD for 3 days (5 days in the case of TOV112d).

cells can stabilise stalled replication forks, which allows more time for repair and helps to prevent DSBs (e.g. loss of EZH2). Finally, cells can restore HR, either through restoration of BRCA1/2 activity, or via re-routing of HR pathways to work without BRCA1/2 support (e.g. down-regulation/loss of Shieldin factors, or 53BP1) [223, 222].

Given the diversity of mechanisms by which cells may achieve resistance, heterogeneity is a significant problem in developing strategies to overcome PARPi resistance. While re-establishment of HR is thought to be the most important mechanism [222], Labrie et al [230] found that different cells even within the same patient may adopt different ways to do so. Similarly, different cell lines show distinct gene expression patterns in response to prolonged PARPi exposure [230]. Currently, the main strategy to combat PARPi resistance is combination with other agents, such as CDK1 or PI3K inhibitors, to suppress re-establishment of HR [223]. However, these approaches suffer from the possibility of double resistance, and increased toxicity [223, 222].

5.2 Characterisation of *in vitro* PARPi response

Given the many steps involved in the action of PARPis, it is perhaps not surprising that the “out-of-the-box” AT algorithms and simple models in Chapter 4 were insuf-

ficient to control the spheroids - a more tailored approach is required. We therefore carried out three sets of experiments aimed at elucidating the treatment response dynamics of the cells. We performed these in 2-D *in vitro* cell culture in which cells adhere to the bottom of the flask or the culture plate. While less realistic than the 3-D spheroids studied in Chapter 4, this system is easier to handle and to image, and allowed me to carry out the experiments myself¹.

We investigated four commonly used human ovarian cancer cell lines: A2780s, TOV112d, Ovar3, and Ovar4 (Figure 5.3). The former two were those which we had investigated in the spheroid cultures in Chapter 4, which represent a BRCA-mutated and a BRCA wildtype cell line, respectively. In addition, we included Ovar3 and Ovar4 because they were derived not from primary tumour but from the ascites of patients, and as such are more representative of metastatic disease [201]. In addition, A2780s and TOV112d were collected from patients who likely had endometrioid cancer whereas Ovar3 and Ovar4 were obtained from patients with high grade serous ovarian cancer [201], which represents around 70% of ovarian cancer patients [16]. Images illustrating the distinct phenotype of each cell line are shown in Figure 5.3. We had originally planned to carry out all experiments with all four cell lines, but because of the Covid-19 pandemic this was not possible. In the following we present the experimental results as far as we were able to complete them.

5.2.1 Experimental methods

Cell culture: Cells were maintained in RPMI (ThermoFisher) supplemented with 10% Fetal Bovine Serum and 1% penicilin/streptomycine. Every 3-4 weeks the medium was additionally supplemented with MycoZap (Lonza) to prevent mycoplasm contamination. At all times cells were kept at 37°C and 5% CO₂.

Drug preparation: To prepare growth medium with the MTD drug concentration of 100μM, we diluted 24mg of Olaparib (AstraZeneca) in 1mL Dimethyl sulfoxide, filtered the solution using a 0.22nm syringe filter, and added it to 555ml of culture medium. Lower concentrations were obtained by dilution of this MTD stock with normal culture medium.

¹We acknowledge Mrs Samantha Byrne and Mr Punit Borad for helping to maintain the cell cultures.

Imaging: Cell growth was monitored once per day using an IncuCyte ZOOM time-lapse microscopy system (Essen BioScience). Phase-contrast white light images were recorded, and analysed using the IncuCyte ZOOM software.

Dose response assays: Cells were seeded in triplicates, either at low (5,000 cells), or high density (60,000 cells), in 48 well flat-bottom plates (Costar Corning), and left to attach overnight in untreated culture medium. During the 9 day experiment the medium was changed every 3 days.

Drug withdrawal experiment: Cells were seeded in triplicates at 10,000 cells per well, in 48 well flat-bottom plates (Costar Corning), and left to attach overnight in untreated culture medium. During the 22 day experiment the medium was changed every 3 days, or when it was time to withdraw drug from a well. Towards the end of the experiment most Ovar4 cells treated at $100\mu\text{M}$ were dead, but the remaining debris caused an inflation of the measurements reported by the image analysis pipeline. We removed such measurements from the analysis.

Evolution of drug resistance: To evolve drug resistant cells, we cultured cells at $25\mu\text{M}$ for 3.5 months in 275ml culture flasks. In the first two months, when cell density dropped below 20% after re-plating, we supported growth for 3-4 days by supplementing the growth medium with Mammary Epithelial Growth medium (at 1:1 ratio).

Dose response of resistant cells: Dose response was measured using the same protocol as above, but starting from 10,000 cells per well. Resistant cells were left to attach overnight in treated medium ($25\mu\text{M}$).

5.2.2 Response to continuous treatment

Firstly, we characterised the dynamics of cell growth in the absence of treatment, and under continuous drug exposure at different drug concentrations. To do so, we seeded the cells, and let them grow for 9 days, monitoring their growth using time-lapse microscopy (Figure 5.4a). In the absence of treatment, the populations expanded in a typical, sigmoidal fashion, overgrowing the plate in 7-9 days. Addition of drug slowed population growth, but a regression in size was only observed once drug concentration exceeded $50\mu\text{M}$. Furthermore, we found that there was a delay

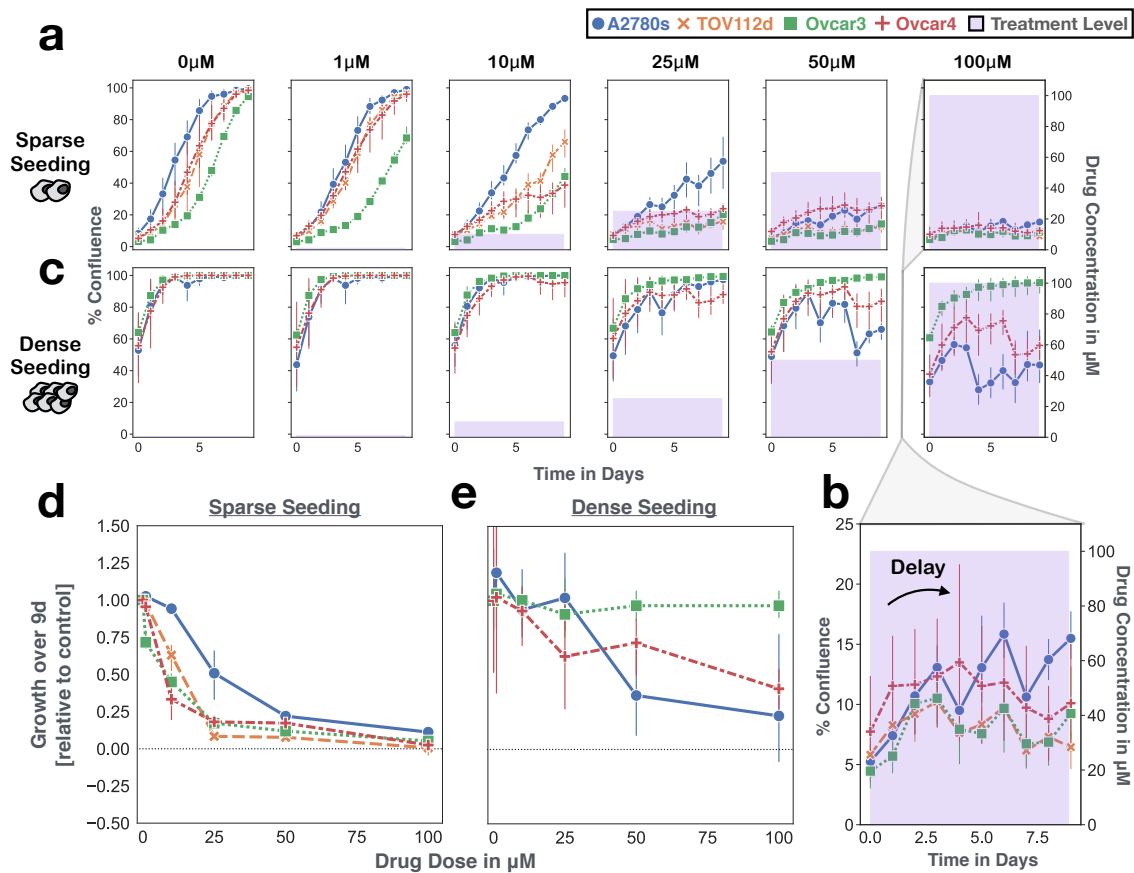


Figure 5.4: Dynamics in response to continuous drug exposure. Error bars denote 95% CIs, based on 3 replicates per condition. (a) Growth over time when cells are seeded sparsely (5-10% confluence). (b) Enlarged illustration of the dynamics at MTD, showing the delay in the drug response. (c) Drug response dynamics when cells are plated more densely (60% confluence), suggesting higher cell densities reduce the drug effect (no TOV112d data available due to experimental error). (d) & (e) Dose-response curves at low and high seeding density based on the data in (a) and (c), quantifying the protective effect of higher densities. Response measured as $\frac{n(9,D)-n(0,D)}{n(9,0)-n(0,0)}$, where $n(t, D)$ denotes the confluence at time t and dose D .

in the drug response. Even at MTD the population initially expanded for 2-3 days before it regressed (Figure 5.4b).

In addition, we hypothesised that, because DNA replication is an important potentiator of PARPi action, higher seeding density may protect cells by restricting proliferation in the population. To test this, we repeated the experiment with cells seeded at 60% confluence. The results confirm our hypothesis (Figure 5.4c). For example, at $25\mu\text{M}$ there is a large difference between the treated and untreated cultures if they were seeded sparsely, but almost no difference when they are seeded more densely (Figures 5.4a & c). Quantifying dose response by plotting the relative growth

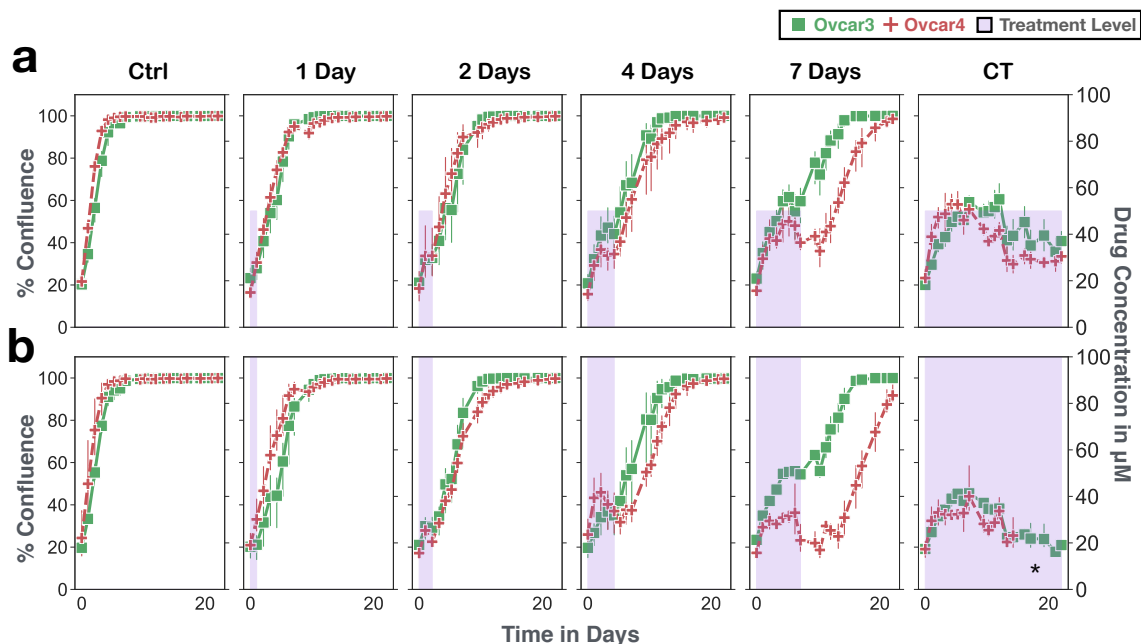


Figure 5.5: Dynamics following treatment withdrawal. Error bars denote 95% CIs, based on 3 replicates per condition. (a) Cells were exposed to treatment ($50\mu\text{M}$) for varying lengths of time, after which it was withdrawn. This shows that cells recover very quickly from drug exposure. (b) Same as (a) but at MTD ($100\mu\text{M}$). * : the final measurements for Ovc4 under CT were unreliable, due to debris from dead cells, and were excluded from the analysis.

under drug compared to the untreated control illustrates how both drug sensitivity (concentration at which there is a 50% growth reduction) and efficacy (maximum observed reduction) are diminished at higher seeding density (Figures 5.4d & e). In addition, this effect is more pronounced in Ovc3 and Ovc4 cells, which grow as defined colonies, than in A2780s cells, which grow more diffusely (see also Figure 5.3).

5.2.3 Response to treatment withdrawal

A key aspect of AT is that treatment is withdrawn, or modulated, over time. Next, we therefore asked how the cells would respond to such treatment withdrawal. We exposed cells to drug for a varying duration of time, after which we removed treatment, and monitored how quickly the population recovered. We repeated this for two of the cell lines (Ovc3 and Ovc4), and two drug concentrations ($50\mu\text{M}$ and $100\mu\text{M}$).

Our results show that the cells recover very rapidly from drug exposure, especially if this exposure has been shorter than 7 days (Figures 5.5a & b). Even if the population has started to regress during treatment, it will expand again within 24 hours after drug withdrawal. Only after 7 days of treatment does there appear to be a delay

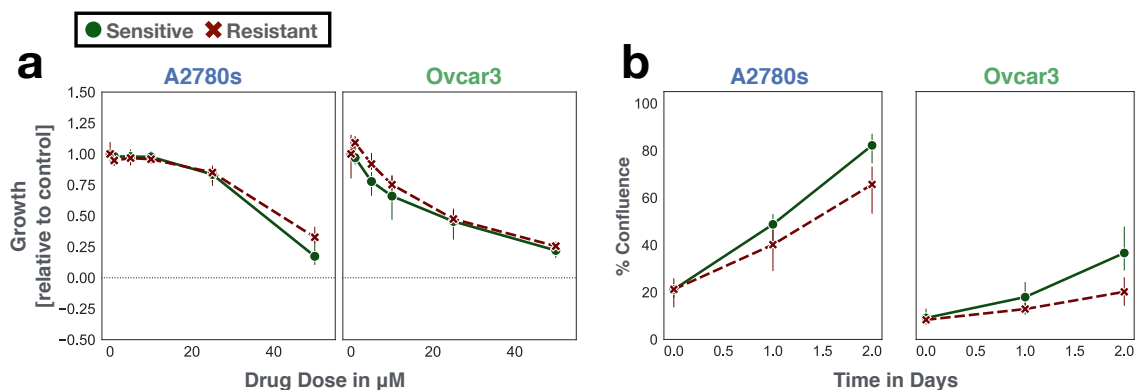


Figure 5.6: Cells show no evidence of resistance even after having been continuously cultured in $25\mu\text{M}$ of drug for 3.5 months. “Resistant” refers to cells being cultured under drug, and “Sensitive” refers to their parental, untreated counterparts, respectively. Error bars denote 95% CIs, based on 3 replicates per condition. (a) Dose-response curves, showing almost no difference in the sensitivity to treatment (calculated as in Figure 5.4e, based on growth over 10 days (A2780s) and 6 days (Ovarcar3), respectively). (b) Growth in the absence of treatment over 48 hours.

before the population is able to resume its expansion. This delay is greater in the more drug sensitive, and faster dividing, Ovarcar4 cell line than in the less sensitive, and more slowly dividing, Ovarcar3 cells (Figures 5.5a & b).

5.2.4 Characterisation of drug resistant cells

The final set of experiments was aimed at characterising Olaparib resistant cells. We sought to artificially evolve drug resistance by culturing cells at $25\mu\text{M}$ for 3.5 months. However, we found that even after this prolonged exposure, there was little change in the drug sensitivity of the cells (Figure 5.6a). That being said, we also observed that in the absence of treatment the growth rate of the “resistant” populations was reduced compared to that of their “sensitive” counterparts (Figure 5.6b). Together this suggests that over the time frame of our experiment, CT had primarily selected for more slowly growing cells and molecular resistance had not yet emerged.

5.2.5 Discussion

The aim of our experiments was to gain an understanding of why adaptive strategies had performed so poorly in the spheroid experiments in Chapter 4. To do so, we simplified the experimental system to a 2-D, adherent cell culture model which is more robust to the issues with growth medium changes and measurement which had confounded our analysis in the spheroids. Overall, we make four key observations.

Firstly, we find that intense, continuous treatment at drug levels in excess of $50\ \mu\text{M}$ is required to induce tumour size regression in our cell lines. This observation is consistent with prior drug sensitivity measurements in the literature. For example, the IC_{50} values (the dose at which the population grows to 50% the size of the untreated control) for Olaparib, reported for 25 different ovarian cancer cell lines in the database by Yang et al [231], range from $13.9\ \mu\text{M}$ to $1,010\ \mu\text{M}$ with a mean of $150\ \mu\text{M}$. This may explain why we observed relatively small differences between the different treatment arms in the initial spheroid experiments in Section 4.2, as there we used a maximum dose of $10\ \mu\text{M}$.

Moreover, at this point we should recall that one of the key assumptions, which asserts the benefit of AT according to the argument from Chapter 2 (Section 2.3.1.4), is that we have complete control over the sensitive population. If this is not the case, then the benefit of AT depends on a trade-off between, the gain in TTP from increased competitive suppression of resistant cells and the loss from reduced suppression of sensitive cell growth. For example, Monro et al [59] found in their simulation study of a 2-population ODE model, akin to that studied in Chapter 2, that de-escalation shortens TTP if treatment can not decrease the size of the sensitive population. As such, our data indicate that the scope for dose reductions or long treatment holidays may be limited, because it is difficult to reverse the sensitive population growth that these will cause.

A further factor which may explain why AT approaches were inferior in our experiments is a 2-3 day delay in drug response. Importantly, during this time the population continues to expand, so that even after 7 days of treatment at, or near, MTD, the net growth is still positive. In contrast, when treatment is withdrawn, regression stops almost immediately. We draw two conclusions from this. Firstly, drug has to be administered for a prolonged period of time before a size reduction can be achieved (at least 7 days, but likely longer), even if given at high doses. Secondly, drug vacations (if any) will have to be shorter than drug on-times. This implies that the time-scales over which treatment decisions are made have to be carefully adapted to the treatment dynamics. In the spheroid experiments in Chapter 4 we were re-evaluating treatment decisions every 3-4 days, which was likely too short an interval for on-times, and too long for off-times.

A characteristic of AT approaches is that they maintain a higher tumour burden than an aggressive CT regime. Our experiments suggest that this may represent an additional challenge for AT, as more densely seeded cell populations appear to experience significant protection from treatment in this setting. This may also explain

why the spheroid cultures were more resilient to treatment at $100\mu\text{M}$ than the 2-D cell cultures, an observation also made elsewhere [232].

Finally, we note that while AT aims to avoid the emergence of molecular drug resistance, we failed to observe evidence for such resistance over the time course of our experiments. That being said, molecular resistance to Olaparib does arise in patients (see Section 5.1.1.3), and has been derived by other groups using similar protocols to ours (e.g. [233, 234, 235]). An explanation for this discrepancy may be the time of exposure. For example, Kim et al [235] cultured their cells for 1.5 years to evolve Olaparib resistance. As such, there may still be a benefit to carefully administered AT to suppress resistance in the long term. But in order to investigate this, a different experimental model to our spheroid or 2-D system will be required, as it is not possible to culture cells for such a long time without changing the culture plate and thereby confounding measurements. One solution would be to start the experiment with a mixture of fluorescently-tagged sensitive and evolved resistant cells to directly evaluate how well resistant cells can be suppressed by competition, similar to the work of Bacevic et al [115]. Furthermore, any findings should be validated in appropriate *in vivo* models, as *in vitro* conditions provide a very limited representation of the conditions in real tumours. Evidence for this from our experiments is the fact that the estimate for the Olaparib dose achieved at MTD in tumours in the clinic is only around $20\mu\text{M}$ [236, 237] yet, unlike what is predicted by our data, this induces tumour size regression in patients.

To sum up, our experiments indicate that while there may be scope for AT in PARPi treatment, careful thought is required as to how dosing is adjusted. In fact, our data show very clearly that if one de-escalates too much or too early, then there is a risk of accelerating recurrence, and this is likely what happened in Chapter 4. That being said, our work also shows that a different experimental setup will be required to assess whether AT can delay progression. Given such experiments will be both time consuming and expensive, we will dedicate the second part of this chapter to developing a mathematical model of PARPi response. Using this model will allow *in silico* exploration of whether or not there is likely to be a benefit of PARPi-driven AT, which will be helpful in planning further experiments.

5.3 A mathematical model of PARPi treatment

To the best of our knowledge, there exists only a single study which has so far modelled PARPi treatment. In this study, the authors used a multi-type branching process

model to explore the evolution of BRCA mutations before and during treatment of ovarian cancer with PARPis and chemotherapy [238]. They assumed that the tumour can be divided into six different, exponentially growing, cell populations, differing in their BRCA mutation status (which determines drug sensitivity), their growth rate (fast or slow), and whether they arise prior to, or during, treatment. Treatment was assumed to reduce the proliferation rate of sensitive cells, and was applied continuously until progression. Yamamoto et al [238] showed that their model could qualitatively recapitulate a number of clinical observations (e.g. tumour heterogeneity at the start of treatment), but they did not directly compare it against any time-series data. In light of this limited amount of prior work, we decided to use the data we had collected to develop and calibrate a mathematical model of PARPi response.

5.3.1 The previously used ODE model is insufficient to describe the experimental data

As a first step, we assessed whether the ODE model we had studied in Chapters 2 & 4 (Equations (2.1)-(2.3)) could describe the *in vitro* dynamics. To do so, we fitted the model to the growth data collected in the absence of treatment and under MTD from Figure 5.4a, and we tested how well the model could predict the response to treatment at a lower dose, or to drug withdrawal (data from Figures 5.4c & 5.5, respectively). Given the short time-scale of our cell culture experiments, and the results of our artificial evolution experiment, we assumed that the contribution of any emerging resistant population was negligible. As such, we simplified the model to the equation for the sensitive population:

$$\frac{dN}{dt} = r \left(1 - \frac{N}{K}\right) (1 - d_D D(t)) N - dN, \quad (5.1)$$

where $N(t)$ denotes the confluence of the cells (in %), r is the cell proliferation rate (in d^{-1}), K is the carrying capacity (in % of the well's area covered by cells), d is the death rate (in d^{-1}), and d_D represents the drug-induced kill fraction (in μM^{-1}). Similar to Chapter 4, we assumed that the drug concentration in the medium, $D(t)$ (in μM), was exactly equal to that administered experimentally, and remained constant between medium changes.

5.3.1.1 Numerical methods

Model simulation and fitting: We fitted the model in a 2-step procedure. Firstly, we estimated the parameters describing growth in the absence of treatment (r , K).

Subsequently, we set r and K to their inferred values and calibrated the treatment term by optimising the value of d_D to fit the data from treatment at $100\mu\text{M}$. We tried also to infer d , but found that because of its strong correlation with the other parameters it was difficult to obtain reliable estimates. As including d did not change the conclusions of our analysis, we decided to remove it from the model for simplicity. Fitting was carried out by minimising the RMSE between the solution to Equation 5.1, and the mean value of the observed confluence across the three experimental replicates for each condition. To solve Equation 5.1 we set the initial conditions equal to the mean initial confluence in the plate, and integrated forward in time using the explicit Runge-Kutta method of order 8 (DOP853) in scipy [179]. Minimisation was carried out using the lmfit package [162] and the scipy implementation of the Levenberg-Marquardt algorithm (`least_squares()`).

Predictions: In order to validate the model we assessed its accuracy in predicting response to experimental conditions different to those on which it had been trained. To do so, we set the initial conditions exactly equal to those in the experiment, and subsequently simulated the model forward using the identical protocol (i.e. drug concentration/changes) to that used in the experiment. In order to capture variation which may arise due to differences in the initial seeding density, we repeated this once for each of the three experimental replicates. When displaying our results (e.g. Figure 5.7b) we show the mean and 95% CIs obtained by bootstrapping these three predictions (2500 bootstraps).

Data: As we had most data available for the Ovar3 and Ovar4 cell lines, and these are thought to be more representative of serous ovarian cancer [201], we focussed on these two cell lines for our analysis. However, the final model we developed also works well for A2780s and TOV112d (see Tables 5.1 & 5.2 for the fitted parameters).

5.3.1.2 Results

In Figure 5.7a we show the results of fitting this simple model to our cell line data. We find that even though logistic growth is generally thought to be a poor descriptor of tumour growth, a result corroborated by our work in Chapter 3, it provides a good representation of the cell growth dynamics in this case (Figure 5.7a; left column). Similarly, its predictions of the % confluence after 9 days of treatment is fairly accurate (Figure 5.7a; right column). That being said, its description of the time-series

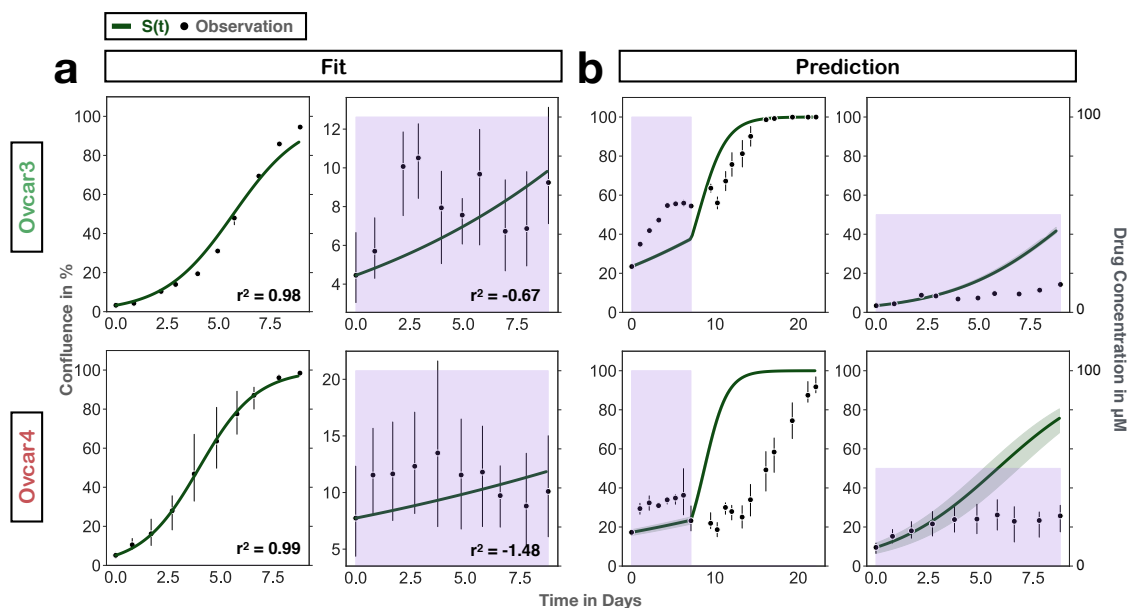


Figure 5.7: The model of sensitive cell growth and treatment response, studied in previous chapters of this thesis (Equation 5.1), is insufficient to describe the experimentally observed PARPi response. Error bars denote 95% CIs (based on 3 experimental replicates). (a) Fitting of the model to the data of untreated and treated growth ($100\mu\text{M}$) from Figure 5.4a. (b) Predictions of the calibrated model for the response to treatment withdrawal after 7 days of exposure, or treatment at a lower dose ($50\mu\text{M}$), in comparison to experimentally observed dynamics. Green lines show means, and green bands denote 95% CIs based on variation in initial seeding numbers (see Section 5.3.1.1 for details).

dynamics under treatment is poor, as it fails to capture the non-monotonic, delayed response shown by the cells.

In order to investigate how reliably one may use the model to draw conclusions about AT, we next tested how accurately it predicted response to treatment withdrawal or modulation. We find that in both cases the accuracy of the predictions is poor (Figure 5.7b). In predicting response to treatment withdrawal, the model under-estimates growth under treatment, and over-estimates growth once treatment is withdrawn. Although it still correctly predicts the % confluence at the end of the experiment, it is questionable whether this would still hold true after several rounds of intermittent treatment. Similarly, when used to predict growth under a lower dose of treatment, the model greatly under-estimates the drug-induced kill (Figure 5.7b).

Overall, we conclude that the simple model we had studied in previous chapters of this thesis is insufficient to accurately describe the experimentally observed drug response. As such, we will dedicate the next section to developing a more accurate

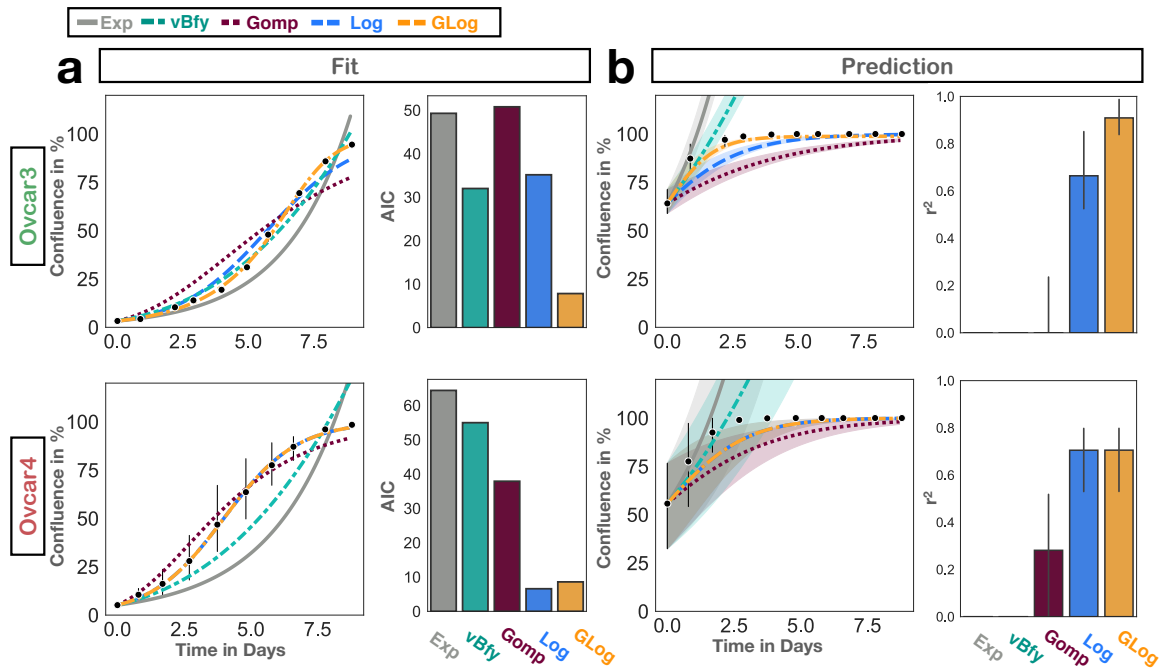


Figure 5.8: Evaluation of 5 different plausible growth models identifies the GLog model as the most descriptive and predictive growth model (see text for abbreviations). Error bars denote 95% CIs. (a) Comparison of the descriptive ability of the 5 models in fitting untreated growth data from cells seeded at low density. (b) Predictions of the calibrated models for the growth curve when cells are seeded at near 60% confluence, in comparison to the experimentally observed dynamics. Lines show means, and bands denote 95% CIs based on variation in initial seeding numbers (see Section 5.3.1.1 for details).

model of PARPi treatment.

5.3.2 Model development

5.3.2.1 The growth model

To begin with, we carefully selected a model for the growth of the cells in the absence of treatment. PARPi action depends on cell division, and one of the key experimental observations was that the cells' treatment response depended on their seeding density. While the logistic model had performed well in fitting growth starting from a low seeding density (Figure 5.7a), we wanted to ensure that we chose a growth model which could accurately represent the growth dynamics of the population at different densities. To do so, we evaluated five plausible, and commonly used, growth models: i) Exponential growth (Exp; $\frac{dN}{dt} = rN$), ii) Von Bertalanffy growth (vBfy; $\frac{dN}{dt} = rN^{\frac{2}{3}}$), iii) Gompertzian growth (Gomp; $\frac{dN}{dt} = r \log\left(\frac{K}{N}\right) N$), iv) Logistic growth (Log;

Table 5.1: Estimated growth parameters assuming a GLog growth model. ΔAIC is the difference in AIC to the best fitting model for that cell line. Values smaller than -2 would indicate that a different model provides a more parsimonious fit to the data, and should be preferred to the GLog model.

Cell Line	ΔAIC	r (in d^{-1})	K (in %)	ν (no dims)
A2780s	-1.53	0.88	100	0.91
TOV112d	0	0.51	99.8	2.16
Ovcar3	0	0.47	98.8	2.6
Ovcar4	-2.0	0.73	100	1.0

$\frac{dN}{dt} = r \left(1 - \frac{N}{K}\right) N$, and v) Generalised logistic growth (GLog; $\frac{dN}{dt} = r \left(1 - \left(\frac{N}{K}\right)^\nu\right) N$). As before, $N(t)$ is the % confluence of the cells in the well, r is the cell growth rate, K is the carrying capacity, and ν is an exponent determining the curvature of the relationship between density and growth rate. Model fitting and prediction were carried out using the same methods as in Section 5.3.1.1.

In Figure 5.8a we show the results of fitting each model to the untreated growth data. Interestingly, we find that the best fitting model depends on the cell line, which might be reflective of the different spatial morphology of the cells (see also Figure 5.3). That being said, the GLog model performs well across all cell lines, even when a penalty for its extra parameter is applied (AIC score; Figure 5.8a, and Table 5.1). A summary of the inferred parameters and the goodness of fit of the GLog model for all four cell lines is shown in Table 5.1.

Next, we tested how well the models could predict the growth dynamics when cells are seeded at a higher density. As testing data we used the untreated control from the drug sensitivity experiment in which we had seeded cells at 60% initial confluence (Figure 5.4c). We find that the GLog model provides the most accurate predictions, although for both cell lines it slightly under-predicts the observed cell growth rate (Figure 5.8b). Based on the convincing results in both fitting and predicting growth data we therefore decided to use the GLog model.

5.3.2.2 The drug response model

Next, we developed a mechanistic, mathematical model of the cells' drug response. Based on our review of the biological basis of PARPi action (Section 5.1.1.1), we made the following assumptions about the mechanisms underlying treatment response:

- PARPis do not kill cells directly, but induce cell cycle arrest.

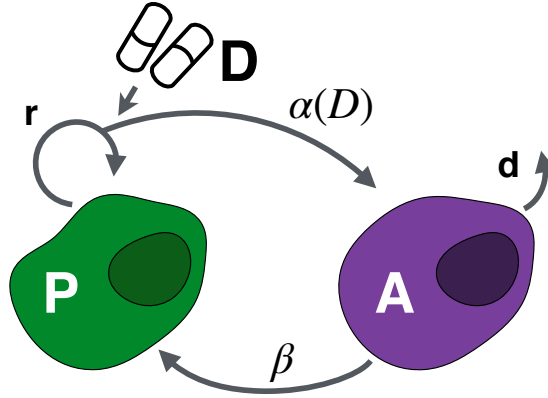


Figure 5.9: The mathematical model of PARPi response (Equations (5.2) - (5.4)). PARPis are assumed to induce cell cycle arrest in actively proliferating cells (P). Arrested cells (A) then attempt to repair the damage, or undergo apoptosis.

- Arrest is induced primarily in those cells which are actively undergoing division, due to the stalling of replication forks as a result of PARP trapping.
- Arrested cells will attempt to repair the damage caused by PARPis. If they succeed, they can recover and continue proliferation.
- Only if a cell cannot complete repair, will apoptosis be induced.

Based on these assumptions, we propose a simple, 2-compartment ODE model (Figure 5.9). We divide the population into actively proliferating cells, and cells in cell cycle arrest, whose % confluence we will represent as continuous functions of time denoted by $P(t)$ and $A(t)$, respectively. Furthermore, we assume that the temporal dynamics of the two populations is governed by the following ODEs:

$$\frac{dP}{dt} = r \left(1 - \left(\frac{P + A}{K} \right)^\nu \right) (1 - \alpha(D))P + \beta A, \quad (5.2)$$

$$\frac{dA}{dt} = \alpha(D) r \left(1 - \left(\frac{P + A}{K} \right)^\nu \right) P - \beta A - dA, \quad (5.3)$$

$$N(t) = P(t) + A(t), \quad (5.4)$$

where $N(t)$ is the total confluence of the population (in %), as observed in the experiment. Here we assume that proliferating cells follow the GLog growth model we identified in the previous section (with the inferred parameter values given in Table 5.1), whereas arrested cells either recover from the drug-induced damage at rate β (in d^{-1}), or die at rate d (in d^{-1}). To represent the fact that PARPis preferentially induce cell cycle arrest in those proliferating cells which are currently undergoing mitosis, we

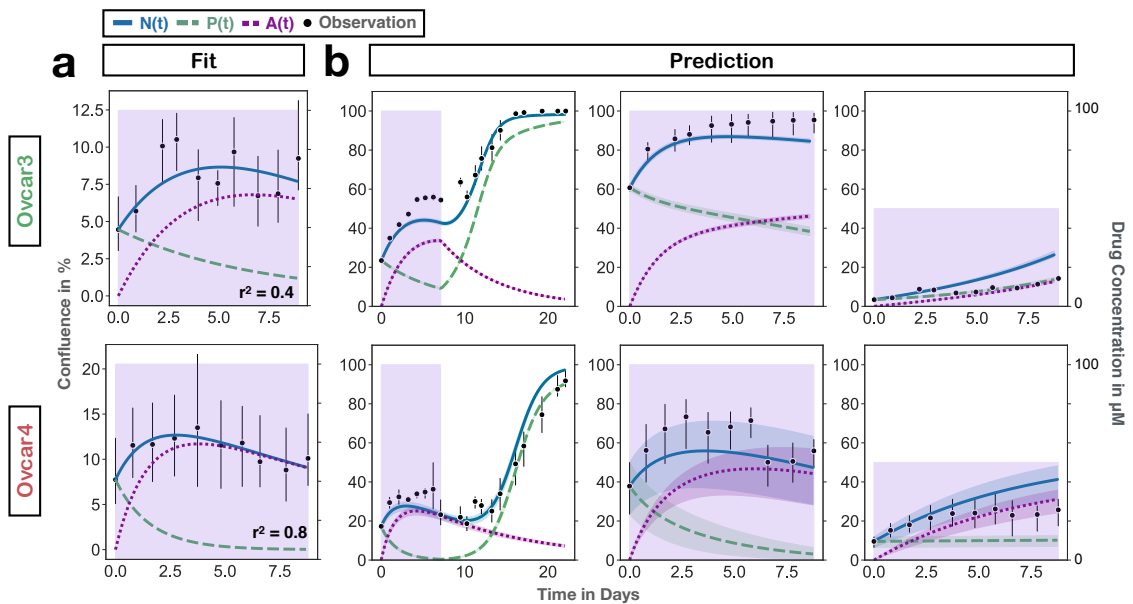


Figure 5.10: Fits and predictions of the proposed 2-population ODE model for PARPi response (Equations (5.2) - (5.4)). Error bars denote 95% CIs. **(a)** Fitting of the model to the experimentally observed dynamics under continuous treatment at $100\mu\text{M}$ (compare also to Figure 5.7a). **(b)** Predictions of the calibrated model for three unseen experimental protocols, in comparison to the experimentally observed dynamics (compare also to Figure 5.7b). Left: treatment withdrawal after 7 days; Centre: continuous treatment at $100\mu\text{M}$, but with cells seeded at higher initial density; Right: continuous treatment at a lower dose ($50\mu\text{M}$). Lines show means, and bands denote 95% CIs based on variation in initial seeding numbers (see Section 5.3.1.1 for details).

assume that the rate of arrest induction is proportional to the population growth rate of $P(t)$, where the proportionality constant is the drug-effect, $\alpha(D)$ (dimension-less). In this section we will, for simplicity, assume that the drug-effect is a linear function of dose ($\alpha(D) = \sigma D$, where σ has units μM^{-1}), although we will revisit this assumption in the next section. Finally, we assume that, because cells were grown in treatment-free medium before the drug was added, the initial population is entirely composed of proliferating cells, so that $P(0)$ is exactly equal to the confluence observed at $t = 0\text{d}$, and $A(0) = 0\%$.

In Figure 5.10a we show the results of fitting the model to the experimentally observed drug response at $100\mu\text{M}$. Fitting was done using the same methods as in Section 5.3.1.1, leaving σ , β , and d as free (cell line-specific) parameters. We find that the model can capture the observed non-monotonic behaviour well, in particular in the Ovc4 cells (compare also to Figure 5.7a). The fits suggest that treatment induces cell cycle arrest in the majority of cells by the end of the experiment. However,

because this process occurs only gradually, there is still a transient initial burst of proliferation, which causes the observed peak.

Next, we tested the ability of the model to predict the dynamics under treatment protocols which are different to those on which it was trained (Figure 5.10b). This shows that the model can predict the dynamics in response to treatment withdrawal with high accuracy. Similarly, the model is also able to capture the protection from higher seeding densities, which we had observed in the experiments. However, for both cell lines we see that it tends to slightly over-predict cell kill during treatment. Conversely, when we examine its predictions for treatment response at a lower drug dose, it significantly under-predicts drug kill (Figure 5.10b). We conclude that while the model provides an accurate representation of the time dynamics in response to treatment, it does not capture correctly the dose-effect relationship.

5.3.2.3 The dose-effect model

For simplicity, we assumed in the previous section that the relationship between the drug concentration, $D(t)$, and its instantaneous effect on the cells, $\alpha(D)$, is linear ($\alpha(D) = \sigma D$). However, as there are many molecular processes involved in the induction of cell cycle arrest by PARPis, this is likely an over-simplification. In order to investigate what the dose-effect relationship may be in our data, we devised the following method to measure it. We fixed β , and d at the values which we had estimated in the previous section, and allowed $\alpha(D)$ to be a constant which was independently defined for each dose ($\alpha(D) = \sigma_D$ for $D \in [0, 1, 10, 25, 50, 100] \mu\text{M}$). Next, we determined σ_D by fitting the model (Equations (5.2) - (5.4), assuming $\alpha(D) = \sigma_D$) to the data from treatment at dose, D , with σ_D as the only parameter allowed to vary (other numerical methods as in Section 5.3.1.1; for the fits, see Figure 5.11a). By doing so, we obtained what we will refer to as the *empirical dose-effect relationship* (Figure 5.11b). This is for each dose, D , the value of $\alpha(D)$ which would give the best model fit to the data, and this is the relationship which analytic expressions for $\alpha(D)$, such as $\alpha(D) = \sigma D$, are trying to approximate. For brevity, we will in the following focus on the results for Ovc3. Analysis of the Ovc4 data yields the same conclusions.

In Figure 5.11b we show the thus measured empirical relationship for the Ovc3 data, as well as the linear model from the previous section ($\alpha(D) = \sigma D$). We see that the empirical relationship is clearly concave, which explains why the linear model was under-estimating cell kill at lower doses in Figure 5.10b. To address this issue we investigated whether a non-linear E_{Max} model (also referred to as the Hill equation),

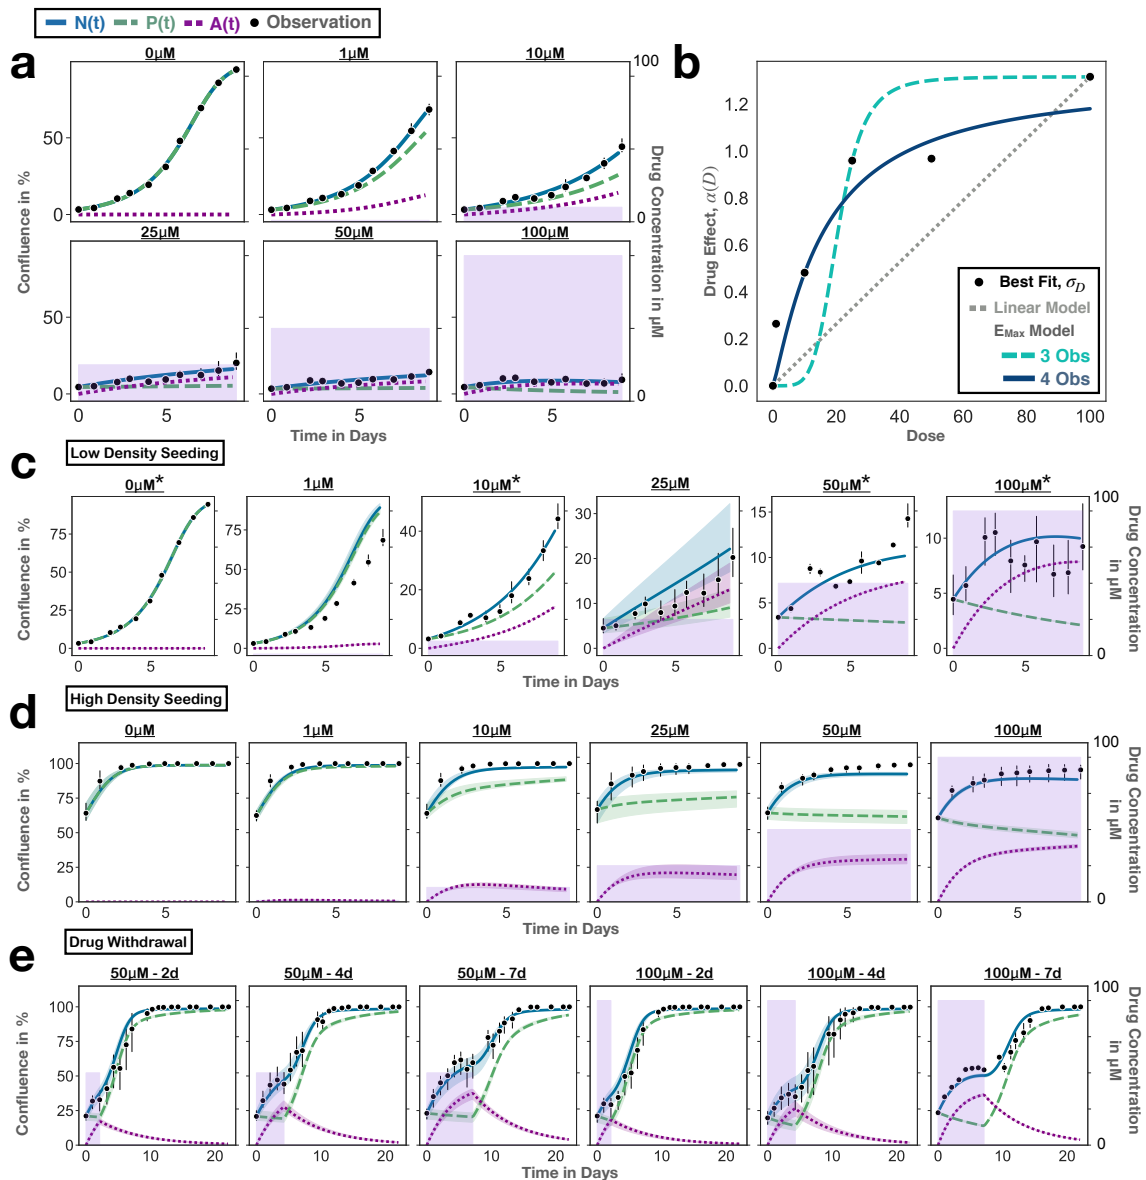


Figure 5.11: Derivation of a more accurate dose-effect model (Ovcar3 data). Error bars denote 95% CIs. (a) Model fits, with the drug-effect $\alpha(D)$ as a dose-specific (and the only free) parameter. This yields one value $\alpha(D) = \sigma_D$ for each concentration D - the empirical dose-effect curve (see text for details). (b) Plot of the empirical dose-effect curve (points), together with the linear model from Section 5.3.2.2 (dotted line), and the E_{Max} model (Equation (5.5)). Shown are two parametrisations of the E_{Max} model: One based on fitting only to the data at 0, 50, and 100 μM (dashed line), and one based on fitting to the data at 0, 10, 50, and 100 μM (solid line), which was the one taken forward for analysis. (c) Fits and predictions of the E_{Max} model for the treatment response when starting from low seeding densities. Starred data were used for calibrating the model. (d) Predictions of the E_{Max} model for the treatment dynamics at a higher seeding density. (e) Predictions of the E_{Max} model for the treatment dynamics upon treatment withdrawal. In the predictions in panels (c)-(e), lines show means, and bands denote 95% CIs based on the variation in initial seeding numbers (see Section 5.3.1.1 for details).

Table 5.2: The fitted parameters for the E_{Max} model (Equations (5.2) - (5.5)). For the model parameters associated with untreated growth parameters, see Table 5.1.

Cell Line	E_{Max} (no dims)	k_{50} (in μM)	n (no dims)	β (in d^{-1})	d (in d^{-1})
A2780s	0.95	14.3	1.8	3×10^{-6}	0.68
TOV112d	1.22	12.0	1.28	1×10^{-6}	0.42
Ovcar3	1.32	16.0	1.17	2×10^{-8}	0.15
Ovcar4	1.97	11.7	0.59	4×10^{-13}	0.08

which is a commonly used model of dose-response in the literature [79], could provide a more accurate approximation. This is given by:

$$\alpha(D) = \frac{E_{\text{Max}}D^n}{k_{50}^n + D^n}, \quad (5.5)$$

where E_{Max} is the maximum drug effect (dimensionless), k_{50} is the drug concentration at which 50% of the drug effect is achieved (in μM), and n is a parameter governing the slope of the curve (dimensionless). As we will use Equation (5.5) only in conjunction with the equations for P and A , we will in the following adopt the name *E_{Max} model* also to refer to the full model (Equations (5.2) - (5.5)). Due to the two extra parameters in this model in comparison to the prior linear model, we employed a 2-step process to fit it (numerical methods as before). Firstly, we calibrated E_{Max} , β , and d by fitting the model to the experimental data from continuous treatment at MTD, allowing these three parameters to vary, and assuming that $\alpha(D) = E_{\text{Max}}$ (for the fit see Figure 5.11c; $100\mu\text{M}$). In the next step, we obtained n and k_{50} by fixing all other parameters, and fitting the model to the data from several drug concentrations simultaneously by minimising the combined RMSE. We experimented with different choices of which, and how many different, concentrations to use in this step, and found that at least two interior points (concentrations between 0 and $100\mu\text{M}$) are required to obtain a good fit. With just a single interior point, the system is under-determined and yields a poor approximation of the dose-effect relationship (Figure 5.11b; dashed line). Furthermore, the particular combination of interior points affects the inferred curve (not shown).

In Figure 5.11b we show the curve inferred using the data from treatment at 0, 10, 50, and $100\mu\text{M}$ (solid line), which provides the most even spread of error which we could obtain, and which we carried forward in our analysis. The associated fits are depicted in Figure 5.11c and show that the E_{Max} model provides an accurate description of the dynamics under all four drug concentrations (denoted by *; for inferred parameters, see Table 5.2). Moreover, the model provides highly accurate




Figure 5.12: Quantitative comparison of the goodness-of-fit of the three different models (Ovcar3 data). For each experimental replicate the model prediction was computed, and compared to the experimental data using the r^2 value. Shown are the means (bars) and 95% CIs (error bars) based on these three replicates per condition. For visualisation purposes, negative r^2 values were set to 0 for the analysis.

predictions of the response of cells treated at other drug concentrations, cells grown at higher densities, cells subjected to treatment withdrawal, and combinations thereof (Figures 5.11c-e). To quantify this, in Figure 5.12 we show the goodness-of-fit of the E_{Max} model, which corroborates that the model predictions match the data closely in most cases (typically achieving an r^2 value of greater than 0.6). Moreover, when we compare this to the linear drug-effect model from Section 5.3.2.2, and the single Equation model from Section 5.3.1, we find that the E_{Max} model is typically the most accurate out of the three models, and in cases where it is not, the difference to the best model is small (Figure 5.12). The same holds true for the other three cell lines, although for A2780s and TOV112d we did not have data to test the predictions of the response to treatment withdrawal (not shown). Overall, we conclude that the E_{Max} model provides a simple, yet descriptive and predictive, model of the PARPi response observed in our experiments. In the final section of this chapter, we will use this model to investigate different drug scheduling protocols with the aim of identifying alternatives which reduce toxicity, and may help to delay the emergence of PARPi resistance.

5.3.3 Application to adaptive PARPi treatment

5.3.3.1 Exploration of strategies for treatment de-escalation

The motivation for constructing this model was to explore whether AT may allow better control of PARPi resistance on time scales which are longer than those we examined in our experiments. While an exhaustive investigation of this issue was

outside the scope of this thesis, we will in the following present preliminary results to motivate and guide further research. One of the key assumptions of AT is that one can control the sensitive population at will, so that one can leverage it to competitively suppress resistance. However, our experimental work in the first part of this chapter indicated that in the case of PARPis this assumption is not necessarily true. Instead, our results showed that great care has to be taken in the timing and degree of treatment de-escalation. To begin with, we therefore explored whether it is possible to maintain control over the sensitive cells whilst also de-escalating treatment.

Experimental evidence that this may be possible is provided by the spheroid experiments we carried out in Chapter 4 (Figure 4.18). While the two adaptive strategies performed more poorly than CT, the difference in tumour size was relatively small, given the adaptive protocols used only about half as much drug as CT. At the same time, the modulation-based strategy (AT1) appeared to perform better than the vacation-based algorithm (AT2), indicating that how treatment is de-escalated matters. To investigate this further, we used the calibrated model to simulate CT, AT1, and AT2 (Figure 5.13a). From this, we make two important observations. Firstly, we find that even though the model was calibrated with 2-D cell culture data, its predictions qualitatively match the treatment responses in the spheroids (compare plots to insets in Figure 5.13a). For example, under CT at $100\mu\text{M}$, the tumour briefly expands before it retracts and remains at a small size, whereas under AT1 its size cycles in response to treatment de- and re-escalation. Although there are differences, likely due to the different growth dynamics of the spheroids, this further corroborates the model's ability to predict response to different treatment schedules.

Secondly, the model predicts that while it is possible to reduce the dose and still achieve tumour regression, withdrawing treatment completely results in rapid re-growth which is challenging to undo (Figure 5.13a). Again this behaviour is qualitatively consistent with that seen in the spheroids under AT1 and AT2 treatment (insets in Figure 5.13a), and formalises the verbal model we developed in the experimental section of this chapter. Moreover, when we simulate CT using only half the MTD ($50\mu\text{M}$), we find that we can still achieve a comparable outcome to that of treatment under MTD, although the dynamics are slightly slowed down (Figure 5.13b). Together this suggests that it is possible to reduce treatment whilst maintaining control over the sensitive population, but that one should avoid treatment vacations.

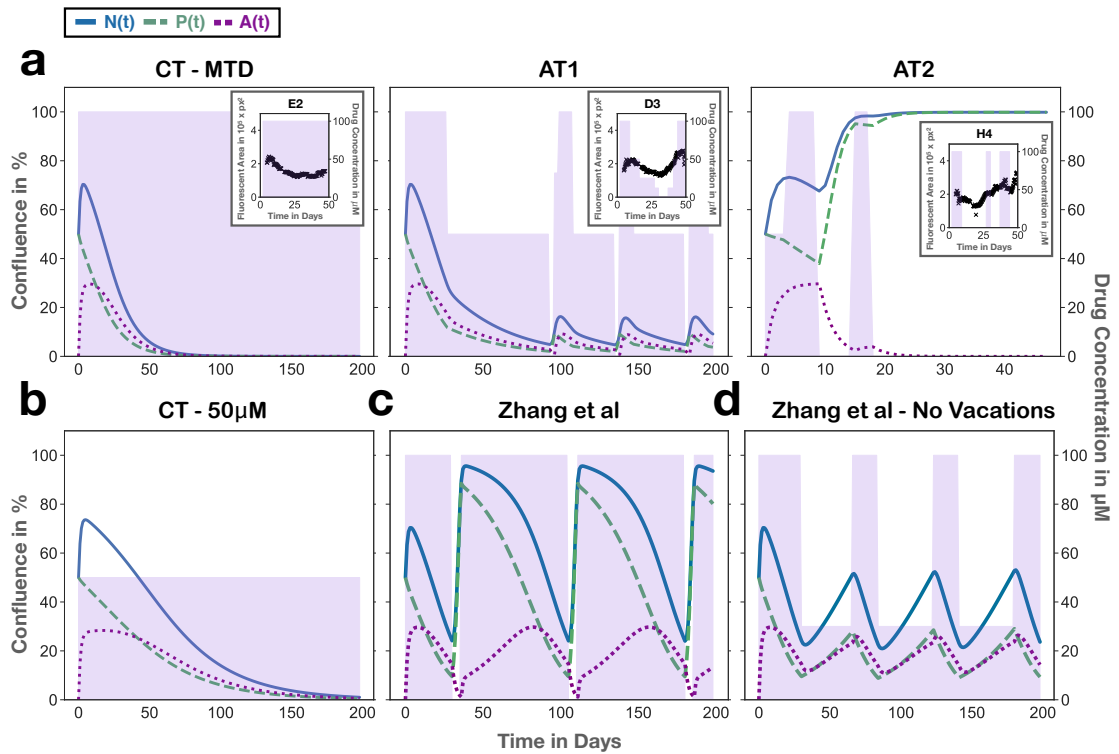


Figure 5.13: Preliminary results on how to de-escalate treatment whilst maintaining control over the drug-sensitive cells (model parametrised based on TOV112d cells). (a) Simulations of the different treatment schedules investigated in the spheroid experiments in Chapter 4. This shows that dose reductions may be possible, but that complete dose withdrawal should be avoided, which may also be why the adaptive strategies performed poorly in the experiments. Insets show example data from the TOV112d spheroids (see also Figure C.4). (b) Simulation showing that CT at half MTD can still drive tumour eradication. (c) Application of the 50% rule investigated in Chapters 2 & 3 (Equation (2.4)), illustrating that it is possible, but challenging, to maintain the higher tumour burden required for competitive suppression of resistance. (d) The proposed PARPi AT algorithm (Equation (5.6); $\tau_{AT} = 50\%$, $D_{Min} = 30\mu M$), demonstrating robust control of the sensitive cell population whilst administering a reduced cumulative dose. Parameters: $P(0) = 50\%$, $A(0) = 0\%$; others as in Tables 5.1 & 5.2; treatment updates every 3 days.

5.3.3.2 The challenges of maintaining a tolerable tumour burden

While the simulations in the previous section assert an important requirement for AT, we note that in all cases where control is achieved the final tumour size is small (Figures 5.13a & b). However, if the density of sensitive cells is so low, it is unlikely that they will be able to exert meaningful competitive suppression on the resistant cells. As such, in the next step, we investigated whether we could maintain the sensitive cell population also at a higher level. To do so, we applied the AT algorithm from

the prostate cancer clinical trial by Zhang et al [97] (Equation (2.4) and Figure 1.3c), which we had investigated in the first part of this thesis. We find that this protocol induces cycling of the tumour population around its initial size, which indicates that maintaining a higher burden is possible in principle (Figure 5.13c). That being said, the amplitude of these oscillations is large, so that the population regularly exceeds its original size significantly. This overshoot is again driven by the withdrawal of treatment, and makes this protocol unlikely to be suitable for practical use.

5.3.3.3 A new AT algorithm for PARPi treatment

Based on our insights from the previous sections, we hypothesised that it may be possible to address this issue by modifying the Zhang et al [97] algorithm. In particular, we propose that once the desired tumour size reduction, τ_{AT} , is achieved, treatment is reduced, but not completely withdrawn:

$$D(t) = \begin{cases} D_{\text{Max}}, & \text{until } N(t) < \tau_{AT} N_0 \\ D_{\text{Min}}, & \text{until } N(t) = N_0, \end{cases} \quad (5.6)$$

where $N_0 = N(0)$, and D_{Min} is chosen so that the tumour slowly regrows ($dN/dt > 0$). In Figure 5.13d we show an example of this modified algorithm, assuming the same tumour parameters as before. Under this new algorithm, the tumour cycles in a fashion that is similar to the dynamics we saw in Chapters 2 & 3, with only short intervals of high intensity treatment, followed by prolonged periods of slow re-growth under lower drug doses. As such, we propose that this modified algorithm may provide a more effective AT protocol for PARPis than those we have trialled so far.

To provide a proof-of-principle illustration, we extended the E_{Max} model to include a fully resistant cell population which is in competition with the proliferating and arrested cells (Figure 5.14a). This provides a simple model of patients who initially respond but subsequently experience disease progression whilst still under treatment. Indeed, simulations of this model under CT at MTD yield the typical picture of competitive release: temporary remission, followed by regrowth despite continuing treatment (Figure 5.14b; TOV112d, Ovc3, and Ovc4). In contrast, application of the newly proposed AT algorithm is predicted to result in reductions in cumulative drug use in three out of the four cell lines, and extensions of TTP in two. Interestingly, there are significant differences in the response between the different cell lines, with A2780s predicted not to respond much under either CT or AT. In addition, we find that we require a large resistance cost ($r_R = 10\% r$), and a very small initial resistance fraction ($R(0) = 5 \times 10^{-3}\%$) to obtain at least temporary remissions in our

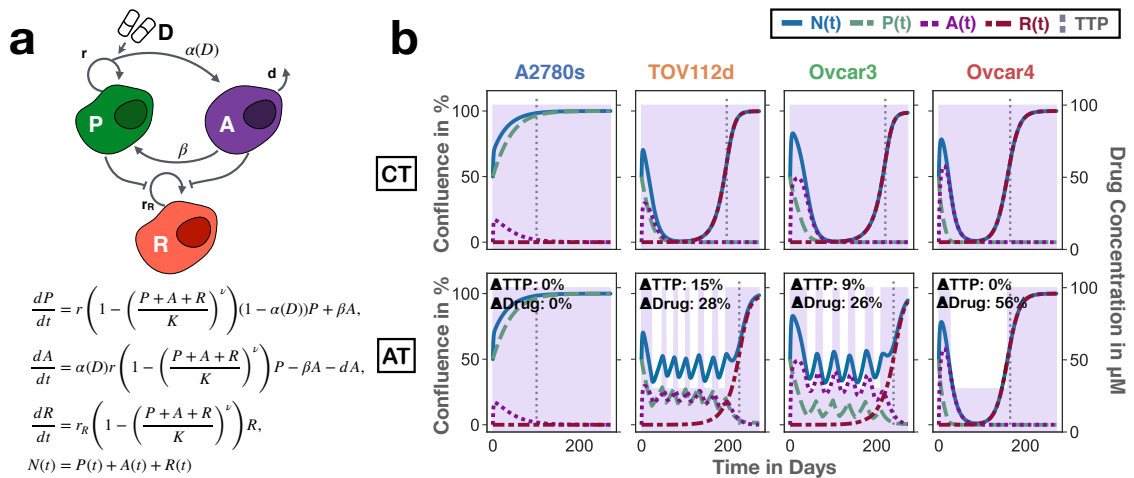


Figure 5.14: Proof-of-principle investigation of the proposed PARPi AT algorithm (Equation (5.6)). **(a)** A simple model of PARPi resistance, assuming it is driven by a fully drug-resistant population, which is initially rare, and competes for resources with the proliferating and arrested cells ($\alpha(D)$ given by Equation (5.5)). **(b)** Simulations of the model in (a) for different cell lines treated with either CT at MTD or the new AT algorithm ($\tau_{AT} = 20\%$, $D_{Min} = 30\mu M$). Progression defined as a 20% increase in size from pre-treatment levels. This shows that the new algorithm may be able to improve toxicity by reducing the cumulative drug use, whilst also potentially helping to slow the expansion of resistance. As such, we propose further research into this direction. Parameters: $N(0) = 50\%$, $A(0) = 0\%$, $R(0) = 5 \times 10^{-3}\%$, $r_R = 10\% r$; others as in Tables 5.1 & 5.2; treatment updates every 3 days.

simulations. As such, further work is required to model resistance and intra-tumoral competition more accurately and to investigate, for example, the choice of D_{Min} (the value of $D_{Min} = 30\mu M$ used here was chosen simply for illustration). Nevertheless these results indicate that the proposed new algorithm is a promising candidate for a potentially better tolerated PARPi schedule which may also help to delay the expansion of resistance, and warrants further investigation.

5.4 Discussion

The aim of this chapter was to investigate why AT had yielded inferior results to CT in the spheroid experiments in Chapter 4. To do so, we carried out 2-D *in vitro* experiments in which we systematically interrogated the cell's treatment response. Our results indicate that the reasons for the inferior performance of AT are two-fold: i) the tested AT algorithms reduced treatment too much or too often, thereby losing control over the drug-sensitive cells, and ii) the time frame of the experiments was too short to detect any benefit in controlling resistance, even if one was present.

To investigate whether it is possible to de-escalate treatment whilst maintaining control of the sensitive population, and whether there may be a benefit to AT in the long term, we leveraged the data we had collected to develop a mechanistic, mathematical model of PARPi treatment. To the best of our knowledge this is the first such model, and it is carefully derived from, and validated with, experimental data. Its key feature is that we explicitly account for the fact that PARPis do not directly induce cell death, but first cause cell cycle arrest. In this way we can explain why the population initially continues to expand during treatment, and why prolonged treatment is required before a size reduction is achieved. Interestingly, while arrested cells can also recover from the drug-induced damage in principle, our model fits suggest that this is not the case in our cell lines (values of β around 0 in Table 5.2). This is contrary to our expectations, given that nominally all cell lines but A2780s have functioning BRCA [201] and should, to some extent, be able to recover. While we did not sequence our cells to verify this, and it is possible that HR capabilities might have been lost due to mutations in other genes, a more plausible explanation is that this is due to an issue with parameter identifiability in the model. As such, an important next step would be to carry out a parameter identifiability analysis to determine an experimental protocol which allows robust inference of all parameter values. In addition, it would be useful to further validate the model by comparing the model's predictions against experimental measurements of PARPi-induced arrest, such as the abundance of γ H2AX foci [222].

A second defining feature of our model is that we assume a non-linear drug-effect relationship. While we did not have direct experimental validation, we found that we were unable to fit our data when assuming a linear model. Moreover, when we devised a simple method to estimate the relationship from the data at hand, we found a non-linear curve, which we could fit with a Hill function. Interestingly, the curvature of the observed relationship is concave, which suggests that there may be synergistic interactions between different PARPi-induced lesions in a cell. That being said, some care should be taken with the biological interpretation of the dose-effect relationship. It does not directly represent a biological mechanism, but rather is a phenomenological summary of the multi-step and multi-scale process that is drug action. In addition, the relationship depends on the particular assumptions of the model, such as that it is only the induction of cell cycle arrest, but not the repair or the rate of apoptosis, which depends on the drug concentration. Going forward it would be useful to obtain experimental data on the relationship between these processes and drug concentration, therefore allowing direct validation of our model

assumptions. Moreover, in our analysis we find that the choice of training data (how many different concentrations, and which concentrations to use) strongly influences the inferred treatment parameters. In particular, our measurements at $50\mu\text{M}$ appear to show less treatment effect than would be expected from the model. This experiment should be repeated to establish whether it is an erroneous measurement, or whether the dose-effect relationship is more complex than currently proposed. Furthermore, the uncertainty in the parameters should be incorporated into the model predictions.

In the next step, we used our model to carry out preliminary investigations into PARPi scheduling. This showed that it is possible to reduce treatment whilst maintaining control over the tumour, but that complete withdrawal should be avoided. The first conclusion is consistent with the wide therapeutic window of Olaparib observed in the clinic (see Section 5.1.1.2). The second conclusion has potentially important implications for the management of drug-related toxicities. The current toxicity management plan for Olaparib allows for either dose reduction or interruption, and, for example, in the Phase III study by Moore et al [17] interruptions (52% of adjustments) were more common than reductions (28% of adjustments). Furthermore, some practical guidelines have recommended interruptions over (long-term) dose reductions in managing side effects [239, 240]. In light of our findings, we strongly advocate further research into the effects of interruptions, and whether short-term dose reductions (if possible) should be preferred instead.

While our work corroborates that CT at MTD is very effective at managing the disease initially, emergence of resistance is a challenge in the clinic. Based on our insights, we propose an alternative adaptive algorithm, which is a simple modification of the 50% rule by Zhang et al [97]. Although significant further research is required, this protocol is potentially appealing for a number of reasons. Its lower cumulative dose may reduce toxicity, whilst it may also slow resistance through competitive suppression. Furthermore, it is plausible that because the periods of high dose treatment are short, one might be able to administer higher doses than current MTD, and thereby target cells which are only partially resistant to treatment, as has been discussed in targeted therapies [56]. In the next step of this research we propose to test whether our model can describe the treatment dynamics in patients, and to extend the model to capture the emergence of resistance and treatment-related toxicity. Extraction of suitable data from Moffitt's research database is already in progress. Based on this we will be able to draw a more accurate picture of the feasibility and potential benefit of the proposed algorithm in the clinic, and decide on future experimental testing.

Chapter 6

Discussion and Future Work

The aims of this thesis were to improve our understanding of which patients may stand to benefit from AT, and to investigate how we might tailor existing protocols to increase this gain. To do so, we performed a series of studies using different mathematical modelling techniques, which we integrated with analysis of clinical data and experimental work. Our work was split into two parts: the first part (Chapters 2 & 3) in which we performed a more theoretical analysis of AT and focussed on prostate cancer, and the second part (Chapters 4 & 5), in which we contributed to preclinical experiments and new models for resistance management in ovarian cancer. Overall, this thesis has made contributions to the theoretical understanding of AT, but has also brought to light practical issues which may be encountered in the translation of theory into practice.

To begin with, in Chapter 2 we aimed to establish a better understanding of the factors which determine whether or not AT will be successful. While a number of studies had shown that AT could significantly extend TTP, the diversity in the theoretical and experimental model systems made it difficult to identify precisely which factors were required. In particular, the role of resistance costs was a subject of controversy. As such, we set out to develop a minimal working example of AT, for which we chose to study a simple extension of the classical Lotka-Volterra competition model. By leveraging the work by Hansen et al [1], we were able to develop a deep and general understanding of AT, which can also be easily extended to other AT models. Since then, Viossat and Noble [126] have formalised the same approach to propose an elegant and comprehensive theoretical framework for AT, and we highly recommend their manuscript to new and experienced modellers in the field (see [127] for the final, published version of this work). Nevertheless, we believe that our manuscript [125] will also play its part, in particular by helping to communicate this insight to experimental and clinical colleagues by illustrating it using phase planes. In fact,

our artistic collage based on these phase planes was chosen as the cover of the print edition of the *Cancer Research* issue in which the paper will appear (February 15th, 2021), which will help raise awareness of this result and the field more generally.

In addition, the second key result from Chapter 2 was that cell death may be an important factor in AT, as it sets the pace of evolution in the tissue. The higher it is, the more strongly suppression by resistance costs or competition will affect resistant cells. While its impact is weaker in the ABM, this conclusion is confirmed also by our subsequent analysis in Chapter 3. There are two implications of this result. On the one hand, it indicates that turnover is a factor which should be accounted for when drawing inference on AT from different *in vitro* or *in vivo* experiments. On the other hand, it suggests a potential target to help improve AT. For example, while Abiraterone was given adaptively in the Zhang et al [97] prostate cancer AT trial (NCT02415621), patients were concurrently receiving a continuous dose of Lupron. Lupron suppresses systemic testosterone production, thereby reducing the cancer's environmental carrying capacity (supply of growth factors) and increasing cell death due to starvation. In light of our findings, it would be useful to investigate what role this secondary drug has had in the success of the trial, to help with the design of future multi-drug strategies.

In Chapter 3 we sought to corroborate our insights, and study the impact that space has on AT. Again, spatial models had been considered before, but as these had involved additional assumptions about nutrients and resistance, it was not clear which results were due to these extra assumptions, and which were due to space. We, thus, built an ABM realisation of our ODE model and repeated our analyses from Chapter 2. This showed that the models are qualitatively consistent, and provides evidence that the principles which we (and Hansen et al [1] and Viossat and Noble [126]) have developed will be also applicable to spatial systems. That being said, our analysis also revealed significant quantitative discrepancies. This is not unexpected, but it shows that an in-depth understanding of intra-tumoral competition for the specific clinical setting will be required to evaluate whether the theoretically predicted benefits of AT will be practically meaningful.

What does intra-tumoral competition during AT actually look like? To address this question, we explicitly quantified competitive suppression of resistance in our model, the first time this has been done in AT. While it is theoretical, and not empirical evidence, our results do help to illustrate and validate the theory underlying AT. Moreover, our analysis shows that a key difference between the non-spatial ODE

model and the spatially explicit ABM is the intra-specific competition between resistant cells. The ODE assumes homogeneous spatial mixing, so that the balance of inter- vs intra-specific competition simply reflects the relative frequencies of the two cell populations. In contrast, in the ABM there is no mixing, so that cells segregate, and competition of resistant cells with each other becomes an important factor. A key implication of this observation is that in solid tumours, in which the extent of cell mixing is limited, the spatial distribution of the resistant cells will likely play a role in how well we can control them with AT. In the future, it would therefore be interesting to compare how different modes of resistance evolution, which will result in different spatial distributions (e.g. mutation-driven vs phenotype-driven vs environmentally-driven), affect outcome under AT. Moreover, the role of non-tumour tissue should be explored, both as an important further competitor for resources, as well as for its role in shaping tumour architecture.

In addition to developing theoretical understanding of AT, the first part of this thesis investigated the cycling dynamics of prostate cancer patients undergoing intermittent androgen deprivation therapy. Our analysis provides evidence for the hypothesised relationship between resistance cost and cellular turnover, and makes two predictions which warrant further investigation. Firstly, we observe a strong correlation between the cycling speed of the patients and the fitted cell death parameter, and secondly, fitting the ABM suggests that these differences may also be reflected in distinct spatial tumour architectures. While the inferred cost and turnover values are larger than suggested by the literature (although this could reflect the harsher conditions in human tissue compared to laboratory settings), and the models are highly simplified depictions of the biology, the strength of the correlation is striking. Ideally, these predictions should be tested by analysis of clinical data, or via wet-lab experiments, although costs and data availability make this more of a long term goal. A more short term validation could be provided by testing if these predictions can be replicated in more realistic *in silico* models of prostate cancer (e.g. [241]). Overall, we advocate further research into previous intermittent treatment studies in prostate cancer - the majority of which found neither a benefit nor a disadvantage of treatment cycling over CT [110, 111, 242] - in order to solidify our understanding of why the Zhang et al AT trial [97] was successful.

In the second part of this thesis we turned our attention to ovarian cancer and the question of whether adaptive PARPi maintenance protocols may help to delay progression in this setting. In particular, in Chapter 4 we tested adaptive strategies in which the treatment decision was not made by a pre-determined “rule-of-thumb”

but was guided by the predictions of one, or more, mathematical models. Notably, we performed this work not in the abstract, but in collaboration with an experimentalist using *in vitro* spheroids. Given the complexity of the problem, it feels like this part of the thesis raised more questions than it provided answers. That being said, we believe that the concept of using not a single, but an ensemble of models, and the challenges we highlighted in the practical testing of AT, represent valuable contributions to the field. Going forward, we propose to further develop our ensemble approach, starting by testing it on *in silico* models from the literature, as outlined in Chapter 4. In this way we will have precise control over the characteristics of the tumour we treat, and we can evaluate how mis-matches in assumptions affect the predicted treatment. In fact, testing a new “optimal” treatment schedule on other *in silico* models, much like the way in which new drug candidates are tested on multiple experimental models, could be a strategy for evaluating the robustness of theoretical findings more generally.

Finally, Chapters 4 & 5 also make contributions to our understanding of Olaparib treatment of ovarian cancer. In particular, our results suggest that while there may be scope for dose reductions, complete treatment withdrawal should be avoided. This observation has potentially important implications for toxicity management, and should be further investigated. Moreover, this work demonstrates that when developing AT protocols in practice the specific treatment and disease dynamics have to be taken into account. We learned this the hard way in Chapter 4, where even an ensemble of different models only has limited success in maintaining control of the tumour. This is partially because none of the models we chose accounted for the delay in drug response, and partially because we only allowed for treatment to be modulated in a binary on/off fashion.

To address this problem, we carried out *in vitro* experiments in Chapter 5 in which we carefully dissected the treatment response, and we leveraged the self-collected data to develop, calibrate, and test a more accurate mathematical model of PARPi response. To the best of our knowledge, this is the first model of PARPi treatment, and we believe that because of its good ability to fit, and predict responses to different treatment schedules, it will be a valuable tool for the study of PARPi scheduling in the future. In fact, by leveraging this model, and integrating the various insights we had gained throughout the thesis, we developed a new AT algorithm which we propose warrants further investigation. Going forward we suggest that our model should be further validated in the wet-lab, for example, by testing its response predictions for more complex schedules. In addition, it should be compared against clinical data to test whether it translates to patients, and so to help us learn more about the

strength of intra-tumoral competition in patients and to integrate the emergence of resistance into the model. Furthermore, inclusion of PARPi-induced toxicity, and drug PK would be valuable model extensions. Together such an holistic model will allow for a more reliable assessment of whether or not AT may be able to help delay PARPi resistance, and will provide guidance for the design of subsequent validation experiments.

To conclude, we want to highlight what we consider to be the two greatest weaknesses of this thesis, and suggest how these could be addressed. Firstly, throughout this thesis we have given only limited consideration to the issue of parameter identifiability when fitting our models to data. While we have always repeated model fitting from different initial conditions to address the issue of possible local optima, we have not quantified the confidence in our parameter estimates. This leaves open the possibility that different parameter sets to the ones identified by us may fit the data equally well. This is problematic in Chapters 2 & 3 because it is possible that the correlation between cost and turnover estimates may in fact be due to structural or practical identifiability issues within the model, rather than being a real phenomenon. While parameter sweeps, such as that in Figure 2.9, show that the two parameters do have distinct impact on the cycling dynamics, it would be important to ascertain the confidence in our parameter estimates by, for example, computing confidence intervals based on the Fisher information matrix [243, 41, 244].

Similarly, in Chapter 4, knowing the possible error in the parameter estimates, and incorporating this uncertainty into the model predictions, would have allowed for more informed decision-making, and would have helped in deciding how to weigh the suggestions of different models in the ensemble. An elegant and powerful way to perform such error propagation would be through the use of mixed effect models. This means certain parameters would be drawn from statistical distributions which are learned from the data and which allow us to predict the range of possible outcomes [24, 25]. Such an approach would also allow us to better leverage the structure of the data in Chapter 4 (multiple independent replicates per condition; different related conditions), and improve parameter identifiability. Finally, mixed effect models can also be combined with Bayesian inference methods. This approach, in which “prior” knowledge from historical data or early time points is gradually updated as more data for the individual patient become available, is regularly employed in model-guided TDM (e.g. [89, 96, 87]), and would be a valuable extension to our model-guided AT framework.

The second key limitation of our work is that most of the models studied in this thesis do not account for the potential risks associated with maintaining a higher tumour burden, an issue also raised in the wider AT literature [126, 170, 129]. Indeed, a higher tumour burden may, for example, result in a higher rate of *de novo* resistance acquisition via mutation, or a higher rate of metastasis. If these rates become sufficiently high, then it may in fact be preferable to reduce tumour burden via MTD treatment [1, 126]. As such, it would be important to better quantify these risks. This could be done by explicitly including mutation in the model, as done by Viossat and Noble [126] or Hansen et al [1, 129] (see also Appendix C), or by extending models of metastasis, such as that by Iwata et al [245]. In addition, as recently highlighted by Mistry [170], by building on the work of the TGI modelling community (see e.g. Benzky [22] or Bruno et al [23] for reviews) we may directly link tumour burden with clinically calibrated and validated statistical models of survival. For example, Cao et al [246] recently developed a joint model to capture the time-course dynamics of CA-125 levels in ovarian cancer patients and predict patient survival. Interfacing such a model with models of intra-tumoral competition may provide a means to more accurately quantify the benefit of AT in ovarian cancer. More generally, building on the tools in the TGI literature may allow us to better connect our models with clinical data, and would encourage interaction with the pharmacometric modelling community and their expertise in model calibration and TDM.

When we began working on this thesis, many of the basic principles of AT were just beginning to be understood. In order to tackle this problem, mathematical (and experimental) models have often had to make highly simplifying assumptions about the tumour biology, the action and delivery of treatment, data availability, and possible risks of de-escalation (this thesis being no exception). We believe that the remarkable progress over the past ten years has now put us in a position from which we can, and should, engage with this extra complexity. Metastatic cancers are highly heterogeneous diseases, requiring simultaneous control of multiple, independently growing, and evolving, tumour lesions. In addition, the growth of these tumours shapes, and is shaped by, the microenvironment and the immune system, and these interactions are altered by treatment. Moreover, cancer therapy typically involves not a single, but a combination of different, drugs. As such, developing frameworks for multi-drug AT, studying how to adapt treatment in the light of multiple target lesions, and investigating how to integrate other aspects of quality of life (e.g. toxicity or psychological factors) should be key questions to address in the future.

Chapter 7

Conclusion

“Can insects become resistant to sprays?” - This was the question with which we started this thesis, and as Melander [2] had predicted, the answer turned out to be a resounding *yes*. Similarly, antibiotic resistance is becoming an increasingly pressing problem [247], and it is a strange turn of events that, at the time of writing this, the emergence of vaccine-resistant Covid19 strains is dominating global headlines.

Also in cancer treatment, resistance is a key challenge. Treatment options exist for most cancers, but all too often they only provide partial and/or temporary remission. In this thesis, we explored a new approach in which the objective of treatment is shifted from minimising burden to delaying resistance by capitalising on competition between tumour cells. The guiding questions throughout our work were: i) when will such an approach be successful?, and ii) how do we adapt treatment to maximise this benefit? Using an integration of mathematical modelling, wet-lab experiments, and clinical data analysis, we showed that key factors in answering these questions are: i) the efficacy of treatment, ii) the initial resistance fraction, iii) the environmental carrying capacity, iv) resistance costs, v) the tumour’s turnover rate, vi) the tumour’s spatial organisation, and vii) potential harm caused by the higher tumour burden. Importantly, our work also showed that these factors interact in a non-linear fashion so that this isn’t a simple “tick-list”, but that careful assessment is required when choosing whether, and how, to apply AT. A century after Melander’s first report, the San Jose scale is approached today through integrated pest management, in which a combination of interventions is strategically tailored to the local environmental and economic context [248]. We believe that, similarly, in cancer we will be able to achieve the best outcomes by adapting treatment modality and, importantly, objective, according to the eco-evolutionary characteristics of the tumour and the wishes of the individual patient. Given the complexity of this problem, inter-disciplinary team science which integrates, and learns, from other disciplines, will be the key to success.

Appendix A

Appendix for Chapter 2

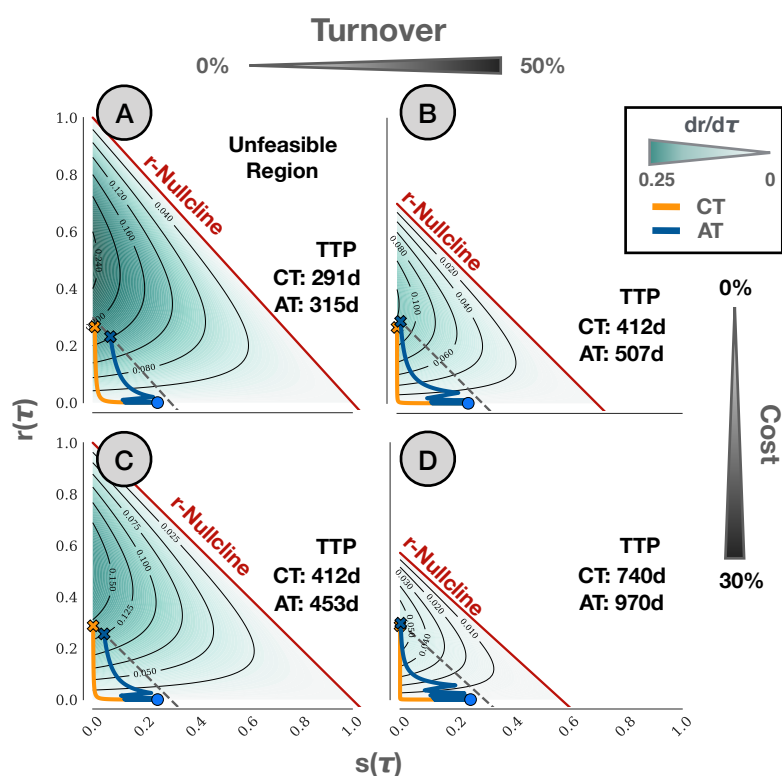


Figure A.1: The resistance growth rate, $dr/d\tau$, as a function of s and r with and without cost and turnover. Overlaid on top are the treatment trajectories for Tumour 1 under AT and CT. A) The closer the tumour (blue dot) is initially to the r -nullcline (red line), the longer AT can reduce $dr/d\tau$ and prevent progression (dashed grey line). B) Turnover shifts the nullcline, thereby reducing the distance of the tumour from the nullcline. C) In the absence of turnover, a cost of resistance merely rescales $dr/d\tau$ but does not change the distance to the nullcline, so has less effect. D) In contrast, in the presence of turnover, a resistance cost both reduces the distance to the nullcline and re-scales the surface.

Appendix B

Appendix for Chapter 3

B.1 Convergence and consistency analysis

Since we are interested in performing quantitative, as well as qualitative, comparisons between the ODE and the ABM, we performed a number of tests to verify the correct implementation of the ABM. Furthermore, we carried out a convergence and a consistency analysis, in order to choose an appropriate model time step and sample size.

B.1.1 Testing the model implementation

In our ABM we are assuming that the rates at which cells attempt to divide are independent of time. As such, if we were to seed a cell in isolation, so that its division is not constrained by neighbouring cells, we would expect the time between divisions (the cell cycle time) to be exponentially distributed [249]. Furthermore, we would expect that as $dt \rightarrow 0$, the mean of this distribution would converge to $1/r_S$ or $1/r_R$, respectively. In Figure B.1 we show that this is indeed the case. When $dt > 1/r_S$, the cell will divide after one time step, so that the only observed value for the cell cycle time is dt (Figure B.1a). Once $dt < 1/r_S$ we see a range of values, which approaches the expected distribution as $dt \rightarrow 0$ (Figure B.1a). To test whether the mean of this distribution converges to the expected value, we plot the mean value of the observed cell cycle time distribution against the time step, dt (Figure B.1b). This shows that the algorithm produces mean cell cycle times that are within 10% of the true value for a large range of time steps and confirms convergence. To ensure the algorithm is accurate across the range of proliferation rates we assessed in Chapter 3, we repeated the above analysis for the smallest value of the proliferation rate in our range of interest ($r_S = 0.0135d^{-1}$, corresponding to a 50% fitness cost). The results are shown in Figure B.1c and confirm the algorithm's robustness. Since we

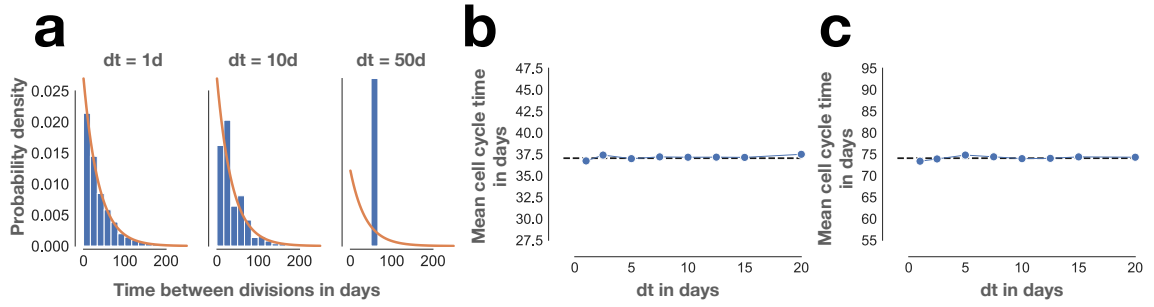


Figure B.1: Testing the implementation of the ABM. (a) Cell cycle time distributions of cells in the ABM for different time steps, dt . Single sensitive cells were seeded and the time to division recorded (shown are the results from 10,000 independent simulations). As dt is decreased, the cell cycle times approach an exponential distribution with mean $1/r_S = 37.03d$ (the probability density function of this distribution is shown as a solid orange line). (b) Mean of the observed cell cycle distributions as a function of the time step, dt ($dt \in \{20., 15, 12.5, 10., 7.5, 5., 2.5, 1.\}$; 10,000 independent simulations per value of dt). This shows that the algorithm produces average cell cycle times which are close to the expected value (dashed line), for a large range of time steps, dt . (c) Same as (b) but with $r_S = 0.0135d^{-1}$, corresponding to a 50% fitness cost. This confirms that the algorithm is accurate across the range of proliferation rates examined in Chapter 3. Parameters: $d_T = 0, D(t) = 0$. Domain size: $l = 100$.

assume constant death rates, we also expect the life time of a cell (time to natural death) to be exponentially distributed with a mean which approaches $1/d_T$ as $dt \rightarrow 0$. Analogous tests to the ones performed above confirm this (not shown).

B.1.2 Convergence analysis

An important aspect of the simulation algorithm is the time step, dt . It needs to be small enough to provide stochastically accurate results yet also be large enough to minimise run time. In order to study the sensitivity of the model results to the time step, and identify a suitable time step for our analysis, we performed a convergence analysis. In this analysis we compared the distributions of treatment results for a set of different time steps. Specifically, for each value of dt , we computed the distribution of relative time gained ($TTP_{\text{Gained}} = \frac{TTP_{\text{AT}} - TTP_{\text{CT}}}{TTP_{\text{CT}}}$) based on 1000 independent replicates, where each replicate was seeded with a different random number seed and a different random initial condition. We repeated this three times to obtain three independent estimates of the TTP_{Gained} distribution per value of dt . We considered four different tumour parameter sets, which were the four edge cases of possible initial

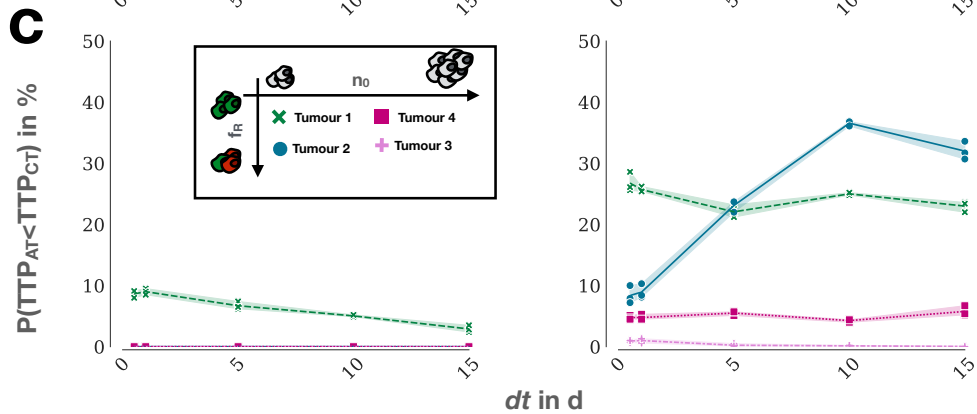




Figure B.2 (*previous page*): Convergence analysis. **(a)** Convergence was assessed based on the distribution of relative time gained by AT ($\frac{TTP_{AT} - TTP_{CT}}{TTP_{CT}}$). For each value of the time step, dt , three independent estimates of this distribution were obtained, each based on 1000 independent runs of the model. Shown here are the distributions obtained for Tumour 2 ($(n_0, f_R) = (25\%, 10\%)$), assuming a 30% turnover rate. As the time step is refined, the shape of the distribution converges. Moreover, there is little variation between the three independent replicates at each time step, suggesting this result is consistent across runs. **(b)** To quantify the effect of dt on the outcome distribution we compare the mean relative time gained. This value remains relatively unchanged by dt , except for Tumour 4 with turnover. **(c)** As a second way of quantifying convergence, we compare the proportion of runs in which CT outperforms AT ($\mathbf{P}(TTP_{AT} < TTP_{CT})$). This is again only moderately affected by dt , except for Tumour 2 in the presence of turnover. We conclude that while the time step does affect the results, under most circumstances it has relatively little impact. In the cases where it does have a strong impact, a value of $dt = 1$ provides representative results, whilst still allowing for fast simulations, and is thus chosen for all analyses in the chapter. Points correspond to individual replicates. Lines and shading indicate the mean value and 95% confidence interval. Tumour 1 ($(n_0, f_R) = (25\%, 0.1\%)$), Tumour 2 ($(n_0, f_R) = (25\%, 10\%)$), Tumour 3 ($(n_0, f_R) = (75\%, 10\%)$), and Tumour 4 ($(n_0, f_R) = (75\%, 0.1\%)$). We assume no resistance cost ($c_R = 0$). Domain size: $l = 100$.

tumour compositions from Chapter 2. We studied each in the absence of turnover ($d_T = 0$) and in the presence of turnover ($d_T = 30\%$). In all cases, we assumed no resistance cost ($c_R = 0$), since a time step which can accurately capture the faster dynamics in the absence of cost will also be suitable for the slower dynamics when a cost is present.

In Figure B.2 we show the results of our analysis. As we decrease the time step, we see convergence of the TTP_{Gained} distribution (Figure B.2a). We quantify convergence using two measures which are of interest in Chapter 3. Firstly, we consider the mean value of the TTP_{Gained} distribution. This shows that the mean value changes relatively little in response to changes in dt . Secondly, we compare the proportion of runs in which CT outperforms AT ($\mathbf{P}(TTP_{AT} < TTP_{CT})$). In 7/8 cases this proportion changes by less than 10% as a result of refining dt (Figure B.2b). The exception to this is Tumour 2 with turnover (Figure B.2b). However, in this case the sudden jump in $\mathbf{P}(TTP_{AT} < TTP_{CT})$ can be explained by the fact that 2954/3000 simulations predict a difference in TTP of less than 10d, which can only be resolved accurately when dt is sufficiently small (see also Figure B.2a).

We conclude that the predictions of the model are reasonably robust to the specific

time step used. We choose to use a time step of $dt = 1d$, as this yields accurate results, while simulations are still relatively fast (100 simulations can be run in about 2 minutes on a desktop computer).

B.1.3 Consistency Analysis

A second factor which is important in the analysis of stochastic processes is the sample size. On one hand, we would like to perform sufficiently many replicates in order to obtain a sample of realisations which is representative of the stochastic process. On the other hand, we are required to again balance this with the computational burden. To identify an appropriate sample size for our analysis we performed a consistency analysis. In this analysis one chooses a candidate sample size, n , and then obtains k independent groups of n replicates. If the k outcome distributions thus obtained are statistically equivalent, then n produces a representative sample. If not, then a larger sample size is required. For a more detailed discussion of this technique we refer to the review by Hamis et al [175].

We again examined the same corner cases of the parameter space which we considered in the convergence analysis and, for each, obtained $k = 10$ groups of n samples, where we considered values for n of $n = 10, 50, 100, 250, 500, 1000$, and 1500 . For each of the n replicates we simulated both CT as well as AT in order to obtain one estimate of the value of the relative TTP_{Gained} . The same random number seed was used for the two matched simulations of CT and AT, but was changed between every replicate. In Figure B.3a we show the thus obtained series of estimates for the outcome distributions for one parameter set. While for small values of n the 10 estimates for the outcome distribution differ significantly, these differences disappear as we increase n . For $n \geq 250$ the distributions closely resemble each other.

Consistency can be quantified using measures of the distance between distributions, such as the A-measure, or the Kullback-Leibler divergence [175]. However, how do we decide that the distributions are sufficiently similar for our application of interest based on these rather abstract metrics? Thus, instead of using one of these statistical measures, we decided to assess consistency based on the metrics which we knew would be of interest in our study. As such, we again compared the mean relative time gained, and the proportion of runs in which AT performs worse than CT. The results of this analysis in the presence of turnover are shown in Figures B.3b & c, respectively, and corroborate that for $n \geq 250$, independent groups of replicates yield very similar results. Based on these results we choose to use a value of $n = 250$ for

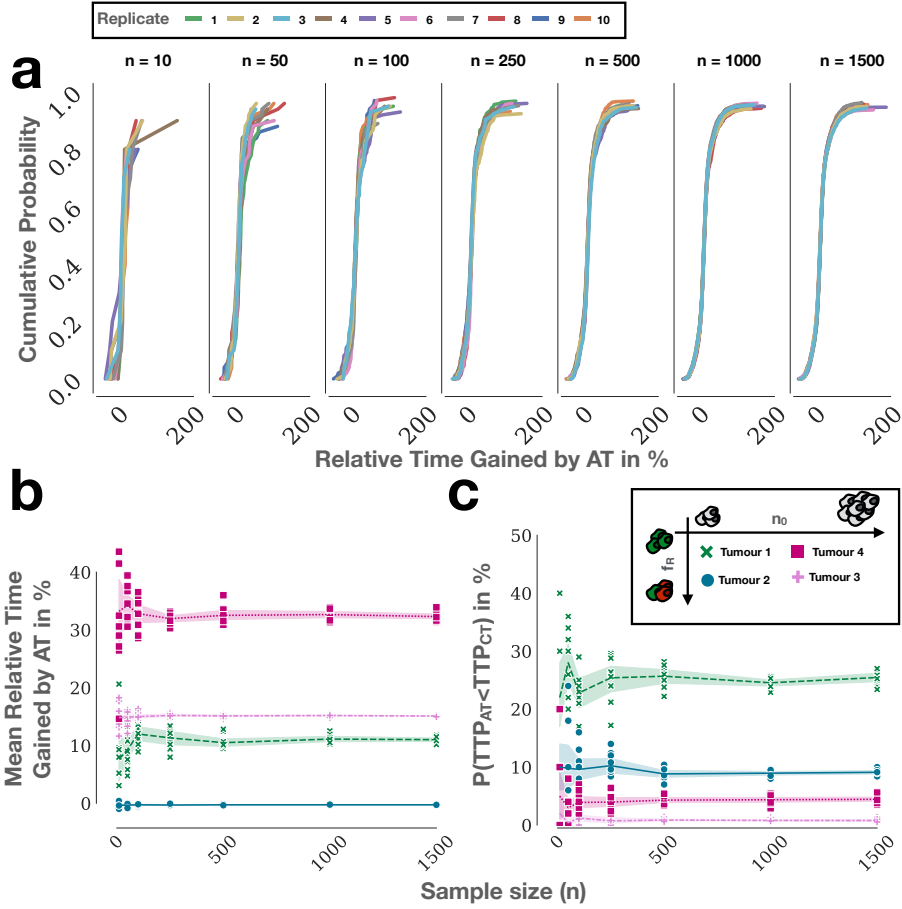


Figure B.3: Representative results of the consistency analysis: (a) Increasing consistency in the distribution of time gained by AT for one parameter set $((n_0, f_R, c_R, d_T) = (25\%, 0.1\%, 0\%, 30\%))$ as the number of independent replicates per sample, n , is increased. For each value of n , 10 independent samples of sample size n were collected. (b) Mean value of the outcome distribution as a function of sample size, n , for four parameter sets (Tumour 1: $(n_0, f_R) = (25\%, 0.1\%)$; Tumour 2: $(n_0, f_R) = (25\%, 10\%)$; Tumour 3: $(n_0, f_R) = (75\%, 10\%)$; Tumour 4: $(n_0, f_R) = (75\%, 0.1\%)$). Lines indicate the mean value, shading a 95% confidence interval. Markers show values of individual replicates. For sample sizes upwards of 250, the different replicates produce very consistent results. (c) Proportion of runs in which AT failed as a function of the sample size, n , for the same four parameter sets. Again for $n \geq 250$ we see consistent values between independent replicates indicating that a sample size $n \geq 250$ will yield a representative outcome distribution.

plots in which we show treatment dynamics, and a value of $n = 1000$ for all quantitative analyses. The reason we use the smaller sample size for the time-series plots is to reduce the computational burden on the plotting routine.

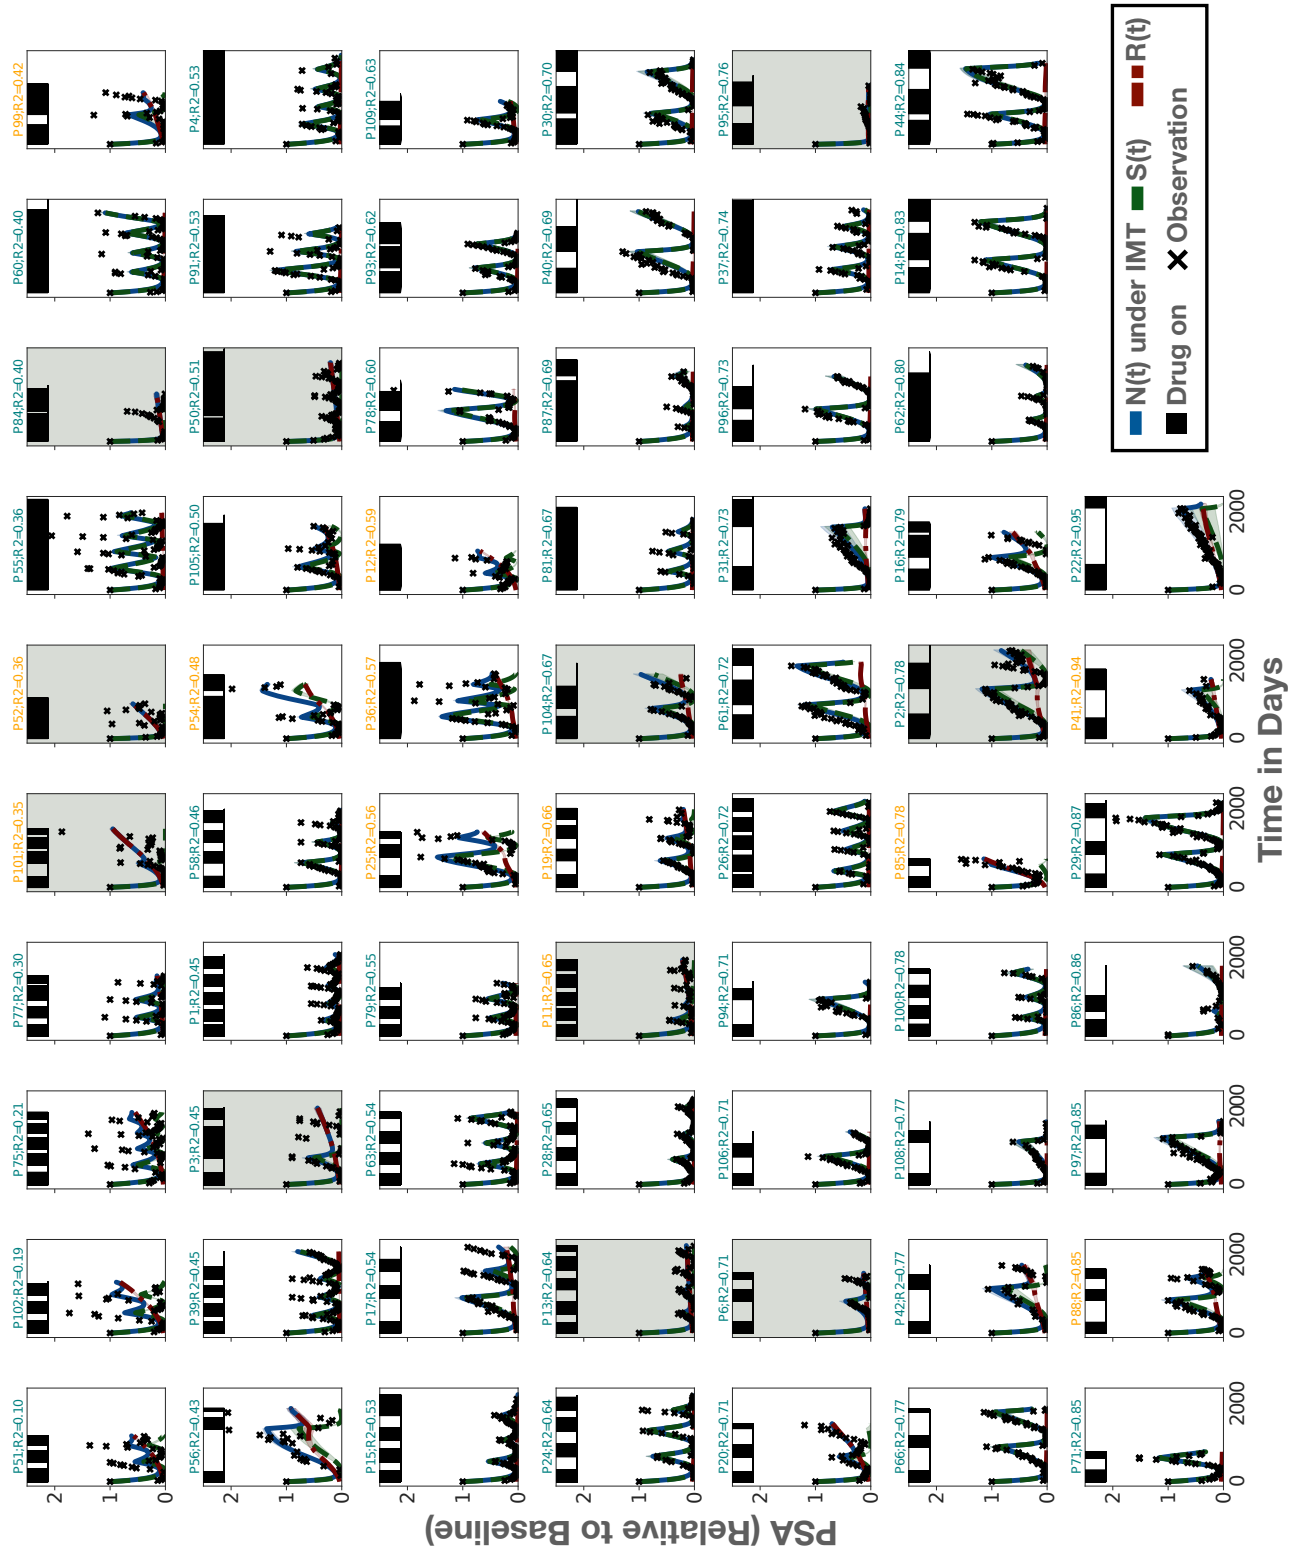


Figure B.4: Overview of the ABM fits of the reduced model (fitting only cost and turnover) for all 67 patients, arranged by their r^2 value (showing the mean and standard deviation of 25 replicates per patient). Title colour indicates whether a patient relapses (orange) or not (green). Patients who were excluded from further analysis due to poor model fits are marked with a grey background.

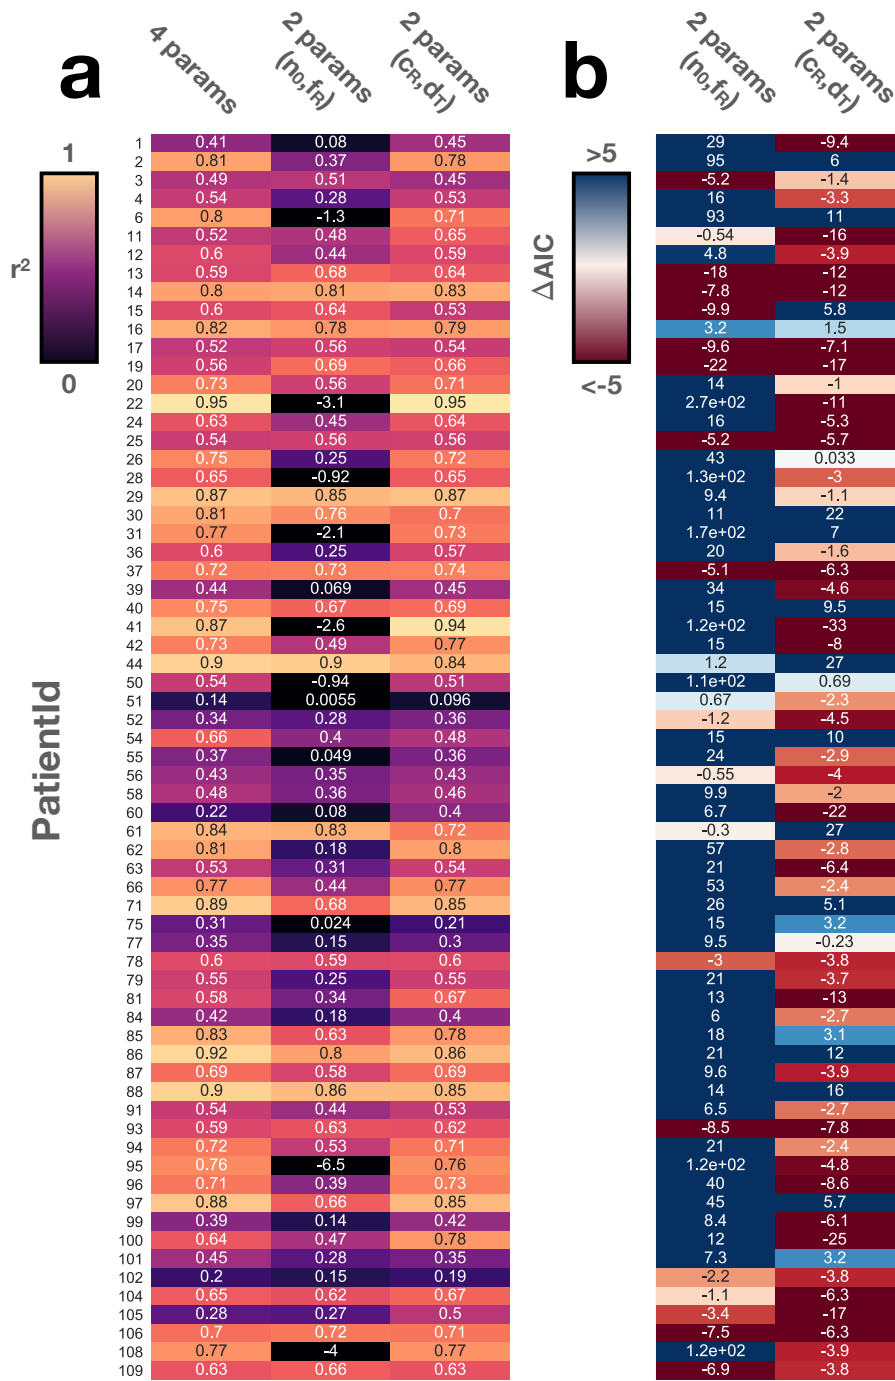
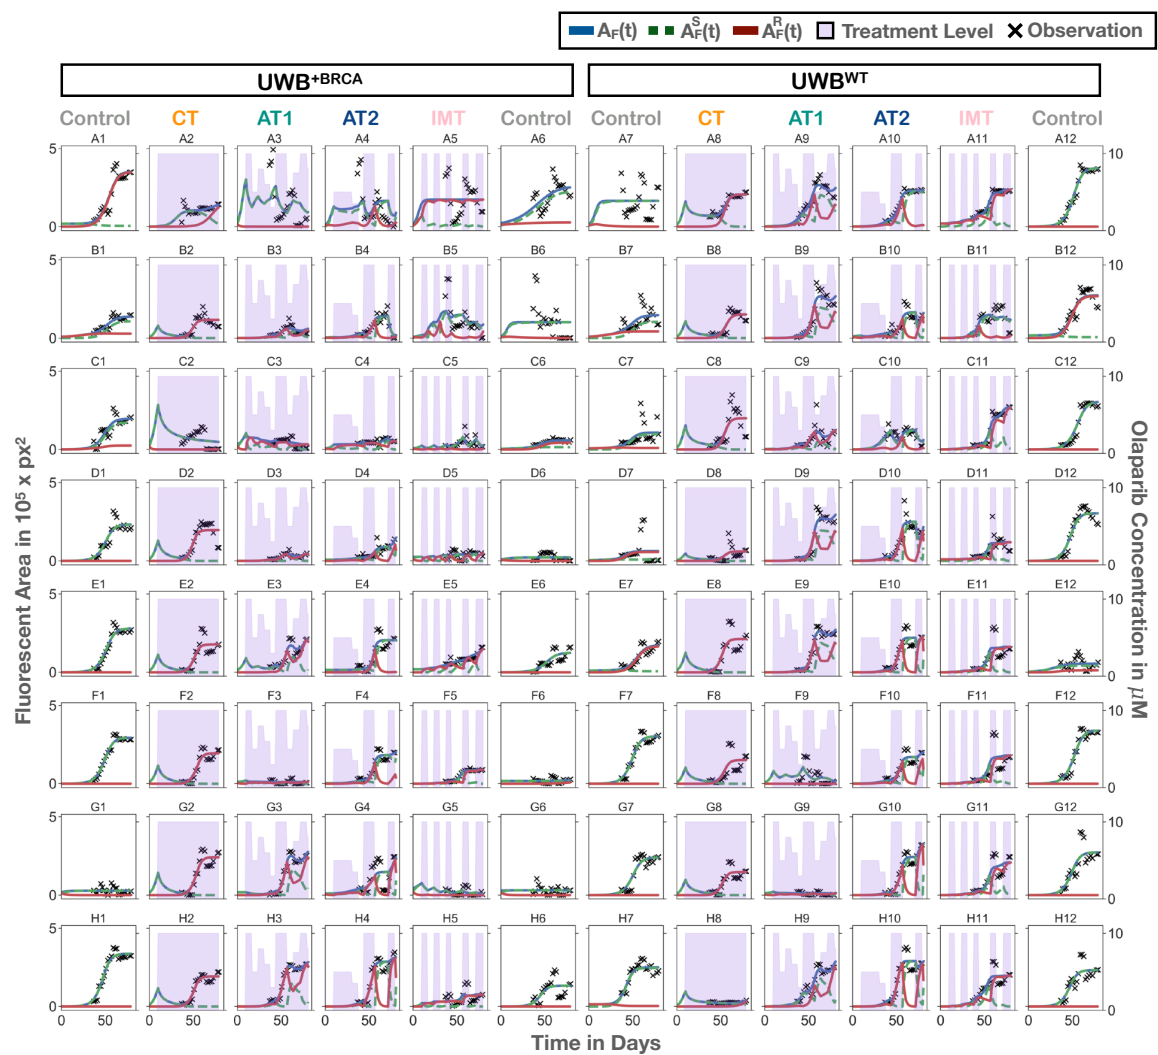


Figure B.5: Comparison of fitting the ABM by allowing either the initial tumour composition (n_0, f_R), the cell kinetics (c_R, d_T), or all four parameters to vary (“4 params”). (a) Comparison of the r^2 values for each patient for each model. We observe that allowing only cost and turnover to be patient-specific can explain the data almost as well as the full 4-parameter model. This is not true for the model in which the initial conditions are patient-specific. (b) Difference in AIC between the 2-parameter models compared to the 4-parameter models for each patient. This corroborates that for most patients the cost-turnover model provides the best description of the data.

Appendix C

Appendix for Chapter 4



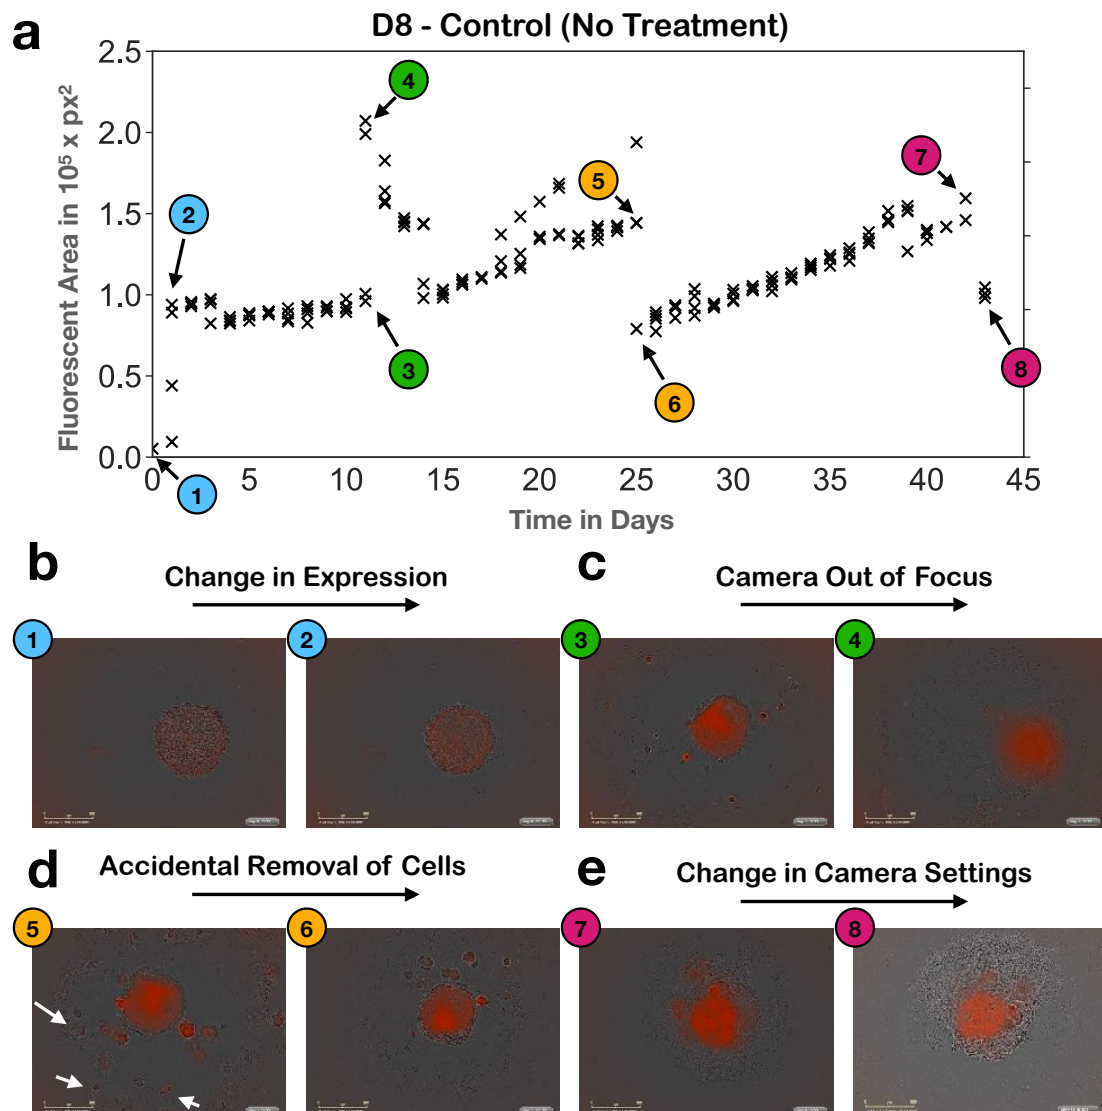


Figure C.2: Issues with imaging and data acquisition in the open-loop optimisation experiment. (a) Fluorescent area provided by the image analysis pipe line, showing discontinuities in the trajectory and tumour size decreases even though no treatment was administered. (b) Some jumps can be explained by temporary loss of fluorescence after medium change which, however, quickly re-established. (c) In other cases the spheroid was accidentally moved and is now out of focus, resulting in an inflated estimate of its size, due to the blurry image. (d) In some instances some cells were accidentally removed when the medium was aspirated (arrows). (e) Finally, due to sharing of the machine with other groups, sometimes the camera settings were changed, causing problems for the segmentation algorithm. Note: while Mr Wicker had written the Python script for the image analysis, this post-hoc analysis was carried out by me.




Figure C.3: Growth trajectories of all spheroids treated with the closed-loop, ensemble algorithm. Black crosses indicate curated data, used for model fitting and analysis. Outliers due to issues with imaging (e.g. blurry images, see Section 4.3.2.5) were removed (red crosses) or corrected (blue crosses).

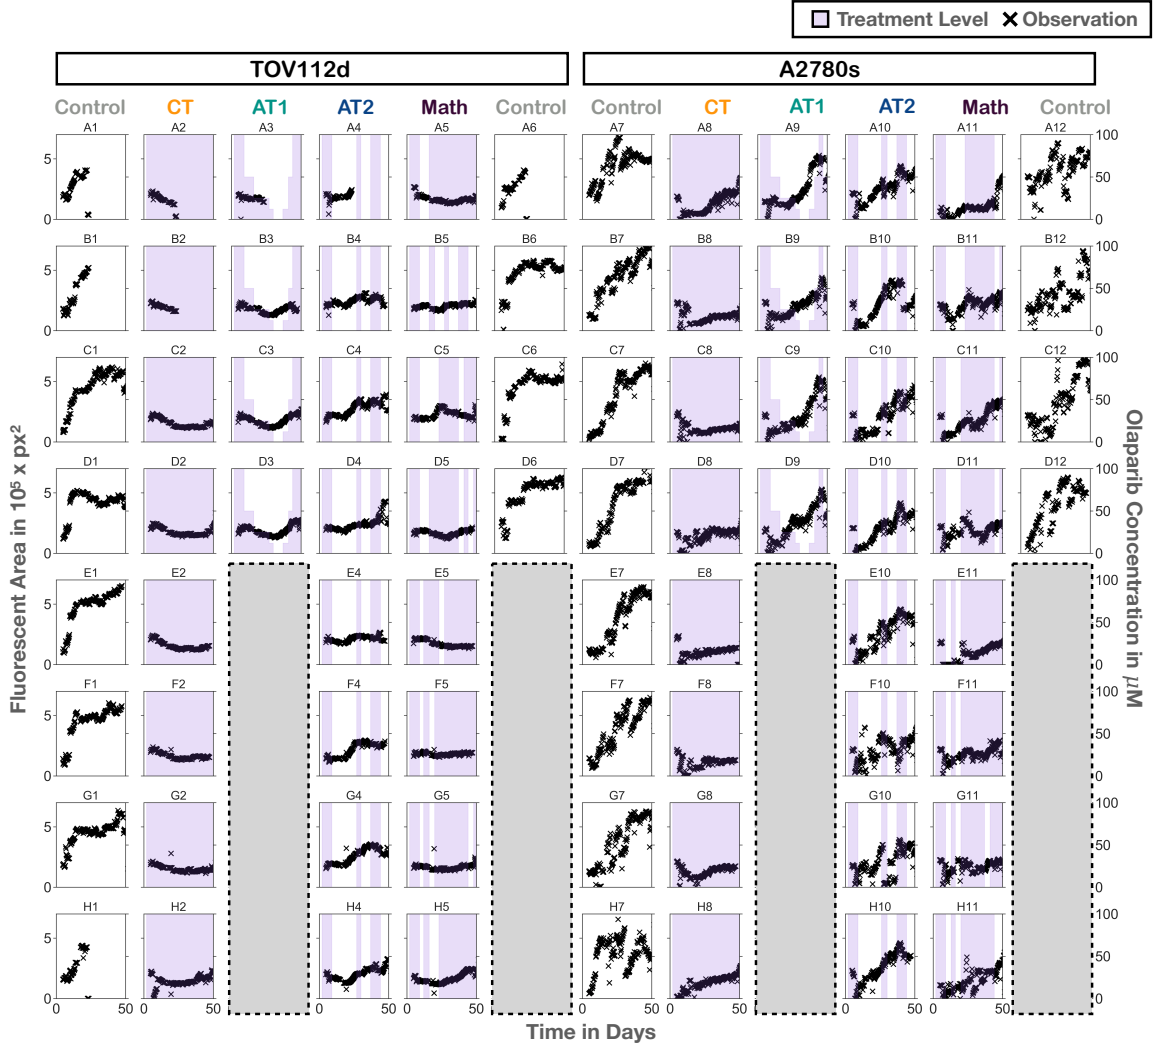


Figure C.4: Growth trajectories of all spheroids in the experiment in which we compared the ensemble framework to CT, AT1, and AT2. Shown are the raw (uncurated) data.

Table C.1: Spheroid growth under different treatment branches. Mean size is the average size across the length of the experiment. Square brackets denote 95% CIs and * indicate a statistically significant difference to the Math schedule at the 5% level in a Student t-test. All values are in units 10^5 px^2 .

Tx Branch	A2780s		TOV112d	
	Final Size	Mean Size	Final Size	Mean Size
Ctrl	5.25* [4.76, 5.74]	3.95* [3.87, 4.03]	4.81* [4.14, 5.49]	4.32* [4.25, 4.39]
CT	1.95 [0.71, 3.18]	1.21* [1.18, 1.24]	1.99 [1.3, 2.68]	1.6* [1.58, 1.62]
AT1	3.32 [2.68, 3.95]	2.34* [2.24, 2.43]	2.66 [1.89, 3.43]	1.8 [1.77, 1.83]
AT2	3.46* [3.02, 3.9]	2.33* [2.27, 2.39]	2.72* [2.2, 3.23]	2.37* [2.34, 2.4]
Math	2.72 [2.1, 3.35]	1.59 [1.55, 1.64]	2.24 [1.91, 2.57]	1.82 [1.8, 1.84]

C.1 The impact of cell plasticity on AT

In their first-in-human clinical trial of AT, Zhang and colleagues [97] chose a 50% reduction in serum PSA levels as the threshold for treatment withdrawal. Treatment was re-started once PSA levels returned to pre-treatment levels, where the assumption was that the changes in PSA reflect corresponding changes in tumour burden. However, we show in Chapters 2 & 3 that this treatment algorithm may not be optimal. Both the ODE model and the ABM predict that smaller burden reductions would have yielded longer tumour control, a finding also made by others [123, 126, 182].

That being said, the reason why tumour control is maximised by the adaptive algorithm with the smallest burden reduction is because of our model assumptions. As we discussed in Chapter 2, we assume that there is no “cost” associated with maintaining drug sensitive cells, but only the benefit of reduced resistant cell growth. It follows that in order to maximise control we should maintain tumour burden at the maximal allowable level. Yet, with an increased tumour burden there may, for example, be an increased chance of the formation of new resistance mutations. Hansen et al [1] have shown that this can mean that maintaining sensitive cells may actually speed up progression and that one should, instead, aim to eradicate as many cells as possible - even if this results in the competitive release of resistance.

To illustrate this let us consider a simple extension of the ODE model studied in Chapter 2 (Equations (2.1)-(2.3)) which includes the possibility of drug-induced switching from sensitive to resistant cells:

$$\frac{dS}{dt} = r_S \left(1 - \frac{S+R}{K}\right) (1 - d_D D(t)) S - d_S S - \alpha D(t) S, \quad (\text{C.1})$$

$$\frac{dR}{dt} = r_R \left(1 - \frac{R+S}{K}\right) R - d_R R + \alpha D(t) S, \quad (\text{C.2})$$

$$N(t) = S(t) + R(t), \quad (\text{C.3})$$

where α denotes the rate of switching. The initial conditions are given as before by:

$$N(0) = N_0, \quad S(0) = S_0, \quad \text{and} \quad R(0) = R_0.$$

We will use the same model parameters as in Chapter 2 (Table 2.1) and assume that there is no resistance cost ($r_S = r_R$).

Such switching may occur either due to genetic mutation, or changes in cell phenotype. For example, Smalley et al [114] found that changes in mRNA expression correlated with resistance to BRAF inhibitors in melanoma and that cells may switch between states upon drug exposure at rates that depend on the cell line or the patient.

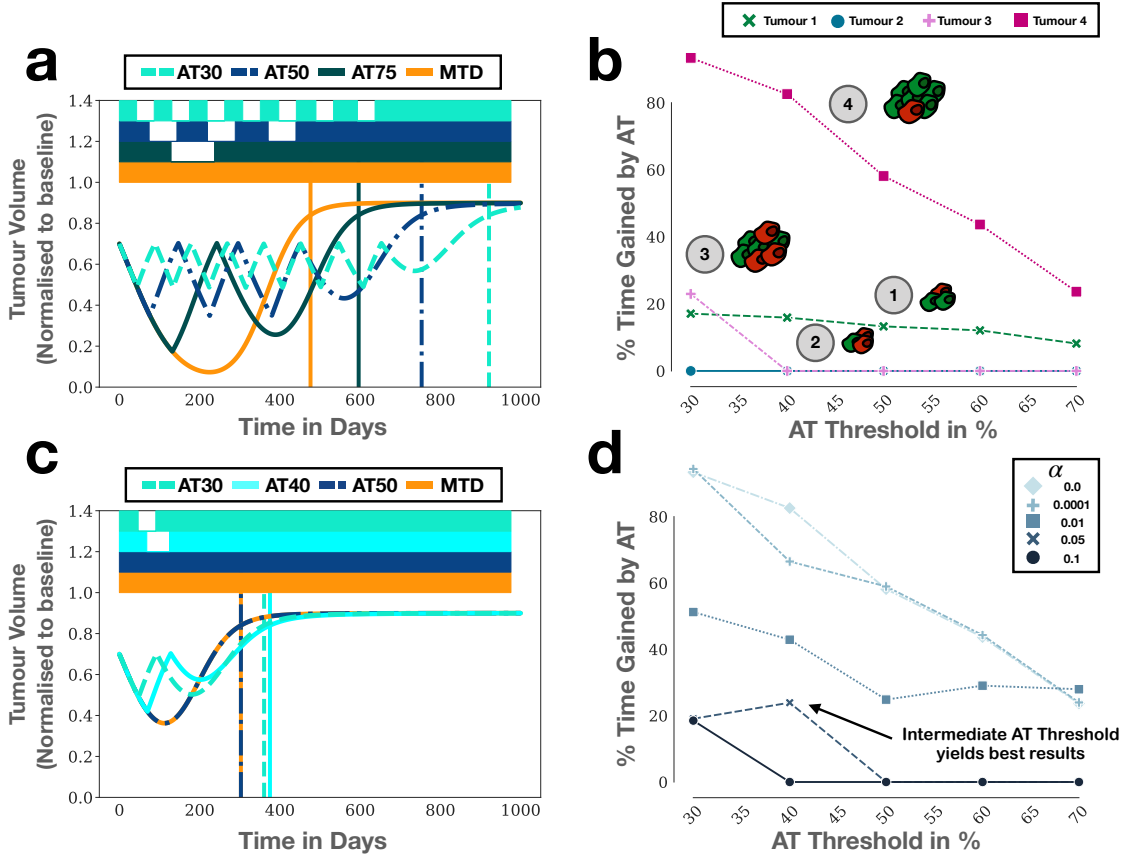


Figure C.5: Optimising the threshold for treatment withdrawal in the AT protocol. (a) Simulations of the total tumour burden, $N(t)$, under three different AT protocols according to the two-population Lotka-Volterra model ($N_0/K = 0.7$; $R_0/N_0 = 10^{-3}$; $r_R = r_S = 0.027\text{d}^{-1}$; $d_R = d_S = 0.0027\text{d}^{-1}$). Vertical lines denote the time point of progression. Horizontal bars illustrate the treatment (filled: treatment on; blank: treatment off). The smaller the threshold, the longer the TTP. (b) Relationship between the gain in TTP by AT compared to CT with the AT threshold for four tumours with different characteristics (Tumour 1: $(N_0/K, R_0/N_0) = (30\%, 0.1\%)$; Tumour 2: $(N_0/K, R_0/N_0) = (30\%, 10\%)$; Tumour 3: $(N_0/K, R_0/N_0) = (30\%, 10\%)$; Tumour 4: $(N_0/K, R_0/N_0) = (70\%, 0.1\%)$; other parameters as in A). This confirms the negative correlation between AT threshold and TTP. (c) Comparison of three AT protocols in the presence of *de novo* resistance formation ($\alpha = 0.05r_R$). Now, an intermediate threshold achieves the highest TTP. (d) Relationship between TTP and AT threshold for different values of α for the tumour in C (values of α are relative to r_S). While the window in which an intermediate AT threshold is preferred is small, α strongly modulates the sensitivity of the gain by AT on the threshold. Overall, our results indicate opportunities for treatment personalisation if R_0/N_0 , N_0/K and α could be known.

Furthermore, they found that they can describe treatment response using a model akin to the one in Equations (C.1)-(C.3).

Let us now compare how the benefit of AT over CT changes as we reduce the threshold at which treatment is withdrawn. In the absence of switching ($\alpha = 0$), we reproduce the results from Chapter 2 that the smaller the threshold, the longer the TTP (Figures C.5a & b). However, when we set $\alpha = 5\%r_S$, this is no longer the case. Now, an intermediate threshold of 40% achieves the longest tumour control (Figures C.5c & d). While the difference is small and the window of α values over which this occurs is narrow, a finding also made by Viossat and Noble [126], this result is conceptually important: It shows that the optimal choice of adaptive schedule changes with the rate at which treatment induces resistance.

References

- [1] Hansen, E., Woods, R. J., & Read, A. F. (2017). How to Use a Chemotherapeutic Agent When Resistance to It Threatens the Patient. *PLoS Biology*, 15(2), e2001110.
- [2] Melander, A. L. (1914). Can Insects Become Resistant to Sprays? *Journal of Economic Entomology*, 7(2), 167–173.
- [3] Palumbi, S. R. (2001). Humans as the world’s greatest evolutionary force. *Science*, 293(5536), 1786–1790.
- [4] De Vita, V. T., & Chu, E. (2008). A History of Cancer Chemotherapy. *American Association for Cancer Research*, 68(21), 8643–8654.
- [5] Mukherjee, S. (2010). *The emperor of all maladies: a biography of cancer*. Simon and Schuster.
- [6] Sporn, M. B. (1996). The war on cancer. *The Lancet*, 347(9012), 1377–81.
- [7] Weinberg, R. (2013). *The Biology of Cancer*. Garland science.
- [8] Niederhuber, J. E., Armitage, J. O., Doroshow, J. H., Kastan, M. B., & Tepper, J. E. (2013). *Abeloff’s Clinical Oncology: Fifth Edition*.
- [9] Crowley, E., Di Nicolantonio, F., Loupakis, F., & Bardelli, A. (2013). Liquid biopsy: Monitoring cancer-genetics in the blood. *Nature Reviews Clinical Oncology*, 10(8), 472–484.
- [10] Khan, K. H., Cunningham, D., Werner, B., Vlachogiannis, G., Spiteri, I., Heide, T., Mateos, J. F., Vatsiou, A., Lampis, A., Damavandi, M. D., Lote, H., Huntingford, I. S., Hedayat, S., Chau, I., Tunariu, N., Mentrasti, G., Trevisani, F., Rao, S., Anandappa, G., Watkins, D., Starling, N., Thomas, J., Peckitt, C., Khan, N., Rugge, M., Begum, R., Hezelova, B., Bryant, A., Jones, T., Proszek, P., Fassan, M., Hahne, J. C., Hubank, M., Braconi, C., Sottoriva, A., & Valeri,

- N. (2018). Longitudinal liquid biopsy and mathematical modeling of clonal evolution forecast time to treatment failure in the prospect-c phase ii colorectal cancer clinical trial. *Cancer Discovery*, 8(10), 1270–1285.
- [11] Liu, M. C., Oxnard, G. R., Klein, E. A., Swanton, C., Seiden, M. V., Liu, M. C., Oxnard, G. R., Klein, E. A., Smith, D., Richards, D., Yeatman, T. J., Cohn, A. L., Lapham, R., Clement, J., Parker, A. S., Tummala, M. K., McIntyre, K., Sekeres, M. A., Bryce, A. H., Siegel, R., Wang, X., Cosgrove, D. P., Abu-Rustum, N. R., Trent, J., Thiel, D. D., Becerra, C., Agrawal, M., Garbo, L. E., Giguere, J. K., Michels, R. M., Harris, R. P., Richey, S. L., McCarthy, T. A., Waterhouse, D. M., Couch, F. J., Wilks, S. T., Krie, A. K., Balaraman, R., Restrepo, A., Meshad, M. W., Rieger-Christ, K., Sullivan, T., Lee, C. M., Greenwald, D. R., Oh, W., Tsao, C. K., Fleshner, N., Kennecke, H. F., Khalil, M. F., Spigel, D. R., Manhas, A. P., Ulrich, B. K., Kover, P. A., Stokoe, C., Courtright, J. G., Yimer, H. A., Larson, T. G., Swanton, C., Seiden, M. V., Cummings, S. R., Absalan, F., Alexander, G., Allen, B., Amini, H., Aravanis, A. M., Bagaria, S., Bazargan, L., Beausang, J. F., Berman, J., Betts, C., Blocker, A., Bredno, J., Calef, R., Cann, G., Carter, J., Chang, C., Chawla, H., Chen, X., Chien, T. C., Civello, D., Davydov, K., Demas, V., Desai, M., Dong, Z., Fayzullina, S., Fields, A. P., Filippova, D., Freese, P., Fung, E. T., Gnerre, S., Gross, S., Halks-Miller, M., Hall, M. P., Hartman, A. R., Hou, C., Hubbell, E., Hunkapiller, N., Jagadeesh, K., Jamshidi, A., Jiang, R., Jung, B., Kim, T. H., Klausner, R. D., Kurtzman, K. N., Lee, M., Lin, W., Lipson, J., Liu, H., Liu, Q., Lopatin, M., Maddala, T., Maher, M. C., Melton, C., Mich, A., Nautilyal, S., Newman, J., Newman, J., Nicula, V., Nicolaou, C., Nikolic, O., Pan, W., Patel, S., Prins, S. A., Rava, R., Ronaghi, N., Sakarya, O., Satya, R. V., Schellenberger, J., Scott, E., Sehnert, A. J., Shaknovich, R., Shanmugam, A., Shashidhar, K. C., Shen, L., Shenoy, A., Shojaee, S., Singh, P., Steffen, K. K., Tang, S., Toung, J. M., Valouev, A., Venn, O., Williams, R. T., Wu, T., Xu, H. H., Yakym, C., Yang, X., Yecies, J., Yip, A. S., Youngren, J., Yue, J., Zhang, J., Zhang, L., Zhang, L. Q., Zhang, N., Curtis, C., & Berry, D. A. (2020). Sensitive and specific multi-cancer detection and localization using methylation signatures in cell-free DNA. *Annals of Oncology*, 31(6), 745–759.
- [12] Zapatero, A., Gómez-Caamaño, A., Cabeza Rodriguez, M. Á., Muinelo-Romay, L., Martin De Vidales, C., Abalo, A., Calvo Crespo, P., Leon Mateos, L., Olivier, C., & Vega Piris, L. V. (2020). Detection and dynamics of circulating

- tumor cells in patients with high-risk prostate cancer treated with radiotherapy and hormones: A prospective phase II study. *Radiation Oncology*, 15(1), 137.
- [13] Lilja, H., Ulmert, D., & Vickers, A. J. (2008). Prostate-specific antigen and prostate cancer: Prediction, detection and monitoring. *Nature Reviews Cancer*, 8(4), 268–278.
- [14] Mathai, R., Vidya, R., Reddy, B., Thomas, L., Udupa, K., Kolesar, J., & Rao, M. (2019). Potential Utility of Liquid Biopsy as a Diagnostic and Prognostic Tool for the Assessment of Solid Tumors: Implications in the Precision Oncology. *Journal of Clinical Medicine*, 8(3), 373.
- [15] Eisenhauer, E. A., Therasse, P., Bogaerts, J., Schwartz, L. H., Sargent, D., Ford, R., Dancey, J., Arbuck, S., Gwyther, S., Mooney, M., Rubinstein, L., Shankar, L., Dodd, L., Kaplan, R., Lacombe, D., & Verweij, J. (2009). New response evaluation criteria in solid tumours: Revised RECIST guideline (version 1.1). *European Journal of Cancer*, 45(2), 228–247.
- [16] Gubbels, J. A., Claussen, N., Kapur, A. K., Connor, J. P., & Patankar, M. S. (2010). The detection, treatment, and biology of epithelial ovarian cancer. *Journal of Ovarian Research*, 3(1), 8.
- [17] Moore, K., Colombo, N., Scambia, G., Kim, B. G., Oaknin, A., Friedlander, M., Lisianskaya, A., Floquet, A., Leary, A., Sonke, G. S., Gourley, C., Banerjee, S., Oza, A., González-Martín, A., Aghajanian, C., Bradley, W., Mathews, C., Liu, J., Lowe, E. S., Bloomfield, R., & DiSilvestro, P. (2018). Maintenance olaparib in patients with newly diagnosed advanced ovarian cancer. *New England Journal of Medicine*, 379(26), 2495–2505.
- [18] Cancer Research UK (2017). Ovarian Cancer Statistics.
URL <http://www.cancerresearchuk.org/health-professional/cancer-statistics/statistics-by-cancer-type/ovarian-cancer>
- [19] Benzekry, S., Pasquier, E., Barbolosi, D., Lacarelle, B., Barlési, F., André, N., & Ciccolini, J. (2015). Metronomic reloaded: Theoretical models bringing chemotherapy into the era of precision medicine. *Seminars in Cancer Biology*, 35, 53–61.

- [20] Araujo, R. P., & McElwain, D. L. S. (2004). A history of the study of solid tumour growth: The contribution of mathematical modelling. *Bulletin of Mathematical Biology*, *66*(5), 1039–1091.
- [21] Altrock, P. M., Liu, L. L., & Michor, F. (2015). The mathematics of cancer: Integrating quantitative models. *Nature Reviews Cancer*, *15*(12), 730–745.
- [22] Benzekry, S. (2020). Artificial Intelligence and Mechanistic Modeling for Clinical Decision Making in Oncology. *Clinical Pharmacology and Therapeutics*, *108*(3), 471–486.
- [23] Bruno, R., Bottino, D., De Alwis, D. P., Fojo, A. T., Guedj, J., Liu, C., Swanson, K. R., Zheng, J., Zheng, Y., & Jin, J. Y. (2020). Progress and opportunities to advance clinical cancer therapeutics using tumor dynamic models. *Clinical Cancer Research*, *26*(8), 1787–1795.
- [24] Mould, D. R., & Upton, R. N. (2013). Basic concepts in population modeling, simulation, and model-based drug development - Part 2: Introduction to pharmacokinetic modeling methods. *CPT: Pharmacometrics and Systems Pharmacology*, *2*(4).
- [25] Lavielle, M. (2014). *Mixed effects models for the population approach: models, tasks, methods and tools*. CRC Press.
- [26] Engelhart, M., Lebedez, D., & Sager, S. (2011). Optimal control for selected cancer chemotherapy ODE models: A view on the potential of optimal schedules and choice of objective function. *Mathematical Biosciences*, *229*(1), 123–134.
- [27] Shi, J., Alagoz, O., Erenay, F. S., & Su, Q. (2014). A survey of optimization models on cancer chemotherapy treatment planning. *Annals of Operations Research*, *221*(1), 331–356.
- [28] Carrère, C. (2017). Optimization of an in vitro chemotherapy to avoid resistant tumours. *Journal of Theoretical Biology*, *413*, 24–33.
- [29] Sharp, J. A., Browning, A. P., Mapder, T., Burrage, K., & Simpson, M. J. (2019). Optimal control of acute myeloid leukaemia. *Journal of Theoretical Biology*, *470*, 30–42.
- [30] Skipper, H. E., Schabel, F. M., & Wilcox, W. S. (1964). Experimental Evaluation of Potential Anticancer Agents. *Cancer chemotherapy reports.*, *35*, 1–111.

- [31] Skipper, H. E. (1965). The Effects of Chemotherapy on the Kinetics of Leukemic Cell Behavior. *Cancer Research*, 25, 1544–1550.
- [32] Freireich, E. J., Karon, M., & Frei III, E. (1964). Quadruple combination therapy (VAMP) for acute lymphocytic leukemia of childhood. In *Proc. Am. Assoc. Cancer Res.*, vol. 5, (p. 20).
- [33] Pinkel, D. P., Simone, J. V., Hustu, H. O., & J, R. (1971). “Total therapy” of childhood acute lymphocytic leukemia. *Pediatric Research*, 5(8), 408–408.
- [34] Steel, G. G., Adams, K., & Barrett, J. C. (1966). Analysis of the cell population kinetics of transplanted tumours of widely-differing growth rate. *British Journal of Cancer*, 20(4), 784–800.
- [35] Steel, G. G. (1968). Cell loss from experimental tumours. *Cell Proliferation*, 1(3), 193–207.
- [36] Malaise, E. P., Chavaudra, N., & Tubiana, M. (1973). The relationship between growth rate, labelling index and histological type of human solid tumours. *European Journal of Cancer (1965)*, 9(4), 305–312.
- [37] Gompertz (1825). On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. *Philosophical Transactions of the Royal Society of London*, 115, 513–583.
- [38] Von Bertalanffy, L. (1957). Quantitative Laws in Metabolism and Growth. *Source: The Quarterly Review of Biology*, 32(3), 217–231.
- [39] Laird, A. K. (1964). Dynamics of tumor growth. *British Journal of Cancer*, 18(3), 490–502.
- [40] Norton, L., & Simon, R. (1977). Growth curve of an experimental solid tumor following radiotherapy. *Journal of the National Cancer Institute*, 58(6), 1735–1741.
- [41] Benzekry, S., Lamont, C., Beheshti, A., Tracz, A., Ebos, J. M., Hlatky, L., & Hahnfeldt, P. (2014). Classical Mathematical Models for Description and Prediction of Experimental Tumor Growth. *PLoS Computational Biology*, 10(8), e1003800:1–19.

- [42] Norton, L., & Simon, R. (1977). Tumor size, sensitivity to therapy, and design of treatment schedules. *Cancer treatment reports*, 61(7), 1307–1317.
- [43] Simon, R., & Norton, L. (2006). The Norton - Simon hypothesis: Designing more effective and less toxic chemotherapeutic regimens. *Nature Clinical Practice Oncology*, 3(8), 406–407.
- [44] Citron, M. L., Berry, D. A., Cirincione, 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., & 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. *Journal of Clinical Oncology*, 21(8), 1431–1439.
- [45] Gaffney, E. A. (2004). The application of mathematical modelling to aspects of adjuvant chemotherapy scheduling. *Journal of Mathematical Biology*, 48(4), 375–422.
- [46] Holohan, C., Van Schaeybroeck, S., Longley, D. B., & Johnston, P. G. (2013). Cancer drug resistance: An evolving paradigm. *Nature Reviews Cancer*, 13(10), 714–726.
- [47] Burrell, R. A., & Swanton, C. (2014). Tumour heterogeneity and the evolution of polyclonal drug resistance. *Molecular Oncology*, 8(6), 1095–1111.
- [48] Goldie, J. H., & Coldman, A. J. (1979). A mathematic model for relating the drug sensitivity of tumors to their spontaneous mutation rate. *Cancer Treatment Reports*, 63(11-12), 1727–1733.
- [49] Goldie, J. H., Coldman, A. J., & Gudauskas, G. A. (1982). Rationale for the use of alternating non-cross-resistant chemotherapy. *Cancer Treatment Reports*, 66(3), 439–449.
- [50] Goldie, J. H., & Coldman, A. J. (1986). A Stochastic Model for the Origin and Treatment of Tumors Containing Drug Resistant Cells. *Bulletin of Mathematical Biology*, 48(4), 279–292.

- [51] Coldman, A. J., & Goldie, J. H. (1983). A model for the resistance of tumor cells to cancer chemotherapeutic agents. *Mathematical Biosciences*, *65*(2), 291–307.
- [52] Goldie, J. H., & Coldman, A. J. (1984). The Genetic Origin of Drug Resistance in Neoplasms: Implications for Systemic Therapy. *Cancer Research*, *44*(9), 3643–3653.
- [53] Bonadonna, G., Zambetti, M., & Valagussa, P. (1995). Sequential or Alternating Doxorubicin and CMF Regimens in Breast Cancer With More Than Three Positive Nodes: Ten-Year Results. *JAMA: The Journal of the American Medical Association*, *273*(7), 542–547.
- [54] Day, R. S. (1986). Treatment sequencing, asymmetry, and uncertainty: Protocol strategies for combination chemotherapy. *Cancer Research*, *46*(8), 3876–3885.
- [55] Murray, J. M., & Coldman, A. J. (2003). The effect of heterogeneity on optimal regimens in cancer chemotherapy. *Mathematical Biosciences*, *185*(1), 73–87.
- [56] Chmielecki, J., Foo, J., Oxnard, G. R., Hutchinson, K., Ohashi, K., Somwar, R., Wang, L., Amato, K. R., Arcila, M., Sos, M. L., Socci, N. D., Viale, A., De Stanchina, E., Ginsberg, M. S., Thomas, R. K., Kris, M. G., Inoue, A., Ladanyi, M., Miller, V. A., Michor, F., & Pao, W. (2011). Optimization of dosing for EGFR-mutant non-small cell lung cancer with evolutionary cancer modeling. *Science Translational Medicine*, *3*(90), 90ra5990ra59.
- [57] Yu, H. A., Sima, C., Feldman, D., Liu, L. L., Vaitheesvaran, B., Cross, J., Rudin, C. M., Kris, M. G., Pao, W., Michor, F., & Riely, G. J. (2017). Phase 1 study of twice weekly pulse dose and daily low-dose erlotinib as initial treatment for patients with EGFR-mutant lung cancers. *Annals of Oncology*, *28*(2), 278–284.
- [58] Martin, R. B., Fisher, M. E., Minchin, R. F., & Teo, K. L. (1992). Optimal control of tumor size used to maximize survival time when cells are resistant to chemotherapy. *Mathematical Biosciences*, *110*(2), 201–219.
- [59] Monro, H. C., & Gaffney, E. A. (2009). Modelling chemotherapy resistance in palliation and failed cure. *Journal of Theoretical Biology*, *257*(2), 292–302.
- [60] Gatenby, R. A., Silva, A. S., Gillies, R. J., & Frieden, B. R. (2009). Adaptive therapy. *Cancer Research*, *69*(11), 4894–4903.

- [61] Vogelstein, B., & Kinzler, K. W. (2004). Cancer genes and the pathways they control. *Nature Medicine*, *10*(8), 789–799.
- [62] Xie, D., & Xie, K. (2015). Pancreatic cancer stromal biology and therapy. *Genes and Diseases*, *2*(2), 133–143.
- [63] Browder, T., Butterfield, C. E., Kräling, B. M., Shi, B., Marshall, B., O’Reilly, M. S., & Folkman, J. (2000). Antiangiogenic scheduling of chemotherapy improves efficacy against experimental drug-resistant cancer. *Cancer Research*, *60*(7), 1878–1886.
- [64] Klement, G., Baruchel, S., Rak, J., Man, S., Clark, K., Hicklin, D. J., Bohlen, P., & Kerbel, R. S. (2000). Continuous low-dose therapy with vinblastine and VEGF receptor-2 antibody induces sustained tumor regression without overt toxicity. *Journal of Clinical Investigation*, *105*(8), R15–R24.
- [65] Hanahan, D., Bergers, G., & Bergsland, E. (2000). Less is, more, regularly: Metronomic dosing of cytotoxic drugs can target tumor angiogenesis in mice. *Journal of Clinical Investigation*, *105*(8), 1045–1047.
- [66] Rajasekaran, T., Ng, Q. S., Tan, D. S. W., Lim, W. T., Ang, M. K., Toh, C. K., Chowbay, B., Kanesvaran, R., & Tan, E. H. (2017). Metronomic chemotherapy: A relook at its basis and rationale. *Cancer Letters*, *388*, 328–333.
- [67] André, N., Carré, M., & Pasquier, E. (2014). Metronomics: Towards personalized chemotherapy? *Nature Reviews Clinical Oncology*, *11*(7), 413–431.
- [68] Hahnfeldt, P., Panigrahy, D., Folkman, J., & Hlatky, L. (1999). Tumor development under angiogenic signaling: A dynamical theory of tumor growth, treatment response, and postvascular dormancy. *Cancer Research*, *59*(19), 4770–4775.
- [69] Hahnfeldt, P., Folkman, J., & Hlatky, L. (2003). Minimizing long-term tumor burden: The logic for metronomic chemotherapeutic dosing and its antiangiogenic basis. *Journal of Theoretical Biology*, *220*(4), 545–554.
- [70] D’Onofrio, A., Ledzewicz, U., Maurer, H., & Schättler, H. (2009). On optimal delivery of combination therapy for tumors. *Mathematical Biosciences*, *222*(1), 13–26.

- [71] Ledzewicz, U., & Schättler, H. (2007). Antiangiogenic therapy in cancer treatment as an optimal control problem. *SIAM Journal on Control and Optimization*, *46*(3), 1052–1079.
- [72] Barbolosi, D., Ciccolini, J., Meille, C., Elharrar, X., Faivre, C., Lacarelle, B., André, N., & Barlesi, F. (2014). Metronomics chemotherapy: Time for computational decision support. *Cancer Chemotherapy and Pharmacology*, *74*(3), 647–652.
- [73] Barlesi, F., Imbs, D. C., Tomasini, P., Greillier, L., Galloux, M., Testot-Ferry, A., Garcia, M., Elharrar, X., Pelletier, A., André, N., Mascaux, C., Lacarelle, B., El Cheikh, R. E., Serre, R., Ciccolini, J., & Barbolosi, D. (2017). Mathematical modeling for Phase I cancer trials: A study of metronomic vinorelbine for advanced non-small cell lung cancer (NSCLC) and mesothelioma patients. *Oncotarget*, *8*(29), 47161–47166.
- [74] Park, D. S., Luddy, K. A., Robertson-Tessi, M., O’Farrelly, C., Gatenby, R. A., & Anderson, A. R. (2020). Searching for Goldilocks: How Evolution and Ecology Can Help Uncover More Effective Patient-Specific Chemotherapies. *Cancer Research*, *80*(23), 5147–5154.
- [75] De Pillis, L. G., Gu, W., & Radunskaya, A. E. (2006). Mixed immunotherapy and chemotherapy of tumors: Modeling, applications and biological interpretations. *Journal of Theoretical Biology*, *238*(4), 841–862.
- [76] Ledzewicz, U., Amini, B., & Schättler, H. (2015). Dynamics and control of a mathematical model for metronomic chemotherapy. *Mathematical Biosciences and Engineering*, *12*(6), 1257–1275.
- [77] Park, D. S., Robertson-Tessi, M., Luddy, K. A., Maini, P. K., Bonsall, M. B., Gatenby, R. A., & Anderson, A. R. (2019). The goldilocks window of personalized chemotherapy: Getting the immune response just right. *Cancer Research*, *79*(20), 5302–5315.
- [78] Benzekry, S., & Hahnfeldt, P. (2013). Maximum tolerated dose versus metronomic scheduling in the treatment of metastatic cancers. *Journal of Theoretical Biology*, *335*, 235–244.

- [79] Kwon, Y. (2001). *Handbook of essential pharmacokinetics, pharmacodynamics and drug metabolism for industrial scientists*. Springer Science & Business Media.
- [80] Beumer, J. H. (2013). Without therapeutic drug monitoring, there is no personalized cancer care. *Clinical Pharmacology and Therapeutics*, *93*(3), 228–230.
- [81] Morrissey, K. M., Marchand, M., Patel, H., Zhang, R., Wu, B., Phyllis Chan, H., Mecke, A., Girish, S., Jin, J. Y., Winter, H. R., & Bruno, R. (2019). Alternative dosing regimens for atezolizumab: an example of model-informed drug development in the postmarketing setting. *Cancer Chemotherapy and Pharmacology*, *84*(6), 1257–1267.
- [82] McDougall, S. R., Anderson, A. R., Chaplain, M. A., & Sherratt, J. A. (2002). Mathematical modelling of flow through vascular networks: Implications for tumour-induced angiogenesis and chemotherapy strategies. *Bulletin of Mathematical Biology*, *64*(4), 673–702.
- [83] Connor, A. J., Nowak, R. P., Lorenzon, E., Thomas, M., Herting, F., Hoert, S., Quaiser, T., Shochat, E., Pitt-Francis, J., Cooper, J., Maini, P. K., & Byrne, H. M. (2015). An integrated approach to quantitative modelling in angiogenesis research. *Journal of the Royal Society Interface*, *12*(110), 20150546.
- [84] Bernabeu, M. O., Köry, J., Grogan, J. A., Markelc, B., Beardo, A., D’Avezac, M., Enjalbert, R., Kaeppler, J., Daly, N., Hetherington, J., Krüger, T., Maini, P. K., Pitt-Francis, J. M., Muschel, R. J., Alarcón, T., & Byrne, H. M. (2020). Abnormal morphology biases hematocrit distribution in tumor vasculature and contributes to heterogeneity in tissue oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(45), 27811–27819.
- [85] Friberg, L. E., Henningsson, A., Maas, H., Nguyen, L., & Karlsson, M. O. (2002). Model of chemotherapy-induced myelosuppression with parameter consistency across drugs. *Journal of Clinical Oncology*, *20*(24), 4713–4721.
- [86] Ferrer, F., Fanciullino, R., Milano, G., & Ciccolini, J. (2020). Towards Rational Cancer Therapeutics: Optimizing Dosing, Delivery, Scheduling, and Combinations. *Clinical Pharmacology and Therapeutics*, *108*(3), 458–470.

- [87] Meille, C., Barbolosi, D., Ciccolini, J., Freyer, G., & Iliadis, A. (2016). Revisiting Dosing Regimen Using Pharmacokinetic/Pharmacodynamic Mathematical Modeling: Densification and Intensification of Combination Cancer Therapy. *Clinical Pharmacokinetics*, *55*(8), 1015–1025.
- [88] Hénin, E., Meille, C., Barbolosi, D., You, B., Guitton, J., Iliadis, A., & Freyer, G. (2016). Revisiting dosing regimen using PK/PD modeling: The MODEL1 phase I/II trial of docetaxel plus epirubicin in metastatic breast cancer patients. *Breast Cancer Research and Treatment*, *156*(2), 331–341.
- [89] Sheiner, L. B., Rosenberg, B., & Melmon, K. L. (1972). Modelling of individual pharmacokinetics for computer-aided drug dosage. *Computers and Biomedical Research*, *5*(5), 441–459.
- [90] Evans, W. E., Relling, M. V., Rodman, J. H., Crom, W. R., Boyett, J. M., & Pui, C.-H. (1998). Conventional Compared with Individualized Chemotherapy for Childhood Acute Lymphoblastic Leukemia. *New England Journal of Medicine*, *338*(8), 499–505.
- [91] Panetta, J. C., Iacono, L. C., Adamson, P. C., & Stewart, C. F. (2003). The Importance of Pharmacokinetic Limited Sampling Models for Childhood Cancer Drug Development. *Clinical Cancer Research*, *9*(14), 5068–5077.
- [92] Barrett, J. S., Mondick, J. T., Narayan, M., Vijayakumar, K., & Vijayakumar, S. (2008). Integration of modeling and simulation into hospital-based decision support systems guiding pediatric pharmacotherapy. *BMC Medical Informatics and Decision Making*, *8*(1), 6:1–12.
- [93] Noble, S. L., Sherer, E., Hannemann, R. E., Ramkrishna, D., Vik, T., & Rundell, A. E. (2010). Using adaptive model predictive control to customize maintenance therapy chemotherapeutic dosing for childhood acute lymphoblastic leukemia. *Journal of Theoretical Biology*, *264*(3), 990–1002.
- [94] Jayachandran, D., Rundell, A. E., Hannemann, R. E., Vik, T. A., & Ramkrishna, D. (2014). Optimal chemotherapy for Leukemia: A model-based strategy for individualized treatment. *PLoS ONE*, *9*(10), 109623.
- [95] Barbolosi, D., Ciccolini, J., Lacarelle, B., Barlési, F., & André, N. (2016). Computational oncology-mathematical modelling of drug regimens for precision medicine. *Nature Reviews Clinical Oncology*, *13*(4), 242–254.

- [96] Ratain, M. J., Schilsky, R. L., Choi, K. E., Guarnieri, C., Grimmer, D., Vogelzang, N. J., Senekjian, E., & Liebner, M. A. (1989). Adaptive control of etoposide administration: Impact of interpatient pharmacodynamic variability. *Clinical Pharmacology and Therapeutics*, *45*(3), 226–233.
- [97] Zhang, J., Cunningham, J. J., Brown, J. S., & Gatenby, R. A. (2017). Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature Communications*, *8*(1), 1816.
- [98] Gatenby, R. A. (2009). A change of strategy in the war on cancer. *Nature*, *459*(7246), 508–509.
- [99] Merlo, L. M. F., Pepper, J. W., Reid, B. J., & Maley, C. C. (2006). Cancer as an evolutionary and ecological process. *Nature Reviews Cancer*, *6*(12), 924–935.
- [100] Hanahan, D., & Weinberg, R. A. (2011). Hallmarks of cancer: The next generation. *Cell*, *144*(5), 646–674.
- [101] Maley, C. C., Aktipis, A., Graham, T. A., Sottoriva, A., Boddy, A. M., Janiszewska, M., Silva, A. S., Gerlinger, M., Yuan, Y., Pienta, K. J., Anderson, K. S., Gatenby, R., Swanton, C., Posada, D., Wu, C. I., Schiffman, J. D., Hwang, E. S., Polyak, K., Anderson, A. R., Brown, J. S., Greaves, M., & Shibata, D. (2017). Classifying the evolutionary and ecological features of neoplasms. *Nature Reviews Cancer*, *17*(10), 605–619.
- [102] Korolev, K. S., Xavier, J. B., & Gore, J. (2014). Turning ecology and evolution against cancer. *Nature Reviews Cancer*, *14*(5), 371–380.
- [103] Amend, S. R., & Pienta, K. J. (2015). Ecology meets cancer biology: The cancer swamp promotes the lethal cancer phenotype. *Oncotarget*, *6*(12), 9669–9678.
- [104] Gatenby, R., & Brown, J. (2018). The evolution and ecology of resistance in cancer therapy. *Cold Spring Harbor Perspectives in Medicine*, *8*(3).
- [105] Frost, I., Smith, W. P., Mitri, S., Millan, A. S., Davit, Y., Osborne, J. M., Pitt-Francis, J. M., MacLean, R. C., & Foster, K. R. (2018). Cooperation, competition and antibiotic resistance in bacterial colonies. *ISME Journal*, *12*(6), 1582–1593.
- [106] Whelan, C. J., & Cunningham, J. J. (2020). Resistance is not the end: lessons from pest management. *Cancer Control*, *27*(1).

- [107] Hansen, E., Karlake, J., Woods, R., Read, A., & Wood, K. (2020). Antibiotics can be used to contain drug-resistant bacteria by maintaining sufficiently large sensitive populations. *PLoS Biology*, *18*(5), e3000713.
- [108] Tabashnik, B. E., Carrière, Y., Dennehy, T. J., Morin, S., Sisterson, M. S., Roush, R. T., Shelton, A. M., & Zhao, J. Z. (2003). Insect resistance to transgenic Bt crops: Lessons from the laboratory and field. *Journal of Economic Entomology*, *96*(4), 1031–1038.
- [109] Bruchofsky, N., Klotz, L., Crook, J., Malone, S., Ludgate, C., Morris, W. J., Gleave, M. E., & Goldenberg, S. L. (2006). Final results of the Canadian prospective Phase II trial of intermittent androgen suppression for men in biochemical recurrence after radiotherapy for locally advanced prostate cancer: Clinical parameters. *Cancer*, *107*(2), 389–395.
- [110] Crook, J. M., O’Callaghan, C. J., Duncan, G., Dearnaley, D. P., Higano, C. S., Horwitz, E. M., Frymire, E., Malone, S., Chin, J., Nabid, A., Warde, P., Corbett, T., Angyalfi, S., Goldenberg, S. L., Gospodarowicz, M. K., Saad, F., Logue, J. P., Hall, E., Schellhammer, P. F., Ding, K., & Klotz, L. (2012). Intermittent androgen suppression for rising PSA level after radiotherapy. *New England Journal of Medicine*, *367*(10), 895–903.
- [111] Hussain, M., Tangen, C. M., Berry, D. L., Higano, C. S., Crawford, E. D., Liu, G., Wilding, G., Prescott, S., Sundaram, S. K., Small, E. J., Dawson, N. A., Donnelly, B. J., Venner, P. M., Vaishampayan, U. N., Schellhammer, P. F., Quinn, D. I., Raghavan, D., Ely, B., Moinpour, C. M., Vogelzang, N. J., Thompson, I. M., Kanaga Sundaram, S., Small, E. J., Dawson, N. A., Donnelly, B. J., Venner, P. M., Vaishampayan, U. N., Schellhammer, P. F., Quinn, D. I., Raghavan, D., Ely, B., Moinpour, C. M., Vogelzang, N. J., & Thompson, I. M. (2013). Intermittent versus continuous androgen deprivation in prostate cancer. *New England Journal of Medicine*, *368*(14), 1314–1325.
- [112] Algazi, A. P., Othus, M., Daud, A. I., Lo, R. S., Mehnert, J. M., Truong, T. G., Conry, R., Kendra, K., Doolittle, G. C., Clark, J. I., Messino, M. J., Moore, D. F., Lao, C., Faller, B. A., Govindarajan, R., Harker-Murray, A., Dreisbach, L., Moon, J., Grossmann, K. F., & Ribas, A. (2020). Continuous versus intermittent BRAF and MEK inhibition in patients with BRAF-mutated melanoma: a randomized phase 2 trial. *Nature Medicine*, *26*(10), 1564–1568.

- [113] Enriquez-Navas, P. M., Yoonseok, K., Das, T., Hassan, S., Silva, A. S., Foroutan, P., Ruiz, E., Martinez, G., Minton, S., Gillies, R. J., & Gatenby, R. A. (2016). Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer. *Science Translational Medicine*, 8(327), 327ra24.
- [114] Smalley, I., Kim, E., Li, J., Spence, P., Wyatt, C. J., Eroglu, Z., Sondak, V. K., Messina, J. L., Babacan, N. A., Maria-Engler, S. S., De Armas, L., Williams, S. L., Gatenby, R. A., Chen, Y. A., Anderson, A. R., & Smalley, K. S. (2019). Leveraging transcriptional dynamics to improve BRAF inhibitor responses in melanoma. *EBioMedicine*, 48, 178–190.
- [115] Bacevic, K., Noble, R., Soffar, A., Ammar, O. W., Boszonyik, B., Prieto, S., Vincent, C., Hochberg, M. E., Krasinska, L., & Fisher, D. (2017). Spatial competition constrains resistance to targeted cancer therapy. *Nature Communications*, 8(1), 1–15.
- [116] Farrokhian, N., Maltas, J., Ellsworth, P., Durmaz, A., Dinh, M., Hitomi, M., Kaznatcheev, A., Marusyk, A., & Scott, J. G. (2020). Dose dependent evolutionary game dynamics modulate competitive release in cancer therapy. *bioRxiv*, 2020.09.18(303966), doi: <https://doi.org/10.1101/2020.09.18.303966>.
- [117] Silva, A. S., Kam, Y., Khin, Z. P., Minton, S. E., Gillies, R. J., & Gatenby, R. A. (2012). Evolutionary approaches to prolong progression-free survival in breast cancer. *Cancer Research*, 72(24), 6362–6370.
- [118] Cunningham, J. J., Brown, J. S., Gatenby, R. A., & Staňková, K. (2018). Optimal control to develop therapeutic strategies for metastatic castrate resistant prostate cancer. *Journal of Theoretical Biology*, 459, 67–78.
- [119] Cunningham, J., Thuijsman, F., Peeters, R., Viossat, Y., Brown, J., Gatenby, R., & Staňková, K. (2020). Optimal control to reach eco-evolutionary stability in metastatic castrate-resistant prostate cancer. *PLOS ONE*, 15(12), e0243386.
- [120] West, J. B., Dinh, M. N., Brown, J. S., Zhang, J., Anderson, A. R., & Gatenby, R. A. (2019). Multidrug cancer therapy in metastatic castrate-resistant prostate cancer: An evolution-based strategy. *Clinical Cancer Research*, 25(14), 4413–4421.

- [121] West, J., Hasnain, Z., Mason, J., & Newton, P. K. (2016). The prisoner’s dilemma as a cancer model. *Convergent Science Physical Oncology*, 2(3), 035002.
- [122] West, J., Ma, Y., & Newton, P. K. (2018). Capitalizing on competition: An evolutionary model of competitive release in metastatic castration resistant prostate cancer treatment. *Journal of Theoretical Biology*, 455, 249–260.
- [123] Brady-Nicholls, R., Nagy, J. D., Gerke, T. A., Zhang, T., Wang, A. Z., Zhang, J., Gatenby, R. A., & Enderling, H. (2020). Prostate-specific antigen dynamics predict individual responses to intermittent androgen deprivation. *Nature Communications*, 11(1), 1–13.
- [124] Gallaher, J. A., Enriquez-Navas, P. M., Luddy, K. A., Gatenby, R. A., & Anderson, A. R. (2018). Spatial heterogeneity and evolutionary dynamics modulate time to recurrence in continuous and adaptive cancer therapies. *Cancer Research*, 78(8), 2127–2139.
- [125] Strobl, M. A., West, J., Viossat, Y., Damaghi, M., Robertson-Tessi, M., Brown, J. S., Gatenby, R. A., Maini, P. K., & Anderson, A. R. (2021). Turnover modulates the need for a cost of resistance in adaptive therapy. *Cancer Research*, 81(4), 1135–1147.
- [126] Viossat, Y., & Noble, R. (2020). The logic of containing tumors. *bioRxiv*, 2020.01.22(915355), doi: <https://doi.org/10.1101/2020.01.22.915355>.
- [127] Viossat, Y., & Noble, R. (2021). A theoretical analysis of tumour containment. *Nature Ecology & Evolution*, (pp. 1–10).
- [128] Hansen, E., & Read, A. F. (2020). Cancer therapy: Attempt cure or manage drug resistance? *Evolutionary Applications*, 13(7), 1660–1672.
- [129] Hansen, E., & Read, A. F. (2020). Modifying adaptive therapy to enhance competitive suppression. *Cancers*, 12(12), 1–13.
- [130] Newton, P. K., & Ma, Y. (2019). Nonlinear adaptive control of competitive release and chemotherapeutic resistance. *Physical Review E*, 99(2), 022404.
- [131] Gluzman, M., Scott, J. G., & Vladimirovsky, A. (2020). Optimizing adaptive cancer therapy: Dynamic programming and evolutionary game theory. *Proceedings of the Royal Society B: Biological Sciences*, 287(1925), 20192454.

- [132] West, J., You, L., Zhang, J., Gatenby, R. A., Brown, J. S., Newton, P. K., & Anderson, A. R. (2020). Towards multidrug adaptive therapy. *Cancer Research*, *80*(7), 1578–1589.
- [133] Yoon, N., Krishnan, N., & Scott, J. (2020). Modeling collaterally sensitive drug cycles: shaping heterogeneity to allow adaptive therapy. *bioRxiv*, *2020.07.02*(184952), doi: <https://doi.org/10.1101/2020.07.02.184952>.
- [134] Ma, Y., & Newton, P. K. (2021). Role of synergy and antagonism in designing multidrug adaptive chemotherapy schedules. *Physical Review E*, *103*(3), 032408.
- [135] Reed, D. R., Metts, J., Pressley, M., Fridley, B. L., Hayashi, M., Isakoff, M. S., Loeb, D. M., Makanji, R., Roberts, R. D., Trucco, M., Wagner, L. M., Alexandrow, M. G., Gatenby, R. A., & Brown, J. S. (2020). An evolutionary framework for treating pediatric sarcomas. *Cancer*, *126*(11), 2577–2587.
- [136] Lenormand, T., Harmand, N., & Gallet, R. (2018). Cost of resistance: an unreasonably expensive concept. *Rethinking Ecology*, *3*, 51–70.
- [137] Broxterman, H. J., Pinedo, H. M., Kuiper, C. M., Kaptein, L. C., Schuurhuis, G. J., & Lankelma, J. (1988). Induction by verapamil of a rapid increase in ATP consumption in multidrug-resistant tumor cells. *The FASEB Journal*, *2*(7), 2278–2282.
- [138] Dobzhansky, T. (1951). *Genetics and the origin of species*. New York: Columbia University Press.
- [139] Andersson, D. I., & Levin, B. R. (1999). The biological cost of antibiotic resistance. *Current Opinion in Microbiology*, *2*(5), 489–493.
- [140] Andersson, D. I., & Hughes, D. (2010). Antibiotic resistance and its cost: Is it possible to reverse resistance? *Nature Reviews Microbiology*, *8*(4), 260–271.
- [141] Jensen, N. F., Stenvang, J., Beck, M. K., Hanáková, B., Belling, K. C., Do, K. N., Viuff, B., Nygård, S. B., Gupta, R., Rasmussen, M. H., Tarpgaard, L. S., Hansen, T. P., Budinská, E., Pfeiffer, P., Bosman, F., Tejpar, S., Roth, A., Delorenzi, M., Andersen, C. L., Rømer, M. U., Brünner, N., & Moreira, J. M. (2015). Establishment and characterization of models of chemotherapy resistance in colorectal cancer: Towards a predictive signature of chemoresistance. *Molecular Oncology*, *9*(6), 1169–1185.

- [142] Kanzawa, F., Sugimoto, Y., Minato, K., Kasahara, K., Bungo, M., Nagakawa, K., Fujiwara, Y., Liu, L. F., & Saijo, N. (1990). Establishment of a camptothecin analogue (CPT-11)-resistant cell line of human non-small cell lung cancer: Characterization and mechanism of resistance. *Cancer Research*, *50*(18), 5919–5924.
- [143] Benavente, S., Huang, S., Armstrong, E. A., Chi, A., Hsu, K. T., Wheeler, D. L., & Harari, P. M. (2009). Establishment and characterization of a model of acquired resistance to epidermal growth factor receptor targeting agents in human cancer cells. *Clinical Cancer Research*, *15*(5), 1585–1592.
- [144] Kaznatcheev, A., Peacock, J., Basanta, D., Marusyk, A., & Scott, J. G. (2019). Fibroblasts and alectinib switch the evolutionary games played by non-small cell lung cancer. *Nature Ecology and Evolution*, *3*(3), 450–456.
- [145] Behrens, B. C., Hamilton, T. C., Masuda, H., Grotzinger, K. R., Whang-Peng, J., Louie, K. G., Knutsen, T., McKoy, W. M., Young, R. C., & Ozols, R. F. (1987). Characterization of a cis-diamminedichloroplatinum(II)-resistant human ovarian cancer cell line and its use in evaluation of platinum analogues. *Cancer research*, *47*(2), 414–418.
- [146] Greene, J. M., Gevertz, J. L., & Sontag, E. D. (2019). Mathematical Approach to Differentiate Spontaneous and Induced Evolution to Drug Resistance During Cancer Treatment. *JCO clinical cancer informatics*, *3*(3), 1–20.
- [147] Grassberger, C., McClatchy, D., Geng, C., Kamran, S. C., Fintelmann, F., Maruvka, Y. E., Piotrowska, Z., Willers, H., Sequist, L. V., Hata, A. N., & Paganetti, H. (2019). Patient-specific tumor growth trajectories determine persistent and resistant cancer cell populations during treatment with targeted therapies. *Cancer Research*, *79*(14), 3776–3788.
- [148] Marušić, M., Bajzer, Freyer, J. P., & Vuk-Pavlović, S. (1994). Analysis of growth of multicellular tumour spheroids by mathematical models. *Cell Proliferation*, *27*(2), 73–94.
- [149] Pisco, A. O., Brock, A., Zhou, J., Moor, A., Mojtahedi, M., Jackson, D., & Huang, S. (2013). Non-Darwinian dynamics in therapy-induced cancer drug resistance. *Nature Communications*, *4*(1), 1–11.

- [150] Vander Velde, R., Yoon, N., Marusyk, V., Durmaz, A., Dhawan, A., Miroshnychenko, D., Lozano-Peral, D., Desai, B., Balynska, O., Poleszhuk, J., Kenian, L., Teng, M., Abazeed, M., Mian, O., Tan, A. C., Haura, E., Scott, J., & Marusyk, A. (2020). Resistance to targeted therapies as a multifactorial, gradual adaptation to inhibitor specific selective pressures. *Nature Communications*, *11*(1), 504837.
- [151] Shen, S., & Clairambault, J. (2020). Open Peer Review Cell plasticity in cancer cell populations [version 1; peer review: 2 approved].
- [152] Therasse, P., Arbuck, S. G., Eisenhauer, E. A., Wanders, J., Kaplan, R. S., Rubinstein, L., Verweij, J., Van Glabbeke, M., Van Oosterom, A. T., Christian, M. C., & Gwyther, S. G. (2000). New guidelines to evaluate the response to treatment in solid tumors. *Journal of the National Cancer Institute*, *92*(3), 205–216.
- [153] Prokopiou, S., Moros, E. G., Poleszczuk, J., Caudell, J., Torres-Roca, J. F., Latifi, K., Lee, J. K., Myerson, R., Harrison, L. B., & Enderling, H. (2015). A proliferation saturation index to predict radiation response and personalize radiotherapy fractionation. *Radiation Oncology*, *10*(159).
- [154] Dormand, J. R., & Prince, P. J. (1980). A family of embedded Runge-Kutta formulae. *Journal of Computational and Applied Mathematics*, *6*(1), 19–26.
- [155] Wodarz, D., & Komarova, N. (2007). Can loss of apoptosis protect against cancer? *Trends in Genetics*, *23*(5), 232–237.
- [156] Steel, G. G. (1967). Cell loss as a factor in the growth rate of human tumours. *European Journal of Cancer (1965)*, *3*(4-5), 381–387.
- [157] Rew, D. A., & Wilson, G. D. (2000). Cell production rates in human tissues and tumours and their significance. Part I: An introduction to the techniques of measurement and their limitations. *European Journal of Surgical Oncology*, *26*(3), 227–238.
- [158] Rew, D. A., & Wilson, G. D. (2000). Cell production rates in human tissues and tumours and their significance. Part II: clinical data. *European Journal of Surgical Oncology*, *26*(4), 405–417.

- [159] Siegel, R. L., Miller, K. D., & Jemal, A. (2020). Cancer statistics, 2020. *CA: A Cancer Journal for Clinicians*, 70(1), 7–30.
- [160] Kratiras, Z., Konstantinidis, C., & Skriapas, K. (2014). A review of continuous vs intermittent androgen deprivation therapy: Redefining the gold standard in the treatment of advanced prostate cancer. Myths, facts and new data on a "perpetual dispute". *International Braz J Urol*, 40(1), 3–15.
- [161] Bruchovsky, N., Rennie, P. S., Goldenberg, S. L., Lawson, D., To, M., & Coldman, A. J. (1990). Effects of Androgen Withdrawal on the Stem Cell Composition of the Shionogi Carcinoma. *Cancer Research*, 50(8), 2275–2282.
- [162] Newville, M., Stensitzki, T., Allen, D. B., & Ingargiola, A. (2014). LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python.
- [163] Krzanowski, W. J. (2010). *An Introduction to Statistical Modelling*. Wiley.
- [164] Rockwood, L. L. (2015). *Introduction to population ecology*. John Wiley & Sons.
- [165] Kingsland, S. (1982). The Refractory Model: The Logistic Curve and the History of Population Ecology. *The Quarterly Review of Biology*, 57(1), 29–52.
- [166] Kuno, E. (1991). Some strange properties of the logistic equation defined with r and K : Inherent defects or artifacts? *Researches on Population Ecology*, 33(1), 33–39.
- [167] Mallet, J. (2012). The struggle for existence: how the notion of carrying capacity, K , obscures the links between demography, Darwinian evolution, and speciation. *Evolutionary Ecology Research*, 14, 627–665.
- [168] Kerr, K. M., & Lamb, D. (1984). Actual growth rate and tumour cell proliferation in human pulmonary neoplasms. *British Journal of Cancer*, 50(3), 343–349.
- [169] Cunningham, J. J. (2019). A call for integrated metastatic management. *Nature Ecology and Evolution*, 3(7), 996–998.
- [170] Mistry, H. B. (2020). Evolutionary based adaptive dosing algorithms: Beware the cost of cumulative risk. *bioRxiv*, 2020.06.23(167056), doi: <https://doi.org/10.1101/2020.06.23.167056>.

- [171] Strobl, M. A. R., Gallaher, J., West, J., Robertson-Tessi, M., Maini, P. K., & Anderson, A. R. A. (2020). Spatial structure impacts adaptive therapy by shaping intra-tumoral competition. *bioRxiv*, 2020.11.03(365163), doi: <https://doi.org/10.1101/2020.11.03.365163>.
- [172] Baker, R. E., & Simpson, M. J. (2010). Correcting mean-field approximations for birth-death-movement processes. *Physical Review E*, 82(4), 041905:1–12.
- [173] Bravo, R. R., Baratchart, E., West, J., Schenck, R. O., Miller, A. K., Gallaher, J., Gatenbee, C. D., Basanta, D., Robertson-Tessi, M., & Anderson, A. R. A. (2020). Hybrid Automata Library: A flexible platform for hybrid modeling with real-time visualization. *PLoS Computational Biology*, 16(3), e1007635.
- [174] Alden, K., Read, M., Timmis, J., Andrews, P. S., Veiga-Fernandes, H., & Coles, M. (2013). Spartan: A Comprehensive Tool for Understanding Uncertainty in Simulations of Biological Systems. *PLoS Computational Biology*, 9(2).
- [175] Hamis, S., Stratiev, S., & Powathil, G. G. (2020). Uncertainty and Sensitivity Analyses Methods for Agent-Based Mathematical Models: An Introductory Review. In *Physics Of Cancer, The: Research Advances*, (p. 1). World Scientific.
- [176] Bradski, G. (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.
- [177] Choi, Y. L., Soda, M., Yamashita, Y., Ueno, T., Takashima, J., Nakajima, T., Yatabe, Y., Takeuchi, K., Hamada, T., Haruta, H., Ishikawa, Y., Kimura, H., Mitsudomi, T., Tanio, Y., & Mano, H. (2010). EML4-ALK mutations in lung cancer that confer resistance to ALK inhibitors. *New England Journal of Medicine*, 363(18), 1734–1739.
- [178] Shi, H., Hugo, W., Kong, X., Hong, A., Koya, R. C., Moriceau, G., Chodon, T., Guo, R., Johnson, D. B., Dahlman, K. B., Kelley, M. C., Kefford, R. F., Chmielowski, B., Glaspy, J. A., Sosman, J. A., Van Baren, N., Long, G. V., Ribas, A., & Lo, R. S. (2014). Acquired resistance and clonal evolution in melanoma during BRAF inhibitor therapy. *Cancer Discovery*, 4(1), 80–93.
- [179] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, b., Feng, Y., Moore,

- E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., & SciPy 1.0 Contributors (2020). {SciPy} 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, *17*, 261–272.
- [180] Gatenbee, C. D., Schenck, R. O., Bravo, R. R., & Anderson, A. R. (2019). EvoFreq: visualization of the Evolutionary Frequencies of sequence and model data. *BMC Bioinformatics*, *20*(1), 710.
- [181] Melnikov, S. V., Stevens, D. L., Fu, X., Kwok, H. S., Zhang, J. T., Shen, Y., Sabina, J., Lee, K., Lee, H., & Söll, D. (2020). Exploiting evolutionary trade-offs for posttreatment management of drug-resistant populations. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(30), 17924–17931.
- [182] Kim, E., Brown, J. S., Eroglu, Z., & Anderson, A. R. A. (2020). Understanding the potential benefits of adaptive therapy for metastatic melanoma. *bioRxiv*, *2020.10.16*(343269), doi: <https://doi.org/10.1101/2020.10.16.343269>.
- [183] Perry, M. C. (1992). *The Chemotherapy Source Book*, vol. 117. Lippincott Williams & Wilkins.
- [184] Florian, J. A., Eiseman, J. L., & Parker, R. S. (2008). Nonlinear model predictive control for dosing daily anticancer agents using a novel saturating-rate cell-cycle model. *Computers in Biology and Medicine*, *38*(3), 339–347.
- [185] Chareyron, S., & Alamir, M. (2009). Mixed immunotherapy and chemotherapy of tumors: Feedback design and model updating schemes. *Journal of Theoretical Biology*, *258*(3), 444–454.
- [186] Harrold, J. M., & Parker, R. S. (2009). Clinically relevant cancer chemotherapy dose scheduling via mixed-integer optimization. *Computers and Chemical Engineering*, *33*(12), 2042–2054.
- [187] Chen, T., Kirkby, N. F., & Jena, R. (2012). Optimal dosing of cancer chemotherapy using model predictive control and moving horizon state/parameter estimation. *Computer Methods and Programs in Biomedicine*, *108*(3), 973–983.
- [188] Angaroni, F., Graudenzi, A., Rossignolo, M., Maspero, D., Calarco, T., Piazza, R., Montangero, S., & Antoniotti, M. (2020). An Optimal Control Framework

- for the Automated Design of Personalized Cancer Treatments. *Frontiers in Bioengineering and Biotechnology*, 8, 523:1–19.
- [189] Beckman, R. A., Schemmann, G. S., & Yeang, C. H. (2012). Impact of genetic dynamics and single-cell heterogeneity on development of nonstandard personalized medicine strategies for cancer. *Proceedings of the National Academy of Sciences of the United States of America*, 109(36), 14586–14591.
- [190] Yeang, C. H., & Beckman, R. A. (2016). Long range personalized cancer treatment strategies incorporating evolutionary dynamics. *Biology Direct*, 11(1), 56.
- [191] Franzese, E., Centonze, S., Diana, A., Carlino, F., Guerrera, L. P., Di Napoli, M., De Vita, F., Pignata, S., Ciardiello, F., & Orditura, M. (2019). PARP inhibitors in ovarian cancer. *Cancer Treatment Reviews*, 73, 1–9.
- [192] Underhill, C., Toulmonde, M., & Bonnefoi, H. (2011). A review of PARP inhibitors: From bench to bedside. *Annals of Oncology*, 22(2), 268–279.
- [193] Tangutoori, S., Baldwin, P., & Sridhar, S. (2015). PARP inhibitors: A new era of targeted therapy. *Maturitas*, 81(1), 5–9.
- [194] Farmer, H., McCabe, N., Lord, C. J., Tutt, A. N. J., Johnson, D. A., Richardson, T. B., Santarosa, M., Dillon, K. J., Hickson, I., Knights, C., Martin, N. M. B., Jackson, S. P., Smith, G. C. M., & Ashworth, A. (2005). Targeting the DNA repair defect in BRCA mutant cells as a therapeutic strategy. *Nature*, 434(7035), 917–921.
- [195] Bryant, H. E., Schultz, N., Thomas, H. D., Parker, K. M., Flower, D., Lopez, E., Kyle, S., Meuth, M., Curtin, N. J., & Helleday, T. (2005). Specific killing of BRCA2-deficient tumours with inhibitors of poly(ADP-ribose) polymerase. *Nature*, 434(7035), 913–917.
- [196] Ledermann, J., Harter, P., Gourley, C., Friedlander, M., Vergote, I., Rustin, G., Scott, C., Meier, W., Shapira-Frommer, R., Safra, T., Matei, D., Macpherson, E., Watkins, C., Carmichael, J., & Matulonis, U. (2012). Olaparib maintenance therapy in platinum-sensitive relapsed ovarian cancer. *New England Journal of Medicine*, 366(15), 1382–1392.

- [197] Banerjee, S., Moore, K. N., Colombo, N., Scambia, G., Kim, B. G., Oaknin, A., Friedlander, M., Lisyanskaya, A., Floquet, A., Leary, A., & Others (2020). Maintenance olaparib for patients (pts) with newly diagnosed, advanced ovarian cancer (OC) and a BRCA mutation (BRCAm): 5-year (y) follow-up (f/u) from SOLO1. *Annals of Oncology*, *31*, S613.
- [198] Friedlander, M., Matulonis, U., Gourley, C., du Bois, A., Vergote, I., Rustin, G., Scott, C., Meier, W., Shapira-Frommer, R., Safra, T., Matei, D., Shirinkin, V., Selle, F., Fielding, A., Lowe, E. S., McMurtry, E. L., Spencer, S., Rowe, P., Mann, H., Parry, D., & Ledermann, J. (2018). Long-term efficacy, tolerability and overall survival in patients with platinum-sensitive, recurrent high-grade serous ovarian cancer treated with maintenance olaparib capsules following response to chemotherapy. *British Journal of Cancer*, *119*(9), 1075–1085.
- [199] Poveda, A., Floquet, A., Ledermann, J. A., Asher, R., Penson, R. T., Oza, A. M., Korach, J., Huzarski, T., Pignata, S., Friedlander, M., Baldoni, A., Park-Simon, T.-W., Sonke, G. S., Lisyanskaya, A. S., Kim, J.-H., Filho, E. A., Vergote, I., Rowe, P., & Pujade-Lauraine, E. (2020). Final overall survival (OS) results from SOLO2/ENGOT-ov21: A phase III trial assessing maintenance olaparib in patients (pts) with platinum-sensitive, relapsed ovarian cancer and a BRCA mutation. *Journal of Clinical Oncology*, *38*(15), 6002–6002.
- [200] Munroe, M., & Kolesar, J. (2016). Olaparib for the treatment of BRCA-mutated advanced ovarian cancer. *American Journal of Health-System Pharmacy*, *73*(14), 1037–1041.
- [201] Koudijs, M. J., Smid, M., Helleman, J., Hoogstraat, M., Brenton, J. D., van IJcken, W. F. J., Heine, A. A. J., Helmijr, J. C. A., Murtaza, M., Ruigrok-Ritstier, K., Piskorz, A. M., Berns, E. M. J. J., Beaufort, C. M., & Besselink, N. (2014). Ovarian Cancer Cell Line Panel (OCCP): Clinical Importance of In Vitro Morphological Subtypes. *PLoS ONE*, *9*(9), e103988.
- [202] Kim, E., Rebecca, V. W., Smalley, K. S., & 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.
- [203] Kozłowska, E., Vallius, T., Hynninen, J., Hietanen, S., Färkkilä, A., & Hautaniemi, S. (2019). Virtual clinical trials identify effective combination therapies in ovarian cancer. *Scientific reports*, *9*(1), 18678.

- [204] Essen BioScience Inc. (2015). *IncuCyte(TM) ZOOM User Manual*.
- [205] Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M., & Gagné, C. (2012). {DEAP}: Evolutionary Algorithms Made Easy. *Journal of Machine Learning Research*, *13*, 2171–2175.
- [206] Maclaren, O. J., Parker, A., Pin, C., Carding, S. R., Watson, A. J., Fletcher, A. G., Byrne, H. M., & Maini, P. K. (2017). A hierarchical Bayesian model for understanding the spatiotemporal dynamics of the intestinal epithelium. *PLoS Computational Biology*, *13*(7), 1–23.
- [207] Browning, A. P., Haridas, P., & Simpson, M. J. (2019). A Bayesian Sequential Learning Framework to Parameterise Continuum Models of Melanoma Invasion into Human Skin. *Bulletin of Mathematical Biology*, *81*(3), 676–698.
- [208] Cheng, Y., Thalhauser, C. J., Smithline, S., Pagidala, J., Miladinov, M., Vezina, H. E., Gupta, M., Leil, T. A., & Schmidt, B. J. (2017). QSP Toolbox: Computational Implementation of Integrated Workflow Components for Deploying Multi-Scale Mechanistic Models. *AAPS Journal*, *19*(4), 1002–1016.
- [209] Niederer, S. A., Aboelkassem, Y., Cantwell, C. D., Corrado, C., Coveney, S., Cherry, E. M., Delhaas, T., Fenton, F. H., Panfilov, A. V., Pathmanathan, P., Plank, G., Riabiz, M., Roney, C. H., Dos Santos, R. W., & Wang, L. (2020). Creation and application of virtual patient cohorts of heart models: Virtual Cohorts of Heart Models. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *378*(2173), 20190558.
- [210] Turing, A. M. (1952). The Chemical Basis of Morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B*, *237*(641), 37–72.
- [211] Murphy, H., Jaafari, H., & Dobrovolsky, H. M. (2016). Differences in predictions of ODE models of tumor growth: A cautionary example. *BMC Cancer*, *16*(1), 163.
- [212] Tanaka, G., Hirata, Y., Goldenberg, S. L., Bruchovsky, N., & Aihara, K. (2010). Mathematical modelling of prostate cancer growth and its application to hormone therapy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *368*(1930), 5029–5044.
- [213] Lambert, B. (2018). *A Student's Guide to Bayesian Statistics*. Sage.

- [214] Hafner, M., Niepel, M., Chung, M., & Sorger, P. K. (2016). Growth rate inhibition metrics correct for confounders in measuring sensitivity to cancer drugs. *Nature Methods*, *13*(6), 521–527.
- [215] Seabold, S., & Perktold, J. (2010). statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*.
- [216] Zucchini, W., & MacDonald, I. L. (2009). *Hidden Markov Models for Time Series: An Introduction Using R*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC Press.
- [217] Scholten, K., & Meng, E. (2018). A review of implantable biosensors for closed-loop glucose control and other drug delivery applications. *International Journal of Pharmaceutics*, *544*(2), 319–334.
- [218] Khodaei, M. J., Candelino, N., Mehrvarz, A., & Jalili, N. (2020). Physiological Closed-Loop Control (PCLC) Systems: Review of a Modern Frontier in Automation. *IEEE Access*, *8*, 23965–24005.
- [219] Dumont, G. A. (2012). Closed-loop control of anesthesia - A review. In *IFAC Proceedings Volumes (IFAC-PapersOnline)*, vol. 45, (pp. 373–378). IFAC Secretariat.
- [220] Kashaninejad, N., Nikmaneshi, M. R., Moghadas, H., Oskouei, A. K., Rismanian, M., Barisam, M., Saidi, M. S., & Firoozabadi, B. (2016). Organ-tumor-on-a-chip for chemosensitivity assay: A critical review. *Micromachines*, *7*(8).
- [221] Permuth-Wey, J., & Sellers, T. A. (2009). Epidemiology of ovarian cancer. In M. Verma (Ed.) *Cancer Epidemiology. Methods in Molecular Biology*, (pp. 413–437). Humana Press.
- [222] Noordermeer, S. M., & van Attikum, H. (2019). PARP Inhibitor Resistance: A Tug-of-War in BRCA-Mutated Cells. *Trends in Cell Biology*, *29*(10), 820–834.
- [223] D’Andrea, A. D. (2018). Mechanisms of PARP inhibitor sensitivity and resistance. *DNA Repair*, *71*, 172–176.
- [224] Murai, J., Huang, S. Y. N., Das, B. B., Renaud, A., Zhang, Y., Doroshow, J. H., Ji, J., Takeda, S., & Pommier, Y. (2012). Trapping of PARP1 and PARP2 by clinical PARP inhibitors. *Cancer Research*, *72*(21), 5588–5599.

- [225] The Cancer Genome Atlas Research Network (2011). Integrated genomic analyses of ovarian carcinoma. *Nature*, 474(7353), 609–615.
- [226] LP, A. P. (2017). Lynparza (olaparib) capsules [prescribing information]. *AstraZeneca Pharmaceuticals LP: Wilmington, DE*.
- [227] Fong, P. C., Boss, D. S., Yap, T. A., Tutt, A., Wu, P., Mergui-Roelvink, M., Mortimer, P., Swaisland, H., Lau, A., O’connor, M. J., Ashworth, A., Carmichael, J., Kaye, S. B., Schellens, J. H. M., & De Bono, J. S. (2009). Inhibition of Poly(ADP-Ribose) Polymerase in Tumors from BRCA Mutation Carriers. *New England Journal of Medicine*, 361(2), 123–134.
- [228] Mateo, J., Moreno, V., Gupta, A., Kaye, S. B., Dean, E., Middleton, M. R., Friedlander, M., Gourley, C., Plummer, R., Rustin, G., Sessa, C., Leunen, K., Ledermann, J., Swaisland, H., Fielding, A., Bannister, W., Nicum, S., & Molife, L. R. (2016). An Adaptive Study to Determine the Optimal Dose of the Tablet Formulation of the PARP Inhibitor Olaparib. *Targeted Oncology*, 11(3), 401–415.
- [229] Audeh, M. W., Carmichael, J., Penson, R. T., Friedlander, M., Powell, B., Bell-McGuinn, K. M., Scott, C., Weitzel, J. N., Oaknin, A., Loman, N., Lu, K., Schmutzler, R. K., Matulonis, U., Wickens, M., & Tutt, A. (2010). Oral poly(ADP-ribose) polymerase inhibitor olaparib in patients with BRCA1 or BRCA2 mutations and recurrent ovarian cancer: A proof-of-concept trial. *The Lancet*, 376(9737), 245–251.
- [230] Labrie, M., Kim, T. B., Ju, Z., Lee, S., Zhao, W., Fang, Y., Lu, Y., Chen, K., Ramirez, P., Frumovitz, M., Meyer, L., Fleming, N. D., Sood, A. K., Coleman, R. L., Mills, G. B., & Westin, S. N. (2019). Adaptive responses in a PARP inhibitor window of opportunity trial illustrate limited functional interlesional heterogeneity and potential combination therapy options. *Oncotarget*, 10(37), 3533–3546.
- [231] Yang, W., Soares, J., Greninger, P., Edelman, E. J., Lightfoot, H., Forbes, S., Bindal, N., Beare, D., Smith, J. A., Thompson, I. R., Ramaswamy, S., Futreal, P. A., Haber, D. A., Stratton, M. R., Benes, C., McDermott, U., & Garnett, M. J. (2013). Genomics of Drug Sensitivity in Cancer (GDSC): A resource for therapeutic biomarker discovery in cancer cells. *Nucleic Acids Research*, 41(D1), D955–D961.

- [232] Sheta, R., Bachvarova, M., Plante, M., Renaud, M. C., Sebastianelli, A., Gregoire, J., Navarro, J. M., Perez, R. B., Masson, J. Y., & Bachvarov, D. (2020). Development of a 3D functional assay and identification of biomarkers, predictive for response of high-grade serous ovarian cancer (HGSOC) patients to poly-ADP ribose polymerase inhibitors (PARPis): targeted therapy. *Journal of Translational Medicine*, *18*(1), 439.
- [233] Vaidyanathan, A., Sawers, L., Gannon, A. L., Chakravarty, P., Scott, A. L., Bray, S. E., Ferguson, M. J., & Smith, G. (2016). ABCB1 (MDR1) induction defines a common resistance mechanism in paclitaxel- and olaparib-resistant ovarian cancer cells. *British Journal of Cancer*, *115*(4), 431–441.
- [234] Scarborough, J. A., McClure, E., Anderson, P., Dhawan, A., Durmaz, A., Lessnick, S. L., Hitomi, M., & Scott, J. G. (2020). Identifying states of collateral sensitivity during the evolution of therapeutic resistance in Ewing’s sarcoma. *iScience*, *23*(7), 101293.
- [235] Kim, H., Xu, H., George, E., Hallberg, D., Kumar, S., Jagannathan, V., Medvedev, S., Kinose, Y., Devins, K., Verma, P., Ly, K., Wang, Y., Greenberg, R. A., Schwartz, L., Johnson, N., Scharpf, R. B., Mills, G. B., Zhang, R., Velculescu, V. E., Brown, E. J., & Simpkins, F. (2020). Combining PARP with ATR inhibition overcomes PARP inhibitor and platinum resistance in ovarian cancer models. *Nature Communications*, *11*(1), 1–16.
- [236] Dean, E., Middleton, M. R., Pwint, T., Swaisland, H., Carmichael, J., Goodege-Kunwar, P., & Ranson, M. (2012). Phase I study to assess the safety and tolerability of olaparib in combination with bevacizumab in patients with advanced solid tumours. *British Journal of Cancer*, *106*(3), 468–474.
- [237] Norris, R. E., Adamson, P. C., Nguyen, V. T., & Fox, E. (2014). Preclinical evaluation of the PARP inhibitor, olaparib, in combination with cytotoxic chemotherapy in pediatric solid tumors. *Pediatric Blood & Cancer*, *61*(1), 145–150.
- [238] Yamamoto, K. N., Hirota, K., Takeda, S., & Haeno, H. (2014). Evolution of Pre-Existing versus Acquired Resistance to Platinum Drugs and PARP Inhibitors in BRCA-Associated Cancers. *PLoS ONE*, *9*(8), e105724.

- [239] Friedlander, M., Banerjee, S., Mileskin, L., Scott, C., Shannon, C., & Goh, J. (2016). Practical guidance on the use of olaparib capsules as maintenance therapy for women with BRCA mutations and platinum-sensitive recurrent ovarian cancer. *Asia-Pacific Journal of Clinical Oncology*, *12*(4), 323–331.
- [240] Lorusso, D., Bologna, A., Cecere, S. C., De Matteis, E., Scandurra, G., Zammagni, C., Arcangeli, V., Artioli, F., Bella, M., Blanco, G., Cardalesi, C., Casartelli, C., De Vivo, R., Di Napoli, M., Gisone, E. B., Lauria, R., Lissone, A. A., Loizzi, V., Maccaroni, E., Mangili, G., Marchetti, C., Martella, F., Naglieri, E., Parolin, V., Ricciardi, G., Ronzino, G., Salutari, V., Scarfone, G., Secondino, S., Spagnoletti, I., Tasca, G., Tognon, G., & Guarneri, V. (2020). Sharing real-world experiences to optimize the management of olaparib toxicities: a practical guidance from an Italian expert panel. *Supportive Care in Cancer*, *28*(5), 2435–2442.
- [241] Frankenstein, Z., Basanta, D., Franco, O. E., Gao, Y., Javier, R. A., Strand, D. W., Lee, M. J., Hayward, S. W., Ayala, G., & Anderson, A. R. (2020). Stromal reactivity differentially drives tumour cell evolution and prostate cancer progression. *Nature Ecology and Evolution*, *4*(6), 870–884.
- [242] Hussain, M., Tangen, C. M., Higano, C. S., Vogelzang, N. J., Thompson, I. M., Berry, D. L., Higano, C. S., Crawford, E. D., Liu, G., Wilding, G., Prescott, S., Sundaram, S. K., Small, E. J., Dawson, N. A., Donnelly, B. J., Venner, P. M., Vaishampayan, U. N., Schellhammer, P. F., Quinn, D. I., Raghavan, D., Ely, B., Moynour, C. M., Vogelzang, N. J., & Thompson, I. M. (2016). Evaluating intermittent androgen-deprivation therapy phase III clinical trials: The devil is in the details. *Journal of Clinical Oncology*, *34*(3), 280–285.
- [243] Seber, G. A., & Wild, C. J. (2003). *Nonlinear Regression*. Wiley-Interscience.
- [244] Daly, A. C., Gavaghan, D., Cooper, J., & Tavener, S. (2018). Inference-based assessment of parameter identifiability in nonlinear biological models. *Journal of the Royal Society Interface*, *15*(144).
- [245] Iwata, K., Kawasaki, K., & Shigesada, N. (2000). A dynamical model for the growth and size distribution of multiple metastatic tumors. *Journal of Theoretical Biology*, *203*(2), 177–186.

- [246] Cao, Y., Jiang, Y., Lin, X., Liu, J., Lu, T., Cheng, W., & Yan, F. (2018). Dynamic prediction of outcome for patients with ovarian cancer: Application of a joint model for longitudinal cancer antigen 125 values. *International Journal of Gynecological Cancer*, 28(1), 85–91.
- [247] Zaman, S. B., Hussain, M. A., Nye, R., Mehta, V., Mamun, K. T., & Hossain, N. (2017). A Review on Antibiotic Resistance: Alarm Bells are Ringing. *Cureus*, 9(6), e1403.
- [248] Husain, M., Jagdeesh Prasad Rathore, I., Sharma, A., Raja, A., Injila Qadri, I., Waheed Wani, I. A., Prasad Rathore, J., Qadri, I., & Waheed Wani, A. (2018). Description and management strategies of important pests of pear: A review. *Journal of Entomology and Zoology Studies*, 6(3), 677–683.
- [249] Goel, N. S., & Richter-Dyn, N. (2016). *Stochastic models in biology*. Elsevier.