Mathematical modelling of autoimmune myocarditis and the effects of immune checkpoint inhibitors



Solveig Antoinette van der Vegt

St. Hilda's College

University of Oxford

Thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2022

To all women in science, past, present and future.

Voor opa.

Acknowledgements

There are too many people to count to whom I owe immeasurable thanks. To the Engineering and Physical Sciences Research Council and F. Hoffman – La Roche Ltd. for funding my research. To my supervisors at Oxford, Prof. Ruth Baker and Prof. Sarah Waters, for having the patience to take on a biologist who wanted to become a mathematician, and for the kindness with which they guided me through all the ups and downs of the project. To my supervisors at Roche, Dr. Liudmila Polonchuk and Dr. Ken Wang, for their support throughout my doctorate despite the challenges of long–distance supervision due to COVID. To everyone at WCMB for making the 4th floor a place where I truly felt like I belonged. To Giulia Celora for being an amazing friend to be on this wild DPhil journey with, from the first day of SABS through the final weeks in the MI. To Trishna Raj for always being there to help me celebrate the highs and recover from the lows. To mama, papa, Yze, Mees and Merlijn for being the most supportive family anyone could wish for. And lastly, to opa, who is not here anymore to see me graduate but who is the reason I fell in love with mathematics. I could not have done this without any of you. I hope I have made you proud.

Abstract

Autoimmune myocarditis is a rare, but frequently fatal, side effect of immune checkpoint inhibitors (ICIs), a class of cancer therapies. Despite extensive experimental work on the causes, development, and progression of this disease, much still remains unknown about the importance of the different immunological pathways involved. To further our understanding of the immunological mechanisms underlying autoimmune myocarditis and the cardiotoxicity associated with ICIs, we build the first mathematical model of this disease. This model is based on the immune cell types and interactions that have been identified by previous experimental research as playing an important role in disease development. We show that this model, when parameterized based on literature insights, can reproduce the qualitative behaviour we expect based on disease characteristics observed in patients. Using bifurcation analysis, we show that the presence of ICIs increases the risk of autoimmune myocarditis developing.

Balancing the positive effects of ICIs in treating the tumour with the detrimental effects of increasing the risk of autoimmune myocarditis, lends itself well to the use of optimal control theory. We show that, in a scenario where the standard–of–care treatment schedule used for ICIs causes the development of autoimmune myocarditis, optimal control theory can be used to establish a wide range of schedules with which the same number of doses of ICIs can be applied while autoimmune myocarditis is prevented. This allows for treatment regimen selection from a set of schedules with a minimized risk of autoimmune myocarditis, so other factors such as effect on the tumour and other possible side–effects can be considered.

Lastly, we show that our model of autoimmune myocarditis can be used as part of a larger modelling framework which explicitly considers the effects of ICIs on the tumour. We use optimal control theory on this extended model to show how this framework could be a first step towards using modelling in optimizing personalized treatment scheduling by considering the risk of autoimmune myocarditis in combination with other factors such as effect on the tumour and other possible side effects. Further extension of this framework can consider more complex tumour models as well as models of other side–effects of concern when ICIs are administered.

Contents

1 Introduction **1**

1.1 Thesis structure 3

2 The immunology underlying autoimmune myocarditis **6**

2.1 Introduction 6

2.2 Immunology 8

2.2.1 Causes 9

2.2.2 Innate immunity 11

2.2.3 Adaptive immunity 15

2.3 Disease outcomes 24

2.4 Mathematical modelling 24

2.4.1 A mathematical model of antigen presentation dynamics . . . 27

2.4.2 A mathematical model of self–tolerance 28

2.4.3 A mathematical model of innate and adaptive immunity . . . 31

2.4.4 A mathematical model of T–cell mediated autoimmunity . . . 33

2.5 Conclusion 35

3 A mathematical model of autoimmune myocarditis **37**

3.1 Introduction 37

3.1.1 Chapter outline 41

3.2 The mathematical model 42

3.2.1 The mathematical model 45

3.2.2 Steady states and linear stability 47

3.2.3 Parameter selection 49

3.2.4 Model simulation with default parameter set 55

3.3 The effects of ICIs 57

3.3.1 Nivolumab 58

3.3.2 Ipilimumab 60

3.4 Patient–specific responses 61
3.5 Parameter identifiability analysis 65
3.5.1 Structural identifiability 66
3.5.2 Practical identifiability 71
3.6 Discussion 77

4 Optimal control of autoimmune myocarditis 81
4.1 Introduction 81
4.2 The model 84
4.2.1 Standard–of–care treatment 90
4.3 Optimal control method for the timing of dosing schedules 95
4.4 Optimized ICI treatment 98
4.4.1 Nivolumab monotherapy 100
4.4.2 Ipilimumab monotherapy 106
4.4.3 Combination therapy 109
4.5 Discussion 112

5 A heart–tumour model of autoimmune myocarditis as a side–effect of immune checkpoint inhibitors 116
5.1 Introduction 116
5.2 Tumour model 118
5.2.1 Model equations 118
5.2.2 Parameter values 120
5.2.3 Optimal control of tumour growth 121
5.3 Two–compartment heart–tumour model 123
5.3.1 Coupling the tumour and heart compartments 124
5.3.2 Steady states 126
5.3.3 Parameter values 127
5.4 Optimized dosing schedules for ICI therapy 132
5.4.1 Nivolumab monotherapy 134
5.4.2 Ipilimumab monotherapy 136
5.4.3 Combination therapy 138
5.5 Discussion 141

6 Conclusion 143
6.1 Future work 147

Appendices **152**

A Exploration of optimal control parameters for autoimmune myocarditis model **153**

A.1 Nivolumab monotherapy 154
A.1.1 Time traces for varying dosing schedules 154
A.1.2 The effect of changes to V_N 154
A.1.3 The effect of extending the treatment window 157
A.1.4 The effect of changes to h 159
A.1.5 The effect of varying γ_3 159
A.1.6 The effect of changes to the payoff function 161
A.1.7 The effect of the timing of the immune perturbation and detection 163
A.2 Ipilimumab monotherapy 164
A.2.1 The effect of changes to V_I 164
A.2.2 The effect of extending the treatment window 165
A.2.3 The effect of h 166
A.2.4 The effect of varying γ_4 167
A.3 Combination therapy 168
A.3.1 The effect of changes to V_N and V_I 168
A.3.2 The effect of extending the treatment window 169
A.3.3 The effect of h 170
A.3.4 The effect of varying γ_3 and γ_4 171

B Exploration of optimal control parameters for heart–tumour model **173**

B.1 Nivolumab monotherapy 174
B.1.1 The effect of changes to V_N 174
B.1.2 The effect of changes to the final dosing day, $t_{\max}$ 175
B.1.3 The effect of changes to h 176
B.1.4 The effect of payoff function weights 177
B.2 Ipilimumab monotherapy 178
B.2.1 The effect of changes to V_I 178
B.2.2 The effect of changes to the final dosing day 180
B.2.3 The effect of h 181
B.2.4 The effect of payoff function weights 181
B.3 Combination therapy 182
B.3.1 The effect of changes to V_N and V_I 182
B.3.2 The effect of extending the treatment window 184

B.3.3 The effect of h 185
B.3.4 The effect of payoff function weights 186

Bibliography **188**

List of Figures

2.1 An overview of the progression of the immune response with two main arms: the innate response and the adaptive response. 9
2.2 Main functions of the four CD4+ T cell subtypes. 21
2.3 Overview of the immunology behind autoimmune myocarditis. 25
2.4 Graphical representation of the model developed by Sontag [142]. . . 27
2.5 Graphical representation of model by Carneiro et al. [22]. 29
2.6 Graphical representation of the model developed by Su et al. [144]. . 32
2.7 Graphical representation of the model developed by Delitala et al. [35]. 35

3.1 An overview of the immunology underpinning the development and progression of autoimmune myocarditis. 38
3.2 Illustration of the steady state and stability concepts related to the hypothesized mechanism of autoimmune myocarditis and the role of ICIs therein. 39
3.3 Graphical representation of the mathematical model of autoimmune myocarditis, indicating the effects of nivolumab and ipilimumab in the network of interactions. 44
3.4 Parameter sensitivity analysis with baseline parameter values as given in Table 3.1. 55
3.5 Time traces of the model variables for different initial conditions with default parameter values. 56
3.6 Bifurcation diagrams for varying values of a_6 and d_4 as all other parameters are kept at default values. 58
3.7 Bifurcation diagrams for varying values of a_6 and d_5 as all other parameters are kept at default values. 60
3.8 Dynamics of dead/damaged cardiomyocytes in response to increasing levels of nivolumab and ipilimumab. 62

3.9 Varying outcomes for different values of a_1 to a dose of nivolumab, ipilimumab, or both, as indicated by the number of damaged/dead cardiomyocytes. 64
3.10 A representative simulation (blue line) and resulting synthetic data set (red circles). 72
3.11 Maximum loglikelihood profiles of the subset of parameters d_1, d_3, d_4, d_5, ϕ in the ICI–free model. 74
3.12 Maximum loglikelihood profiles for the parameter subset a_1, a_2, a_4, a_5. 76
3.13 Maximum loglikelihood profiles for the parameter subset a_1, a_2, d_1 . . 77

4.1 Time traces of the model variables in an ICI–free environment. An immune perturbation of 100 pathogenic T cells occurs on day 9. . . . 91
4.2 When nivolumab is administered every two weeks with $V_N = 0.25$ grams, autoimmune myocarditis develops. 92
4.3 When ipilimumab is administered every three weeks with $V_I = 1.0$ gram, autoimmune myocarditis develops. 93
4.4 When combination therapy is administered with $V_N = 0.25$ grams and $V_I = 0.1$ grams on a standard–of–care treatment schedule, autoimmune myocarditis develops. 95
4.5 The dose schedule for nivolumab monotherapy and payoff evolve over iterations of the optimization algorithm. 102
4.6 Time traces of state variables when the optimized dosing schedule for nivolumab monotherapy is applied. 102
4.7 Time traces of pathogenic CD4+ T cells for nivolumab dosing schedules which do and do not cause autoimmune myocarditis to develop. . . . 104
4.8 Evolution of the dosing schedule schedule and payoff over the iterations as the ipilimumab monotherapy treatment is optimized. 107
4.9 Time traces of state variables when the optimized dosing schedule for ipilimumab monotherapy is applied. 108
4.10 Evolution of the dosing schedule and payoff over the iterations as we optimize combination therapy treatment. 110
4.11 Time traces of state variables when the optimized dosing schedule for combination therapy is applied. 111

5.1 The RHS of Eq. (5.1) plotted as a function of the size of the tumour. 120
5.2 Time traces of tumour growth under ICI–free conditions, or with standard–of–care nivolumab, ipilimumab or combination therapy. 121

5.3 Timeline of the optimization problem for ICI treatment schedule design. 123
5.4 Time traces of state variables under ICI–free conditions. 128
5.5 Time traces with standard–of–care nivolumab, ipilimumab or combination therapy. 130
5.6 Dosing schedule of nivolumab monotherapy and payoff evolving over the iterations of the optimization algorithm. 135
5.7 Time traces of model variables when the optimized dosing schedule for nivolumab monotherapy is applied. 136
5.8 Dosing schedule of ipilimumab monotherapy and payoff evolving over the iterations of the optimization algorithm. 137
5.9 Time traces of model variables when the optimized dosing schedule for ipilimumab monotherapy is applied. 138
5.10 Dosing schedule of combination therapy and payoff evolving over the iterations of the optimization algorithm. 139
5.11 Time traces of model variables when the optimized dosing schedule for combination therapy is applied. 140

A.1 Time traces of state variables when nivolumab dosing schedules are obtained after 50, 100 or 150 iterations. 155
A.2 The payoff of the optimized treatment schedule for nivolumab monotherapy over different values of V_N. 156
A.3 Time traces for the optimized dosing schedule for $V_N = 0.4$ grams. . . 157
A.4 Optimized nivolumab dosing schedule and payoff for different lengths of the treatment window. 158
A.5 Evolution of the nivolumab dosing schedule and payoff over iterations when the number of dead/damaged cardiomyocytes is integrated over time in the payoff function, and time traces of state variables when the optimized dosing schedule for nivolumab monotherapy is applied. . . 162
A.6 Payoff of the optimized dosing schedule as a function of V_I. 165
A.7 The optimized dosing schedule and payoff as a function of the length of the treatment window for ipilimumab monotherapy. 166
A.8 Treatment outcome as a function of the values of V_N and V_I in combination therapy. 169
A.9 The optimized dosing schedule and payoff as a function of the length of the treatment window of combination therapy. 170

B.1 The payoff of the optimized nivolumab dosing schedule as V_N is varied. 174

B.2 Optimized nivolumab dosing schedule (left) and payoff (right) for different lengths of the treatment window. 176
B.3 The payoff of the optimized ipilimumab dosing schedule as V_I is varied. 179
B.4 Optimized ipilimumab dosing schedule and payoff for different lengths of the treatment window. 180
B.5 Treatment outcome of the optimized combination therapy dosing schedule as V_N and V_I are varied simultaneously. 183
B.6 Optimized combination therapy dosing schedule and payoff for different lengths of the treatment window. 184

Chapter 1

Introduction

As patient survival increases for many forms of cancer, the long–term effects of both the cancer itself and cancer treatments, including immune checkpoint inhibitors (ICIs), are becoming of increasing importance [66]. Immune–related adverse events, and autoimmunity in particular, are an extremely common side–effect of ICIs, a class of treatments for which approximately 40% of US cancer patients are eligible [58]. Autoimmune side–effects are experienced by 86–96% of cancer patients treated with ICIs, and in 17–59% of patients these side–effects are severe to life–threatening [59, 149].

A rare but frequently fatal side–effect of ICIs is autoimmune myocarditis, or inflammation of cardiac muscle tissue. Current data indicates that autoimmune myocarditis occurs in approximately 0.04–1.4% of cancer patients treated with ICIs, which is a significantly higher incidence rate compared to the general population, where approximately 0.022% of individuals develop this disease [20, 50, 95, 108, 170]. However, these figures likely underestimate the actual incidence rate of myocarditis as definitive diagnosis remains difficult [171].

Although the reported incidence rate is low, the high fatality rate of autoimmune myocarditis, standing at 25–50%, warrants concern [95, 170]. In common with other

inflammatory diseases such as rheumatoid arthritis and asthma, the pathology comes from the failure of inflammation to resolve itself, instead becoming a chronic condition [130]. In autoimmune myocarditis, this chronic stage is driven by CD4+ T cells and leads to irreversible tissue damage and possibly heart failure [7, 15]. Despite significant experimental research devoted to gaining an improved understanding of autoimmune myocarditis, the immunological mechanisms underlying the development and progression of this disease are not well understood [15, 20].

The inherent complexity of the immune system makes it difficult to determine the major pathways and triggers critical to the development and progression of myocarditis, and the roles that ICIs might play in these processes. This is where mathematical modelling can be helpful. It provides a framework for understanding how multiple highly complex and nonlinear processes combine to produce the pathology observed in patients. Although, to our knowledge, no mathematical model of myocarditis exists to date, using mathematics to increase understanding of pathological processes has led to significant progress in other contexts, such as cancer [16, 38, 86], neurological disorders [14, 21, 57], and cystic fibrosis [8, 127], to name but a few. In a system as overwhelmingly complex as the human immune system, making a distinction between critical and secondary processes in the development and progression of autoimmune myocarditis is essential in moving research forward efficiently.

In addition to a lack of understanding of the immunological mechanisms underlying ICI–driven autoimmune myocarditis, no pre–clinical *in vitro* cardiotoxicity assay currently exists to assess whether potential new ICIs might increase the risk of autoimmune myocarditis. Indeed, the development of such an assay critically depends on the identification of the cell types that play a critical role in the disease.

Mathematical models also lend themselves well to asking "what if" questions about complex systems. The consequences of perturbations to the system due to external interventions or natural fluctuations, for example, can be easily explored with

a mathematical model. In the context of autoimmune myocarditis, it is of significant interest to investigate how different ICI doses and dosing regimens affect the system, something that would require expensive and time–consuming experimental work and clinical trials to do in patients. Additionally, experimental work, especially when using animals and/or humans, has to consider ethical standards and guidelines, which can limit the scope, efficiency and cost-effectiveness of this type of research. With mathematical modelling, however, different treatment variations can be implemented and explored quickly and efficiently through simulations. Furthermore, techniques such as optimal control theory are available to aid in the design of optimal dosing strategies.

Our aim in this thesis is to develop the first mathematical model of autoimmune myocarditis and the effects of ICIs on the development and progression of this disease. This model can then be used as a starting point for progress towards three main objectives: (1) to better understand the immunology driving the pathology of autoimmune myocarditis; (2) to determine the minimal set of immune cell types required to design a biologically realistic pre–clinical *in vitro* assay to test compounds for cardiotoxicity; and (3) to improve the design of cancer treatments by combining mathematical modelling with techniques from optimal control theory to take into account both the effects of treatment on the tumour as well as on other organs and tissues.

1.1 Thesis structure

In **Chapter 2**, we set out the biological background needed to develop a model of autoimmune myocarditis. As we are developing the first model of this disease, we start with a comprehensive literature review of the current immunological understanding. The immense complexity of the immune system poses a significant challenge when

developing a first model of autoimmune myocarditis. Thus, we focus here on providing a broad synthesis of the immunology behind this disease and a discussion of the effects of ICIs, which provides all the relevant biological detail for this project while remaining accessible to mathematicians and other non–experts in immunology. We present a network representation of the relevant cell types and interactions which will serve as a starting point for the design of the mathematical model.

In **Chapter 3**, we present the first mathematical model of autoimmune myocarditis and the role of ICIs, motivated by the immunological discussion in Chapter 2. In the absence of experimental data to determine model parameters, we show that this model can be parameterized such that it displays the dynamic and steady state behaviours that would be expected in a model of autoimmune myocarditis. We furthermore show that the effects of ICIs on the model predictions are consistent with observations in patients. As we are interested in using this model to design an *in vitro* assay, it is important that the model parameters are identifiable when data from this experimental system becomes available, and hence we analyse the structural and practical identifiability of the model parameters.

In **Chapter 4**, we turn to optimal control theory to provide a framework for optimizing dosing regimens for ICI treatment such that autoimmune myocarditis does not develop. The algorithm we use optimizes the time points at which a given dose of ICIs is administered by employing a gradient descent method to minimize the damage to cardiac tissue. This chapter is a first step towards explicitly including the side–effects of cancer treatment in the design of these therapies.

In **Chapter 5**, we extend the work in Chapter 4 by adding a tumour compartment to the model. When we optimize dosing schedules, the added compartment allows us to consider both the positive effects that ICI treatment has on inhibiting tumour growth, and the negative effects of ICIs in increasing the risk of autoimmune myocarditis developing simultaneously. The dosing schedules found by applying the

same optimal control algorithm to this extended model display significant qualitative differences compared to those identified in Chapter 4, emphasizing the importance of a holistic approach to cancer therapy modelling where the effects of therapies on both the tumour *and* on other organs and tissues are explicitly considered.

In **Chapter 6**, we conclude by summarizing the main findings and results presented in this thesis. We emphasize the novelty of this work in explicitly modelling the side–effects of cancer treatment. We furthermore highlight how future work can build upon these results to make cancer therapy more effective and safer, both through aiding in the design of an *in vitro* cardiotoxicity assay, and through emphasizing the need for a holistic approach to mathematical modelling of cancer therapy.

Chapter 2

The immunology underlying autoimmune myocarditis

The work contained in this chapter has been published in the article "A model–informed approach to assess the risk of immune checkpoint inhibitor–induced autoimmune myocarditis" by S.A. van der Vegt, Y.J. Wang, L. Polonchuk, K. Wang, S.L. Waters and R.E. Baker in *Frontiers in Pharmacology: Experimental Pharmacology and Drug Discovery*, 2022 [153]. I conducted the literature review and wrote the manuscript, which was edited in collaboration with the other authors.

2.1 Introduction

In this chapter, we review the current understanding of the causes (Section 2.2.1) and pathways (Section 2.2.2–2.3) involved in the development and progression of autoimmune myocarditis, and how these are affected by the presence of ICIs.[1] We do so in order to pave the way for the development of mathematical models that can be used to identify the most critical cell types and interactions involved in this pathology. Although various factors can cause myocarditis, it is generally believed that a transition

[1]Frequently used abbreviations in this chapter are listed in Table 2.1.

to T–cell mediated autoimmunity eventually takes place, and we therefore focus on these pathways in our discussion of adaptive immunity (Section 2.2.3) [15, 117]. As a starting point for a mathematical model of autoimmune myocarditis, we summarize the relevant immunology in a network graph containing the main cell types and interactions (see Fig. 2.3). In Section 2.4, we discuss several mathematical models each describing a different part of the immune system. None of these models suffice on their own to describe the immunology specific to autoimmune myocarditis, but they can serve as inspiration for the first mathematical model of this disease.

The literature on the immunology underlying autoimmune myocarditis discussed below includes studies with mice, as well as data from human patients. Generally, care should be taken when translating conclusions from animal studies to humans. The general structure of the murine immune system is remarkably similar to that of the human immune system, and many great insights into human immunology have been gained from studies using mice. However, there are a number of significant differences, including, most notably for this work, in the T cell pathway [94]. Conclusions based on work with mice can be, but are thus not necessarily, true for humans. In the case of autoimmune myocarditis, we must consider insights from mice in addition to human patient data as the data available from humans alone is insufficient to formulate a mathematical model. When it comes to validation of such a model in the future, however, the source of the data with which the model is validated should be carefully considered. If we wish to ensure that our insights apply to human patients, validation with human patient data will be an essential first step. As more insights become available into autoimmune myocarditis in humans and how the underlying mechanisms might differ from those in mice, the model should be tested against these insights continuously and adjusted where necessary.

The immune response can be roughly divided into a fast, non–specific innate immune response and a slowly developing, antigen–specific adaptive immune response

Abbreviation	Meaning
ICI	Immune checkpoint inhibitor
APC	Antigen presenting cell
DC	Dendritic cell
(i/n)Treg cell	(induced/natural) T regulatory cell
Th cell	T helper cell
TCR	T cell receptor
PMN	Polymorphonuclear leukocyte
M1/M2	Macrophages of phenotype 1/2
MyHCα	Myosin heavy chain α
PD–(L)1	Programmed death–(ligand) 1
CTLA4	Cytotoxic T–lymphocyte–associated protein 4
MHC	Major histocompatibility complex
IFN	Interferon
IL	Interleukin
TGFβ	Transforming growth factor β
NET	Neutrophil extracellular trap
ECM	Extracellular matrix

Table 2.1: Frequently used abbreviations.

[133]. The innate immune response is driven by macrophages, polymorphonuclear leukocytes (PMNs) and fibroblasts, amongst other cell types, and responds directly to anything that is not recognized as self. The adaptive immune response develops through antigen–presenting cells (APCs) which recognize non–self antigens and present them to T cells and B cells. Fig. 2.1 gives a broad overview of the immune response.

2.2 Immunology

The immune response is a powerful defense mechanism that, through its different pathways, launches a quick and specific attack against foreign pathogens. In autoimmune myocarditis, the full force of the immune system is directed against a self–antigen, a molecule that is a natural part of the body and to which the immune system is normally tolerant. This can have potentially lethal consequences [20, 85].

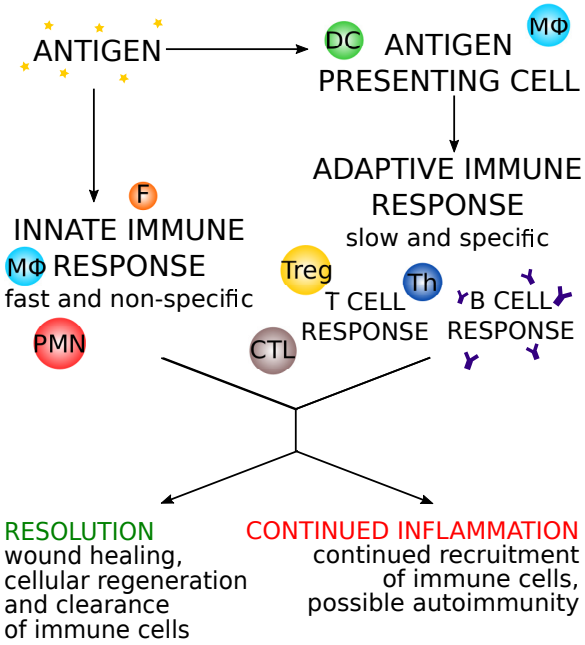


Figure 2.1: An overview of the progression of the immune response with two main arms: the innate response and the adaptive response. DC: dendritic cells, MΦ: macrophages/monocytes, PMN: polymorphonuclear leukocytes, F: fibroblast, Th: T helper cell, Treg: T regulatory cell, CTL: cytotoxic T cell.

Below, we discuss the different causes of autoimmune myocarditis, as well as the development and possible outcomes of this disease.

2.2.1 Causes

The antigens stimulating the autoimmune response in myocarditis can be released through a number of processes, including infection and spontaneous or drug–induced cell damage. Whatever the initial aetiology, myocarditis will, if unresolved, transition to a T–cell mediated autoimmune disease [20]. There are several factors that can predispose an individual to developing autoimmune myocarditis, including certain human leukocyte antigen genes, being male, and deficiency in thymic training or T regulatory (Treg) cell induction [15].

One of the main proteins believed to be a self–antigen targeted in autoimmune myocarditis is myosin heavy–chain α (MyHCα) [20, 85]. This intracellular protein,

although not abundant, is unique to cardiomyocytes and is what is known as a cryptic antigen, meaning it is not or very little expressed in the thymus or peripheral lymphoid stromal cells [85, 122]. Peripheral lymphoid stromal cells are important for mediating peripheral tolerance [85]. The thymus is a lymphoid organ where T cells mature, and which provides central tolerance. Presentation of a molecule in the thymus is required for the immune system to recognize it as "self" and tolerate it, *i.e.* not respond to it. As MyHCα is only present at very low levels the thymus, T cells with TCRs specific to this antigen are not completely removed and MyHCα–specific T cells are known to be in circulation in humans [20, 52, 80]. This can lead to problems in several ways. For example, certain proteins of the commensal gut bacterial species *Bacteroides* mimic MyHCα, and an immune reaction against these bacteria can lead to severe autoimmune myocarditis when a second inflammatory trigger exposes the MyHCα in cardiac tissue [52]. Similarly, it has been found in certain patients that the T cells mobilized by cancer treatment are specific to an antigen that is present in the tumour as well as the heart and that this cross–reaction might constitute one pathway via which ICIs induce myocarditis [117]. Additionally, experiments have shown that knocking out the gene for programmed cell death protein 1 (PD–1) leads to increasing levels of autoantibodies against cardiac antigens, and that knocking out the gene for cytotoxic T–lymphocyte–associated protein 4 (CTLA4) causes increased T cell infiltration across a number of organs, including the heart [103, 170]. Both PD–1 and CTLA4 are involved in key regulatory pathways that control the immune response and prevent autoimmunity. They are also targets of ICIs; for example they are blocked by nivolumab and ipilimumab, both of which have been shown to have autoimmune myocarditis as a possible side effect [68].

Thus, several potential triggers, including ICIs, genetic characteristics and molecular mimicry, can increase the likelihood of an autoimmune reaction against cryptic cardiac antigens such as MyHCα, and thus the development of myocarditis.

2.2.2 Innate immunity

The immune response can be broadly divided into two pathways: the fast, non–specific innate immune response and the slower, pathogen–specific adaptive immune response. The former (also known as the acute inflammatory response) relies mainly on tissue–resident immune cells, and recruited PMNs and macrophages, whereas the latter (also known as chronic inflammation) is executed by activated T cells and B cells that provide antigen–specific defence and memory functions [65, 133].

The innate immune response is instrumental in the first few hours and days after infection or injury. It is activated in response to cues of tissue injury, or recognition of conserved motifs of foreign pathogens [6, 133]. A major pathway in the innate immune response is the complement cascade which produces both cytokines that attract PMNs, and opsonic and lytic factors that lead to the destruction of invading pathogens [133]. Macrophages play an important role in mediating the innate immune response, clearing pathogens and resolving inflammation [15]. Activated resident fibroblasts are important for maintaining tissue integrity and wound healing [137]. Below we discuss the roles of macrophages, PMNs, and fibroblasts further.

2.2.2.1 Macrophages

Macrophages are among the most abundant cell types to accumulate in cardiac tissue in autoimmune myocarditis [31]. Tissue–resident cardiac macrophages originate from embryonic cells and, during homeostasis, are primarily maintained through local proliferation [42]. Furthermore, throughout life cells of the monocyte–macrophage lineage are created in the bone marrow and released into the blood stream as monocytes, which can infiltrate cardiac tissue during inflammation [42, 133, 135]. When an invading pathogen or excessive cell death causes the release of pathogen–associated or damage–associated molecular patterns, tissue–resident macrophages act as sentinel cells. They initiate the recruitment of PMNs such as neutrophils by increasing their

chemokine output, thus promoting the development of the innate immune response [75]. During inflammation, tissue macrophages expand through proliferation and infiltrating macrophages are recruited from the blood stream and activated by T helper (Th) cells [42, 67].

Activated macrophages sit on a phenotypic spectrum between M1 and M2, which are pro– and anti–inflammatory phenotypes, respectively. The M1 phenotype is induced by Th1–secreted cytokines such as IFNγ and tumour necrosis factor α. Differentiation towards the M2 phenotype is stimulated by the Th2–derived cytokines interleukin (IL)–4 and IL–13, as well as IL–10, transforming growth factor (TGF)β and other anti–inflammatory compounds [65, 135]. Macrophages retain a high level of phenotypic plasticity over their life time, allowing them to respond to environmental cues. Their origin somewhat steers their behaviour, however, with tissue–resident macrophages maintaining an essential cardioprotective role, and infiltrating macrophages being important for coordinating inflammation while being less involved in antigen sampling and removal of dead cardiomyocytes [36, 42].

Macrophages have three main functions: (i) phagocytosis of pathogens, dead or infected cells and cell debris; (ii) antigen presentation; and (iii) secretion of immune–modulating cytokines [135]. The level to which macrophages perform each of these functions depends on where on the phenotypic spectrum they sit. M1 macrophages stimulate the Th1 cell–mediated immune response by secreting IL–12, IL–15 and IL–18, and have great phagocytic and antigen–presenting abilities. M2 macrophages modulate the Th2 response, are involved in tissue remodeling and wound healing, and immune tolerance, and they generally decrease the number of active T cells [15, 133, 164]. They secrete TGFβ and IL–10, both anti–inflammatory cytokines. A proper balance between the two phenotypes is essential for an effective immune response [9].

2.2.2.2 Polymorphonuclear leukocytes

Polymorphonuclear leukocytes are a group of immune cells that include neutrophils, eosinophils, basophils and mast cells [133]. Neutrophils and eosinophils in particular have been implicated in myocarditis [37, 156].

Neutrophils phagocytose pathogens, after which granules containing highly toxic reactive oxygen species merge with the vacuole containing the pathogen. This process can contribute to inflammation when the contents of the granules leak out and end up on normal tissue. Another major effector pathway for neutrophils is through neutrophil extracellular traps (NETs). NETs are a scaffold of decondensed (nuclear and mitochondrial) chromatin loaded with cytosolic and granule proteins. They trap and kill most kinds of foreign pathogens and inhibit the dissemination of some. However, NETs are known to lead to tissue injury and play a role in immune–mediated diseases, including myocarditis [91, 110, 156]. Furthermore, neutrophils can release numerous toxic compounds, such as reactive oxygen species, that can damage cardiac tissue and even induce apoptosis in cardiomyocytes [49, 101]. Lastly, neutrophils provide positive feedback by stimulating the maturation and/or activation of macrophages, dendritic cells (DCs), Th cells and natural killer cells, amongst others [71].

Eosinophils can likewise degranulate to attack foreign pathogens [37]. They are also heavily involved in regulating both the innate and adaptive immune responses through the secretion of multiple cytokines with immune–modulating functions. They can, for example, stimulate the differentiation and recruitment of Th2 cells, and act as APCs [158].

Both neutrophils and eosinophils can cause significant damage to the healthy tissue surrounding the inflamed site [37, 101]. Because cardiomyocytes have little to no regenerative capabilities, this can have severe consequences for heart function as damaged cells will be replaced by stiff collagenous fibres and scar tissue will develop [61, 133].

2.2.2.3 Fibroblasts

The microenvironment of cells contains extracellular matrix (ECM) proteins, and blood and lymph vessels, but is primarily defined by stromal cells. The most common type of stromal cells are fibroblasts, which play an important role in maintaining and modulating the ECM, and in the production of cytokines to recruit other inflammatory cells [133, 143]. For example, IL–17 induces the production and excretion of myeloid cytokines, particularly granulocyte–macrophage colony–stimulating factor, in cardiac fibroblasts. These cytokines recruit neutrophils and induce a pro–inflammatory, pro–fibrotic phenotype in macrophages [62, 163].

In myocarditis, cardiac resident fibroblasts are activated through damage–associated molecular patterns released by damaged or dead cardiomyocytes, and by macrophages through the cytokines TGFβ and tumour necrosis factor α [15, 62, 118, 137]. Because fibroblasts are critical for maintaining tissue homeostasis and wound healing through ECM remodelling, there must be a precise balance in the timing and level of their activation. One of the first steps in innate immunity is to activate tissue–resident fibroblasts, at which point they become myofibroblasts. The early presence of myofibroblasts can decrease the severity of myocarditis as they stimulate wound healing. However, if fibroblasts remain active for too long they can cause tissue fibrosis and pathological scarring. This is especially detrimental to organs like the heart where fibrosis inhibits proper muscle function and can lead to heart failure [7, 143, 151].

High levels of complement proteins, PMNs, clotting factors and edema fluid at the site of infection or injury are characteristic of acute inflammation. If the wound can be healed or the pathogen can be cleared, the innate immune response is resolved and either tissue regeneration or fibrotic scarring restores structural integrity to the site of inflammation. If resolution is not possible, a transition from acute to chronic inflammation takes place [133].

2.2.3 Adaptive immunity

The adaptive immune response can take up to seven days to fully develop, hence the need for innate immunity to keep the infection in check initially. It consists primarily of a T–cell mediated and a B–cell mediated pathway [86]. Upon activation, B cells differentiate into either plasma cells, which secrete high amounts of antibodies, or memory cells, which are long–lived cells that can quickly produce large amounts of pathogen–specific antibodies if an infection returns [39]. Myocarditis is a CD4+ T–cell mediated disease and these cells are essential for its development and progression, but B cells are mostly absent in inflamed cardiac tissue [4, 31, 93]. In addition to CD4+ T cells, CD8+ T cells have also been found to infiltrate cardiac muscle tissue during inflammation in monkeys and humans suffering from autoimmune myocarditis [68, 108]. However, studies with mice and observational data obtained from humans show that CD4+ T cells are the primary drivers of autoimmune myocarditis [15, 93]. Below we will discuss APCs and different types of T cells, focusing primarily on CD4+ T cells.

2.2.3.1 Antigen–presenting cells

APCs, including DCs, monocytes and macrophages, are essential in connecting the innate and adaptive immune responses [51]. APCs patrol the body picking up pathogens, related antigens and cell debris on the way, or encounter them in their resident tissue [56]. They use both cell surface pattern recognition receptors, and intracellular receptors to recognize pathogens. Once they have ingested a foreign pathogen or other antigen, they break down the proteins and present them on major histocompatibility (MHC) II complexes on their cell surface, ready to prime naïve T cells, which are mature but inactive [55, 84, 109, 133]. Circulating DCs are immature, a state characterized by high phagocytic ability and low expression of costimulatory surface proteins [72, 92, 150]. Once they have picked up an antigen, they mature and

drain to the lymph nodes.

Maturing DCs stop their uptake of antigens and specialize in antigen presentation and T cell interactions instead [92, 109]. An important part of this process is an increase in lysosomal function [150]. Immature DCs are inefficient in the way they take up and process antigens. Maturation includes the activation of a proton pump which acidifies the lysozome, aiding proteolysis and efficient formation of peptide–MHC II complexes [150]. Maturing DCs further experience a large increase in synthesis of cytokines and co–stimulatory molecules. One of the cytokines that maturing DCs excrete is interferon (IFN) γ, which stimulates the maturation process in an autocrine fashion [109]. Other cytokines will function as "third signals" once a DC starts interacting with T cells and steering their differentiation [20]. Co–stimulatory proteins like CD40, CD80 and CD86 are transported to the surface of the DC and will likewise play an important role in T cell interactions. Other functional and phenotypical changes include the formation of dendrites, reorganization of the cytoskeleton and the expression of adhesive molecules for migration [92]. APCs need to drain from the periphery where they have picked up the antigen, to the secondary lymph tissue or spleen [56, 92, 109]. Once there, they migrate to the T cell zone of the lymph node where they interact with approximately 500–5000 T cells per hour. The minimum number of DCs per lymph node required to elicit a T cell response is estimated to be around 100. An immune stimulus that does not activate enough APCs to meet this threshold will thus not elicit a T cell response [26, 63, 102].

2.2.3.2 T cells

Development T cell progenitors originate in the bone marrow and migrate to the thymus, where they mature [76]. During this process, random genetic rearrangements give each cell a unique set of α and β chains in their TCR, specifying a unique target for them to recognize and respond to. Approximately 10^8 different recombinant TCRs can be generated in this way [77]. Because T cells need to be prevented from

recognizing the body's own proteins and causing an autoimmune reaction, any T cells with TCRs that recognize self–antigens are deleted through negative selection [100, 133, 139]. In order for a protein to be recognized as "self" and not elicit an immune response, it must therefore be presented in the thymus [15, 85]. As MyHCα is only present at low levels, however, T cells with TCRs specific to this antigen are present in the body [85, 122].

Additionally, T cells with TCRs that have too low or too high an affinity for peptide–MHC (I or II) complexes are selected against, leaving only cells with mid–range affinities to differentiate into CD4+ and CD8+ T cells, primarily [84]. CD4+ and CD8+ refer to the specific cell surface receptors present. When activated, CD4+ and CD8+ T cells have distinct phenotypic fates, with CD4+ T cells becoming largely Th and Treg cells and CD8+ T cells becoming cytotoxic T cells. Thymic output of T cells is mainly important during the first couple of decades of a person's life and ceases completely after 40–50 years [76]. Because the production of naïve T cells is a combination of thymic output and peripheral proliferation, it is difficult to estimate exactly how many new cells are produced by the thymus each day but an upper limit is considered to be around 10^8 T cells (both CD4+ and CD8+) per day in humans [155].

All these checks and balances result in a tightly regulated pool of approximately 10^{11} naïve CD4+ T cells, with on average 10^5 cells for each particular epitope – the specific part of the antigen that the T cell recognizes [77, 97]. Exactly how many epitope–specific naïve T cells are in circulation depends on the frequency with which the epitope in question is encountered [77]. The general size and diversity of the pool of circulating T cells is regulated by a number of soluble factors, most importantly IL–7 [97].

Activation Mature, naïve T cells leave the thymus and circulate through the body, where they will survive for five to 10 years [76]. When they pass through the lymph nodes, they remain there for about half a day before going back into circulation. While in the lymph nodes they scan approximately 160–200 APCs per hour [102]. In order for a T cell to be activated, its TCR must recognize the antigen and one of the cell's co–stimulatory receptors, most importantly CD28, must simultaneously bind one of its ligands, e.g. CD80 or CD86, which are upregulated on the surface of activated APCs [55, 64, 84, 133]. Together, these two signals induce activation, clonal expansion and differentiation of T cells. If a T cell is stimulated by an antigen but does not bind any co–stimulatory molecules it will go into anergy, a state in which the cell is hyporesponsive [160].

In addition to negative selection in the thymus, activation of T cells is an important immune checkpoint and plays a significant role in tolerance of self–antigens. Negative co–stimulatory factors such as CTLA4 and PD–1 are essential immune regulators that prevent the activation of self–specific T cells, and instead induce anergy. Anergic T cells are tolerant towards the self–antigen in question, and are characterized by their inability to proliferate or produce cytokines [1, 100].

A T cell that recognizes the antigen presented by an APC has to stay in contact with that APC for approximately one day to be activated [102]. Clonal expansion then takes another three to five days before the activated T cells can differentiate and go back into circulation to attack the source of the antigen [67].

During clonal expansion CD4+ and CD8+ T cells divide on average seven and 15 times, respectively [102]. The daughter cells that an activated T cell produces during clonal expansion will differentiate based on soluble mediators [64, 102]. These "third signals" (in addition to the antigen, the "first" signal, and co–stimulatory surface proteins, the "second" signal) consist mainly of cytokines produced by APCs and innate immune cells [64, 84, 113]. Once differentiated, T cells can produce additional

cytokines that create positive feedback loops, marginally affecting differentiation [84]. All together, these cytokines are essential for the development of pro–inflammatory T cells [20].

The effects of ICIs on T cell activation The ICIs ipilimumab and nivolumab both target surface proteins involved in the regulation of T cell activation [60]. Nivolumab is an antibody against PD–1, which is found on the surface of T cells, particularly Th cells [126]. T cell activity is inhibited when PD–1 binds one of its ligands, PD–L1 or PD–L2. PD–L1 is expressed on the surface of peripheral cells, such as cardiomyocytes, in response to pro–inflammatory cytokines. PD–L2 is found on the surface of DCs in response to antigen uptake, and the binding of PD–L2 to PD–1 protects the DC against the cytotoxic effects of the activated T cell. It also prevents the overactivation of T cells [15, 60, 80, 83]. Binding of PD–1 to either of its ligands induces anergy in the T cell. Blocking PD–1, such as with nivolumab, leads to the stimulation of proliferation, activation and effector functions in CD4+ T cells, thus increasing anti–tumour T cell activity. However, it also increases the risk of off–target inflammation in other organs, such as the heart in the case of autoimmune myocarditis [74].

Ipilimumab is an antibody targeting CTLA4, which is present on the cell surface of activated T cells. Treg cells in particular constitutively express high levels of CTLA4 [15, 60, 80, 117, 131]. CTLA4 binds to CD80 and CD86 with a higher affinity than CD28, which inhibits T cell activation in a number of ways. First, naïve T cells must compete directly with Treg cells for the binding of CD80 and CD86 on DCs [113]. Second, and most importantly, Treg cells remove CD80 and CD86 from DCs through the process of CTLA4–mediated transendocytosis [107]. Furthermore, the binding of CTLA4 to CD80 or CD86 stimulates the production of molecules that make the DC toxic to nearby T cells [125]. Lastly, CTLA4 inhibits the activation of naïve T cells

by stimulating intracellular phosphatase activity which, in turn, inhibits downstream signaling of TCRs and CD28. CTLA4 thus interferes with the functioning of CD28 and the activation of T cells in two ways: directly inhibiting the activity of the DC carrying CD80/CD86; and indirectly inhibiting the activity of T cells through DCs [113, 124]. Treatment with ipilimumab inhibits these regulatory mechanisms and thus leads to a significant increase in T cell activity, both against the tumour and potentially against other organs, causing side–effects such as autoimmune myocarditis.

Differentiation The differentiation of CD4+ T cells is crucially dependent on the cytokines present. IFNγ and IL–12 stimulate differentiation towards the Th1 phenotype [133]. IL–12 is largely produced by DCs, and induces production of IFNγ in natural killer cells [84]. IL–2 and IL–4 stimulate differentiation towards the Th2 phenotype. Furthermore, IL–6 plays a dual role in Th2 differentiation as it both stimulates differentiation towards a Th2 phenotype through IL–4 dependent mechanisms, and inhibits differentiation towards a Th1 phenotype by interfering with signaling downstream of IFNγ [84]. IL–6, IL–21, IL–23 and TGFβ stimulate the differentiation of Th17 cells in three stages [84, 128, 133, 172]. The first stage is coordinated by IL–6 and TGFβ and consists of differentiation. It is important that TGFβ is present at low levels and that IL–6 is present simultaneously, otherwise differentiation is diverted towards a regulatory phenotype. The second stage entails self–amplification and is regulated by IL–21, followed by a stabilization stage which involves APC–produced IL–23 [84]. Differentiation towards an induced Treg (iTreg) phenotype requires high levels of TGFβ and the presence of IL–2 [69, 84, 128, 133]. The main sources of TGFβ are macrophages and T cells [29]. IL–2 is mainly produced by activated CD4+ T cells, although naïve CD8+ T cells, DCs and cells in the thymus also contribute [99].

Function: T helper cells Once differentiated, each of the T cell species has a distinct role (see Fig. 2.2 for an overview of CD4+ T cells). Th1 cells mainly pro-

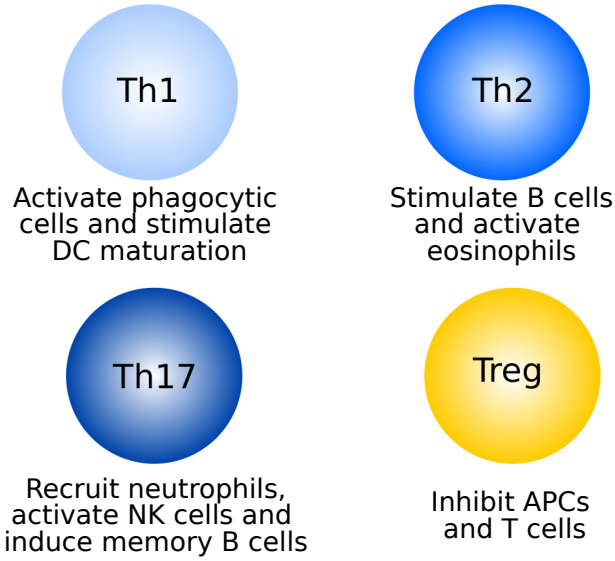


Figure 2.2: Main functions of the four CD4+ T cell subtypes.

duce IFNγ, through which they activate macrophages and other phagocytic cells, self–amplify differentiation towards the Th1 phenotype, and support DC maturation [12, 84, 147, 172]. Th17 cells activate and recruit neutrophils to the site of inflammation. They do this through their main cytokine, IL–17 [133]. The self–amplifying cytokine of Th17 cells is IL–21, which furthermore activates NK cells and induces memory functions in B cells [84]. In the context of myocarditis, the activation of Th1 and Th17 cells against cardiac self–antigens leads to the development of giant cell myocarditis, a rapidly progressing subtype of the disease characterized by the presence of multinucleated giant cells [34, 70]. Th17 cells and IL–17 have been shown to be dispensable in the early stages of inflammation, but are the main drivers towards dilated cardiomyopathy [10, 15, 163].

Th2 cells excrete a variety of cytokines including IL–4, which stimulates the differentiation of more Th2 cells, inhibits differentiation towards the Th1 phenotype and, together with IL–9, stimulates the differentiation and activation of B cells to plasma cells, mediating the antibody response. IL–9 is also involved in the activation of eosinophils, together with IL–5. In addition to these two main functions (stimulating B cells and eosinophils), Th2–secreted cytokines stimulate the differentiation of

macrophages towards the M2 phenotype, activate other PMNs and inhibit Th1 cells [84, 164, 172]. Cardiac antigen–specific Th2 cells, through recruitment of eosinophils, drive the development of eosinophilic myocarditis, one of the most severe forms of the disease [10, 105, 106].

In addition to the heart–specific Th cells whose roles are discussed above, there is a separate function for heart non–specific CD4+ effector T cells, such as Th cells. These cells accumulate throughout disease progression and become the major infiltrating subset at a later stage, after the peak of inflammation has passed. Their function lies in protecting the heart from post–inflammatory fibrotic remodelling, thus decreasing the risk of severe negative outcomes such as dilated cardiomyopathy and cardiac dysfunction, which are further discussed in Section 2.3 [169].

Function: T regulatory cells Treg cells are anti–inflammatory cells, performing key regulatory functions. Whatever the causative factor, it is critical that the immune response is properly regulated as excessive activation of cytotoxic cells could lead to further tissue damage through, for example, leakage of toxic compounds from immune cells [122, 133, 165]. Treg cells inhibit T cells through direct T cell–T cell interactions, through soluble factors and through APCs [131, 133]. There are two types of Treg cells: natural Treg cells (nTreg), which develop as Treg cells in the thymus as a separate lineage to other T cells, and iTreg cells, which differentiate as described above. nTreg cells are generally specific to self–antigens, whereas iTreg cells mediate the immune response against non–self–antigens [2, 84]. Treg cells inhibit effector T cells, and DCs in a number of ways. The CTLA–4 mediated inhibition of T cell activation has been discussed above. Furthermore, Treg cells interfere with T cell development during proliferation and differentiation, and inhibit the secretion of cytokines by fully developed cells [131, 141]. Soluble mediators like IL–10 and TGFβ directly inhibit effector T cells, primarily by decreasing proliferation and cytokine

release [33, 48, 123]. Other molecules, including fibrinogen–like protein–2, granzyme A/B and adenosine, induce apoptosis in effector T cells, affect their secretion profiles and/or inhibit the maturation of DCs [78, 115]. Direct cell–cell contact between Treg cells and Th cells or DCs induces apoptosis in the former and reduced antigen presentation, inhibited maturation and decreased co–stimulation in the latter. Lastly, Treg cells can inhibit Th cells through direct, local competition for growth factors like IL–2, and by competition for space on DCs [115, 125]. The level of Treg cells is inversely correlated with the intensity of inflammation in autoimmune myocarditis [93].

As previously mentioned, Treg cells constitutively express CTLA4 on their cell surface, making them a target of ipilimumab. In addition to increasing T cell activation by providing a steric block preventing CTLA4 from binding CD80 or CD86, it has been proposed that another effect of ipilimumab is the depletion of Treg cells all together [145, 157]. Although there is no definitive evidence for this, the loss of Treg cells would lead to a significant decrease in inhibition of inflammation, increasing the risk of diseases such as autoimmune myocarditis.

T cells, PMNs, macrophages, APCs and fibroblasts thus play their own specific roles in the adaptive immune response. Many details have been left out here for the sake of providing an overview, but even so the complexity of the system is apparent. To better understand the pathology of myocarditis, it is necessary to determine which cell types, cytokines and interactions, soluble or contact–mediated, are most important. Mathematical modelling can help with this. If a mathematical model can reproduce the pathology of autoimmune myocarditis observed *in vivo* and *in vitro*, this suggests that the mechanisms included in the model could potentially play an important role in the development of this disease. This insight can then direct further research to better understand which cell types are involved in the immunological mechanisms underlying autoimmune myocarditis.

As a first step in the development of a mathematical model, we abstract the immunology into a diagram showing the main cell types and interactions discussed in Section 2.2 (see Fig. 2.3). Based on the literature reviewed above, the pathways included in this graph all play a significant role in the development and progression of autoimmune myocarditis and should thus be included, either explicitly or implicitly, in a mathematical model describing this disease.

2.3 Disease outcomes

Myocarditis has been described as a three–phase disease. The first phase is myocardial injury due to infection or other causes. The second phase consists of autoimmune–mediated damage [90]. At this point myocarditis can resolve spontaneously, but in susceptible individuals it can progress to worse conditions [80]. This is phase three: dilated cardiomyopathy [90]. Approximately 9–16% of myocarditis patients progress to this phase, which can lead to heart failure and ultimately death [163]. Although the phases are described as a chronological progression, there is likely overlap between them and inflammation can be ongoing after dilated cardiomyopathy has occurred [90]. The median time between the first dose of ICIs and the onset of myocarditis is approximately 30 days, although some patients present as quickly as three days after their first dose [108].

2.4 Mathematical modelling

Although no model of autoimmune myocarditis exists, many other diseases and aspects of the immune system have been modelled extensively. These existing models provide a starting point for the modelling of autoimmune myocarditis, but none suffice on their own to describe the specific mechanisms of this disease. When modelling a new system for the first time, it is important to start simple to gain a good un-

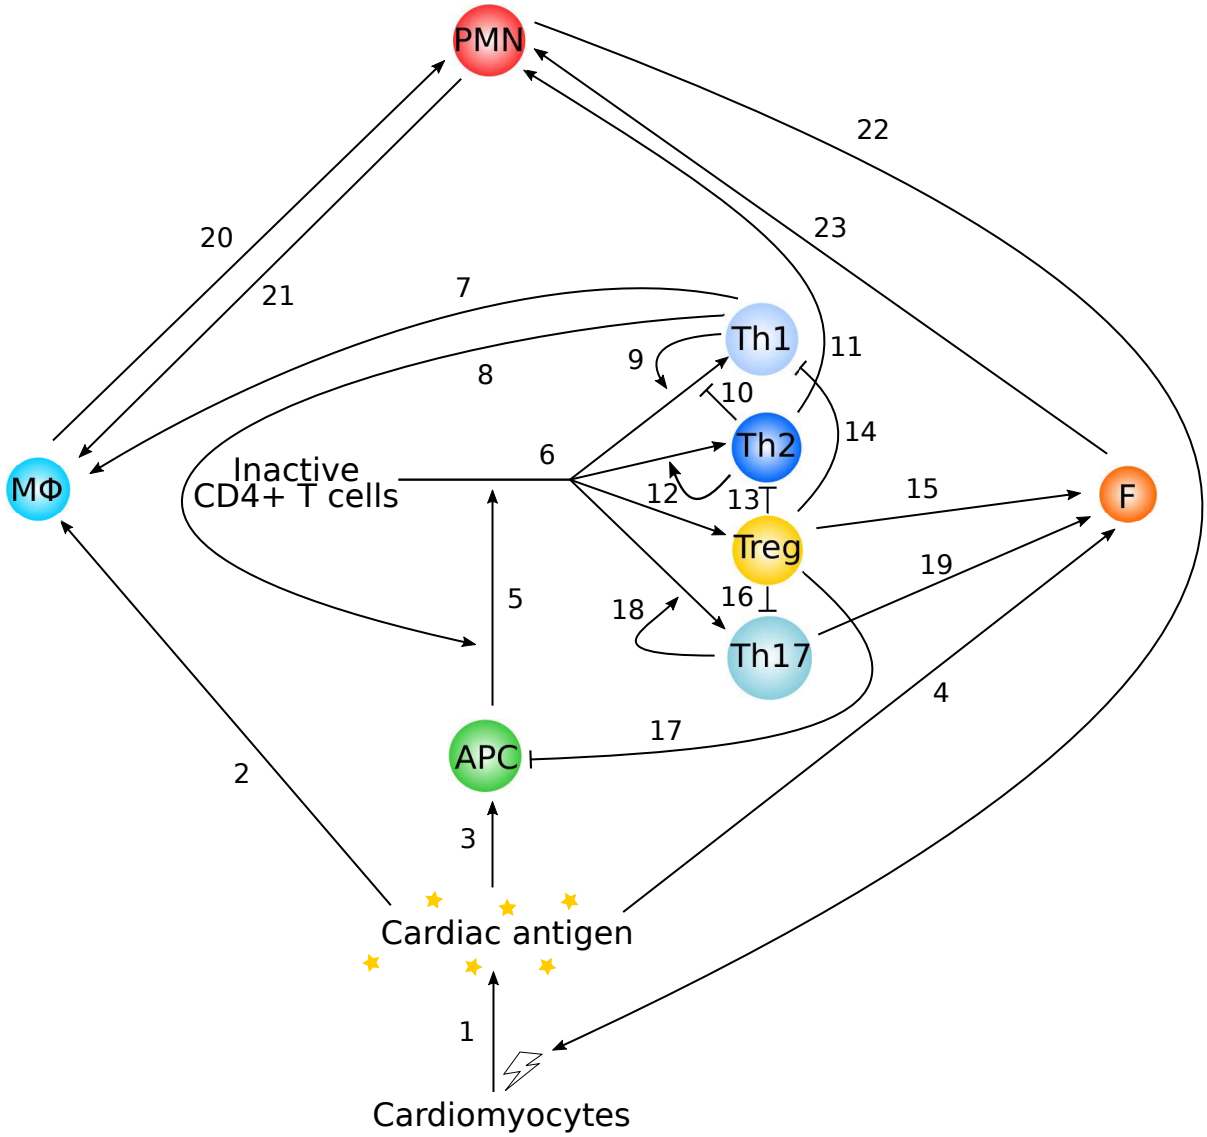


Figure 2.3: Overview of the immunology behind autoimmune myocarditis as described in Section 2.2. Arrow–heads indicate stimulation, bar–heads indicate inhibition. Interaction 1 is supported in Section 2.2.1, interactions 2 and 20 are supported in Section 2.2.2.1, interactions 3 and 5 are supported in Section 2.2.3.1, interactions 4 and 23 are supported in Section 2.2.2.3, interaction 6 is supported in Sections 2.2.3.2 and 2.2.3.2, interactions 7 – 17 are supported in Section 2.2.3.2, interactions 18 and 19 are supported in Section 2.2.3.2, and interactions 21 and 22 are supported in Section 2.2.2.2. The omission of CD8+ T cells, despite their presence in cardiac infiltrate, is supported by the fact that CD4+ T cells are thought to be the primary drivers of disease (see Section 2.2.3).

derstanding of the basics of the system. This simple model can then be extended to include more biological detail, or consider spatial effects, amongst other possibilities. However, given the complexity of autoimmune myocarditis, and the fact that no model of it has yet been developed, we do not discuss these more complex models here, but refer the reader to the relevant references [28, 45, 46, 89].

Examples of simple mathematical models of (parts of) the immune system include ordinary differential equation (ODE)–based models of a generalized immune response that can be used to describe the autoimmune cycle in an analytically tractable way [142], but also more complex ODE models which zoom in on particular immunological processes, such as clonal expansion of T cells and self–tolerance [22]. More detail quickly leads to more complex mathematics, as seen in models of T–cell mediated autoimmunity consisting of both ODEs and integro–differential equations and PDE models describing the minimally required set of immune cells to simulate the immune system [35, 144].

These more detailed models lack the mathematical simplicity that is desirable in a first model of autoimmune myocarditis, and none describe the immunology that is specific to this disease. This is the gap that our work addresses. Below we further discuss the relevance and shortcomings of the above–mentioned models. Where feasible we include the mathematical definition of the models and the interested reader is referred to the relevant papers for further mathematical detail. Although we are interested in the modelling of cancer therapy, specifically ICIs, we omit here a discussion of models of tumour–immune–ICI interactions since our primary focus is on modelling autoimmune myocarditis and the effects of ICIs on the development and progression of this disease.

2.4.1 A mathematical model of antigen presentation dynamics

Sontag investigated the role of antigen presentation dynamics on the initiation of the immune response [142]. They developed an ODE model to describe the interactions between an immune challenge (u), such as a tumour or pathogen, an effector species (y), which can represent pro–inflammatory immune cells, such as T helper cells or polymorphonuclear leukocytes, and a regulatory species (x), which can be interpreted biologically as anti–inflammatory cells. In this model, the immune challenge stimulates both the regulatory and effector species. The regulatory species inhibits the effector species and the effector species stimulates its own proliferation and inhibits the immune challenge (see Figure 2.4). The three state variables evolve over time according to

$$\begin{aligned}
\frac{\mathrm{d}u}{\mathrm{d}t} &= (\lambda - \kappa y)u, \\
\frac{\mathrm{d}x}{\mathrm{d}t} &= -\delta_x x + \beta u, \\
\frac{\mathrm{d}y}{\mathrm{d}t} &= h\left(\frac{u}{x}\right) + f(y),
\end{aligned}$$

where λ, κ, δ_x and β are positive parameters, $h(u/x)$ is a continuous, strictly increasing function with $h(0) = 0$, and $f(y)$ is a nonlinear function which ensures bistability in the system.

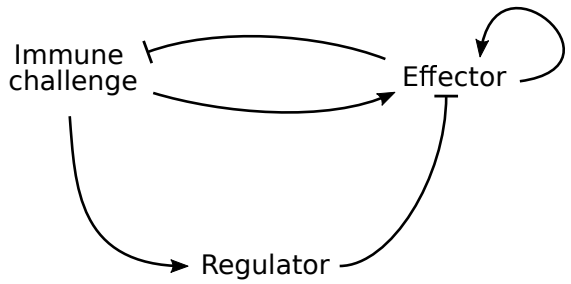


Figure 2.4: Graphical representation of the model developed by Sontag [142]. Arrow–heads indicate a stimulating relationship, bar–heads indicate inhibition.

Sontag showed that antigen dynamics govern the strength of the immune response. They also show, through a series of simplifications, that for certain choices of h and f their model recapitulates the "two–zone tolerance" that is clinically observed in cancer. In "two–zone tolerance" there are two ranges of per capita growth rate of tumour cells in which the immune system cannot shrink the tumour, and the tumour either grows or remains stable. A growth rate outside these ranges leads to elimination of the tumour. Although the motivation for the model is the study of cancer progression, Sontag's model can be applied to other immunological settings as the tumour can easily be replaced by any other pathogen or self–antigen without invalidating the model assumptions. While Sontag's model is relevant as a high–level description of an autoimmune reaction, it lacks the immunological detail needed to apply it to a specific disease. If the aim is to determine the minimally required set of immune cell types and processes for the development and progression of autoimmune myocarditis, specific cell types need to be included as opposed to generic effector or regulatory species.

2.4.2 A mathematical model of self–tolerance

Carneiro et al. consider the clonal expansion of T cells, with a focus on immunological self–tolerance [22]. They present two models, adapted from previously published work, to investigate two hypotheses: (i) self–tolerance is maintained through activation thresholds in T cells (Fig. 2.5 (A)); and (ii) specific regulatory T cells maintain self–tolerance through inhibition of proliferation of self–reactive effector T cells (Fig. 2.5 (B)). Both of the models are relatively simple in that they consist of only a few ODEs each, but the functional forms within them are complicated, making the system analytically intractable. The first model defines the evolution of T cell density (T) over time according to

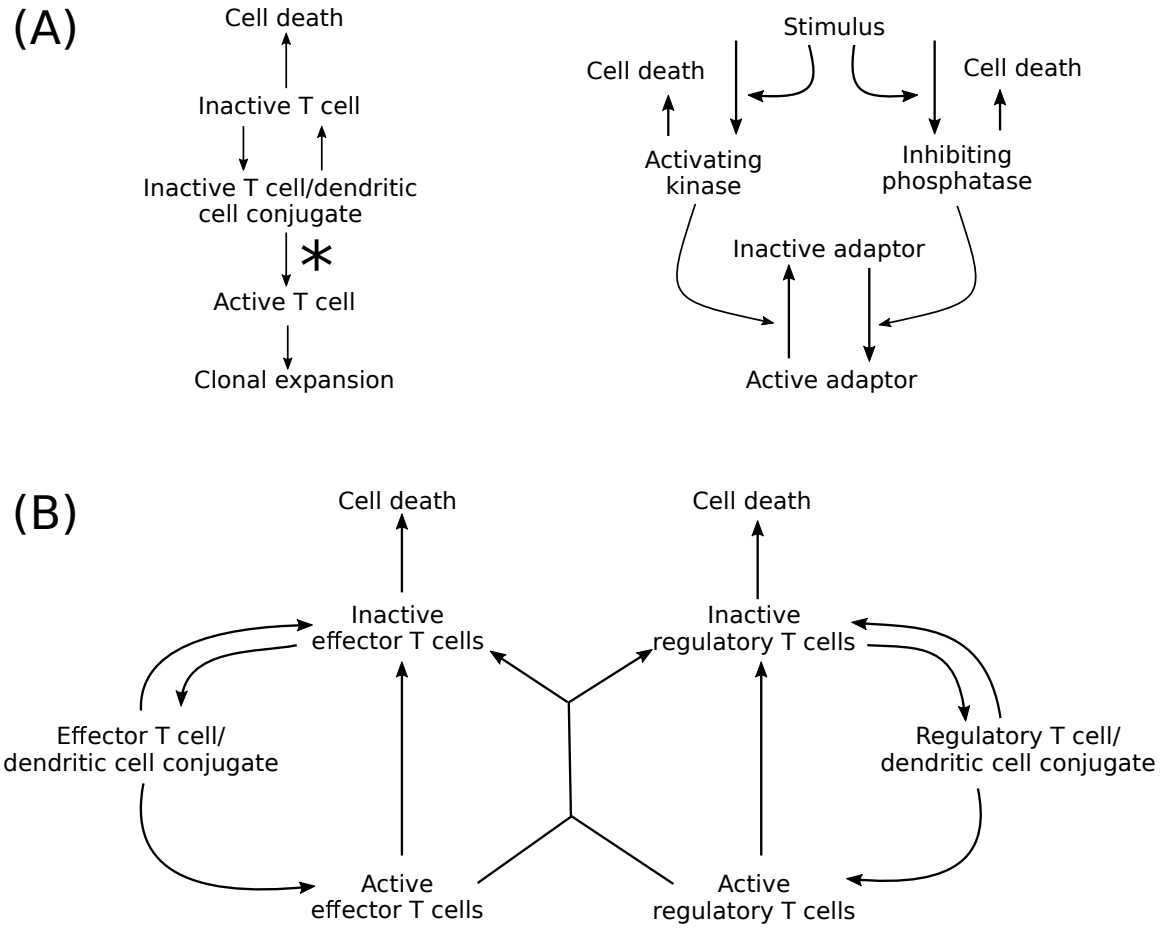


Figure 2.5: (A) graphical representation of model (i) by Carneiro et al. [22]. The molecular dynamics on the right, specifically the ratio of kinase to phosphatase, influence the rate at which T cells leave the conjugate in an activated state, signified by the asterisks on the left. (B) graphical representation of model (ii) by Carneiro et al. [22].

$$\frac{\mathrm{d}T}{\mathrm{d}t} = d\alpha C - \delta(T - C),$$

where d and δ are positive parameters, α is the probability of productive conjugation, and C is the quasi-steady state conjugate density defined as

$$C \cong \frac{c(T + A) + d - \sqrt{-4ATc^2 + (c(T + A) + d)^2}}{2c},$$

where c and A are positive parameters. These two equations capture the dynamics represented graphically in the left-hand graph in Fig. 2.5 (A). The probability α depends on the dynamics depicted in the right-hand graph in Fig. 2.5 (A), and is defined mathematically as

$$\alpha = \int_{P_0}^{K_0(1+\sigma)} \rho_c \, \mathrm{d}P,$$

$$\rho_c = \begin{cases} N|P_0(1+\sigma) - P|^{(d/r_p)-1}|P_0 - P|^{c_E/r_p}, & \text{if } P_0 \leq P \leq P_0(1+\sigma) \\ 0, & \text{otherwise} \end{cases}$$

$$\frac{\mathrm{d}P}{\mathrm{d}t} = r_P(P_0(1+\sigma) - P),$$

where P is phosphatase activity, ρ_c is the phosphatase probability density function in conjugate cells, P_0, K_0, r_p, c_E and N are positive constants and σ is the magnitude of the stimulus which is zero if the cell is free and σ if the cell is conjugated.

The second model, shown in Fig. 2.5 (B), describes the evolution of Treg cells (R) and effector cells (E) according to

$$\begin{aligned}
\frac{\mathrm{d}R}{\mathrm{d}t} &= sE_AR_A - \delta R, \\
\frac{\mathrm{d}R}{\mathrm{d}t} &= pE_A - \delta E, \\
R_A &= \frac{\alpha d R_C}{r + sE_A}, \\
E_A &= \frac{r(p+r) + ds\alpha(R_C - E_C) - \sqrt{4d(p+r)rs\alpha E_C + \left(-r(p+r) - ds\alpha(R_C - E_C)\right)^2}}{2(p+r)s}, \\
R_C &= \frac{RA}{\frac{d}{c} + R + E}, \\
E_C &= \frac{EA}{\frac{d}{c} + R + E},
\end{aligned}$$

where s, δ, p, f, α, A, c, r and d are positive constants.

Through phase plane and bifurcation analyses, Carneiro et al. find that their model for hypothesis (i) only replicates part of the experimental literature, while their model for hypothesis (ii) does not have such limitations. Their main conclusion is that cross–talk between regulatory and effector T cells is sufficient to maintain self–tolerance. They note, however, that mechanisms (i) and (ii) are not mutually exclusive and both have been observed in experiments. Because the inhibition of self–tolerance is a critical part of the development of autoimmune myocarditis as a side–effect of ICIs [60], a model such as this is an excellent starting point for modelling autoimmune myocarditis as one of the most important processes is already well represented mathematically. However, the T cell activation threshold and the role of regulatory T cells in self–tolerance are only two components of the wealth of interactions within the immune system that lead to autoimmune myocarditis. To keep the model of autoimmune myocarditis tractable, each component model needs to be as simple as possible while remaining biologically realistic, and simplifying this model of self–tolerance for the purpose of inclusion in a model for autoimmune myocarditis will thus be essential.

2.4.3 A mathematical model of innate and adaptive immunity

Although simple relative to the complexity of the full immune system, the still somewhat complicated PDE model developed by Su et al. describes the temporal and spatial interactions of various immune cell types and chemokines in response to a source of diffusing antigens (see Fig. 2.6 for details of the non–spatial interactions) [144]. The cell types included are active and apoptotic neutrophils, immature and mature DCs, resting and activated macrophages, effector T cells, and regulatory T cells. Generic cytokines and chemokines are used to describe their contact–independent in-

teractions in response to antigens or pathogen–associated molecular patterns. Based on this model, the authors aim to identify a minimally required set of immune system components to describe the immune response. They simulate the model leaving out either regulatory T cells or dendritic cells and, based on the resulting model dynamics, conclude that all cell types included in their model belong to this minimally required set.

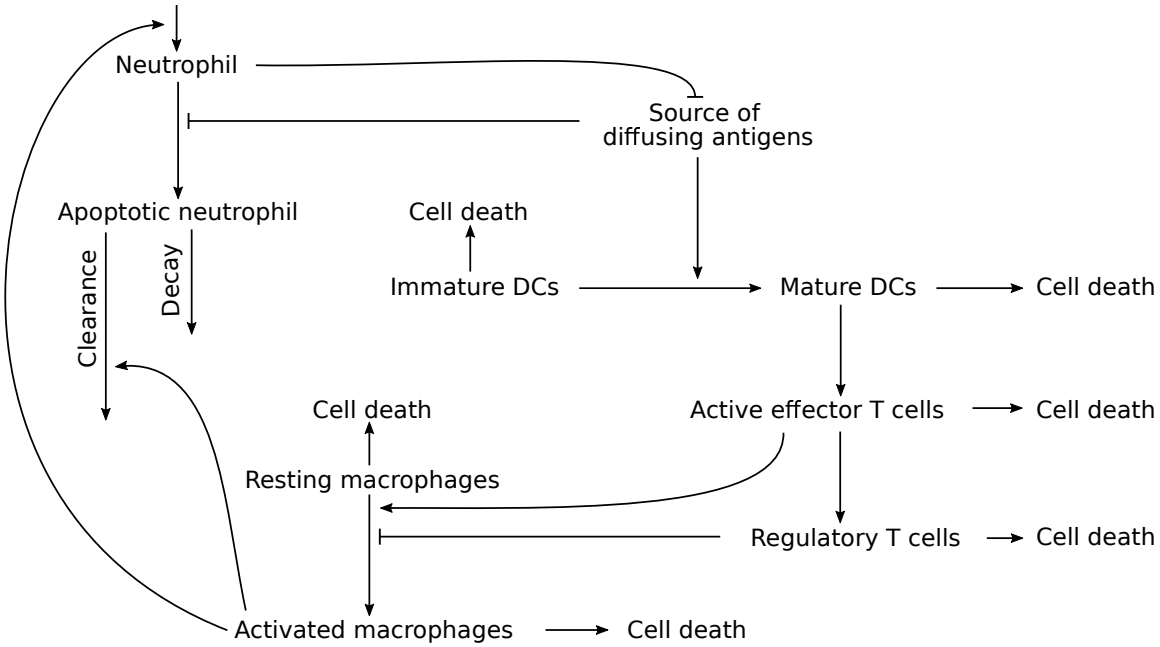


Figure 2.6: Graphical representation of the model developed by Su et al. [144]. Arrow heads indicate stimulation, bar heads indicate inhibition.

The model by Su et al. is relevant to modelling autoimmune myocarditis because it allows for both the resolution of inflammation as well as persistence of the immune response. This behaviour would be expected of a mathematical model of autoimmune myocarditis as both of these scenarios are known outcomes of the disease. However, some of the model assumptions would need to be adapted for the model to be specific for autoimmune myocarditis. First and foremost, the complexity of the model makes it an undesirable approach for a first model of autoimmune myocarditis. With migration of different cell types between the lymph tissue and the site of inflammation being significant aspects of the adaptive immune response, the spatial context is important

in inflammation. However, it adds a level of complexity that precludes the use of analytical approaches to provide initial insights into the model behaviour. It would be preferable to use an ODE–based model, which is more amenable to analysis, in the first instance. Furthermore, Su et al.'s model does not include all immune cell types and interactions that are known to play a significant role in autoimmune myocarditis. Thus a revision of the biology included will be necessary to adapt this model to the context of autoimmune myocarditis.

2.4.4 A mathematical model of T–cell mediated autoimmunity

Lastly, we highlight a model that focuses specifically on T–cell mediated autoimmunity. Delitala et al. use a system of ODEs and integro–differential equations to model the dynamics of host cells, foreign cells, antigen–presenting cells, and naïve and activated T cells (see Fig. 2.7) [35]. The model has six numbered state variables which represent host cells (population 1), foreign cells (population 2), APCs not presenting an antigen (population 3), APCs presenting an antigen (population 4), naïve T cells (population 5), and activated T cells (population 6). All populations except for population 3 are structured by a continuous variable $u \in U = [0, 1]$. For populations 1 and 2, u measures the event of antigenic expression, for population 4 it identifies the antigen being presented, for population 5 it represents the cognate antigen and for population 6 it represents the target antigen. The evolution of population 3 (n_3) is modelled as follows:

$$\frac{\mathrm{d}n_3}{\mathrm{d}t} = -n_3 \sum_{k=1}^{2} \int_U \gamma_k(u^*) f_k(t, u^*) \mathrm{d}u^* + \mu_4 n_4^2,$$

where μ_4 is a positive parameter, the functions $\gamma_1(u^*)$ and $\gamma_2(u^*)$ give the probabilities for antigen presentation, and the functions $f_k(t, u^*),\ k = 1, 2, 4, 5, 6$, define the

dynamics of the populations 1, 2, 4, 5, and 6 (see below). For populations 1, 2, 4, 5, and 6, the number density of the population is computed as

$$n_i(t) = \int_U f_i(t,u)\mathrm{d}u,$$

where, writing $f_i(t,u)$ as f_i for simplicity,

$$\begin{aligned}
\frac{\partial f_1}{\partial t} &= \kappa^C f_1 - \mu^C f_1\ (n_1+n_2) - \xi_1 f_1 \int_U e^{-\theta^C(u-u^*)^2} f_6\ \mathrm{d}u^*,\\
\frac{\partial f_2}{\partial t} &= \kappa^C f_2 - \mu^C f_2\ (n_1+n_2) - \xi_2 f_2 \int_U e^{-\theta^C(u-u^*)^2} f_6\ \mathrm{d}u^*,\\
\frac{\partial f_4}{\partial t} &= n_3 \sum_{k=1}^{2} \gamma_k(u) f_k - \mu_4 f_4 n_4,\\
\frac{\partial f_5}{\partial t} &= \kappa_5 f_5 - \mu_5 f_5 n_5 - \gamma_5 f_5 \int_U e^{-\theta^C(u-u^*)^2} f_4\ \mathrm{d}u^*,\\
\frac{\partial f_6}{\partial t} &= f_4 \int_U e^{-\theta^C(u-u^*)^2} \gamma_5(u^*) f_5\ \mathrm{d}u^* + \kappa_6 f_6 - \mu_6 f_6 n_6 + f_6 \sum_{k=1}^{2} \lambda_k \int_U e^{-\theta^C(u-u^*)^2} f_k\ \mathrm{d}u^*,
\end{aligned}$$

where κ^C, μ^C, ξ_1, ξ_2, θ^C, μ_4, μ_5, μ_6, κ_5, κ_6, and λ_k are positive constants, and γ_5 defines the probability of recognition and activation of naïve T cells.

Delitala et al. use their model as a virtual laboratory to test the effects of molecular mimicry, genetic susceptibility to autoimmunity, and somatic mutations in the *Fas* gene that cause over–proliferation of T cells and decreased action against foreign cells. The model by Delitala et al. can reproduce the effects of molecular mimicry, as seen in experimental work. Furthermore, it suggests that a combination of somatic mutations and molecular mimicry can reproduce the trends seen in patients with autoimmune lymphoproliferative syndrome, indicating that somatic mutations can be a significant trigger for autoimmune disease.

A model of autoimmune myocarditis and the role of ICIs in its development could take inspiration from a model such as this as it already includes a significant portion of the processes involved in the adaptive immune response. Things to consider when

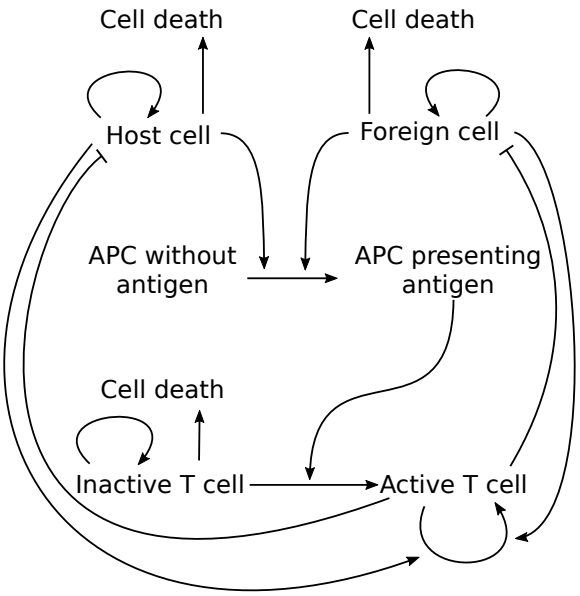


Figure 2.7: Graphical representation of the model developed by Delitala et al. [35]. Arrows pointing out indicate cell death.

adapting this model to describe autoimmune myocarditis are the regulatory roles that T cell phenotypes can play, and the role of the innate immune response in the development and progression of the disease.

Although the mathematical models discussed above each describe a part of the immunological story behind autoimmune myocarditis, none are sufficient to describe the disease on their own. Inspiration from these models, however, will provide a solid starting point for the mathematical modelling of this disease.

2.5 Conclusion

With the increasing survival rate of many types of cancers comes the need to carefully consider the risks and side effects of cancer treatment, especially as these may have health consequences long after the acute phase [66, 114]. Autoimmune myocarditis, like many other side effects of ICIs, has complex and incompletely understood underlying mechanisms, making it difficult to understand which pathways are most critical.

Though it is well established that autoimmune myocarditis is a CD4+ T–cell mediated disease, there are many cell types and processes within this pathway, including the exact effects of ICIs, whose significance in the development and progression of disease are poorly determined [15, 117]. By setting out the biological context, we have collected the information necessary to start modelling autoimmune myocarditis and the effects of ICIs on the development of this disease. With a mathematical model, we can explore and disentangle the immense complexity of the immune system to a set of cell types and interactions that are essential for the development and progression of myocarditis. A minimal model that is able to recapitulate the disease phenotype observed in patients will suggest which processes are likely to be important, help guide and focus experimental efforts, and provide new insights into the mechanisms driving this disease. For example, models will be important for the development of an *in vitro* assay that can test potential new therapeutic avenues for myocarditis–related cardiotoxicity. No such experimental set–up exists to our knowledge, and understanding which cell types are critical and which are of secondary importance will be essential for its development. Furthermore, mathematical models can support experimental work in exploring different treatment regimes through the use of methods such as optimal control theory. Thus, using mathematical modelling to understand autoimmune myocarditis as a side–effect of ICIs can help in the process of making cancer treatments safer and minimizing their long–term effects. In the next chapter, we present the first such model of autoimmune myocarditis.

Chapter 3

A mathematical model of autoimmune myocarditis

The work in this chapter has been published in the article "Mathematical modelling of autoimmune myocarditis and the effects of immune checkpoint inhibitors" by S.A. van der Vegt, L. Polonchuk, K. Wang, S.L. Waters and R.E. Baker in the *Journal of Theoretical Biology*, 537:111002, 2022 [152]. For this article, I conducted the research and wrote the manuscript, which was edited in collaboration with the other authors.

3.1 Introduction

To better direct future research to avenues that are most promising, we need to determine which immunological processes are essential for the development and progression of autoimmune myocarditis, which are of secondary importance, and the role ICIs play in driving the disease. In Chapter 2, we discussed in detail the current understanding of the biology behind the disease, and presented a network representation of the cell types and interactions that underpin the disease (reproduced in Fig. 3.1). Although the biology presented in Fig. 3.1 is a simplification of the myriad of immunological processes underlying autoimmune myocarditis, it is still highly complex and the rela-

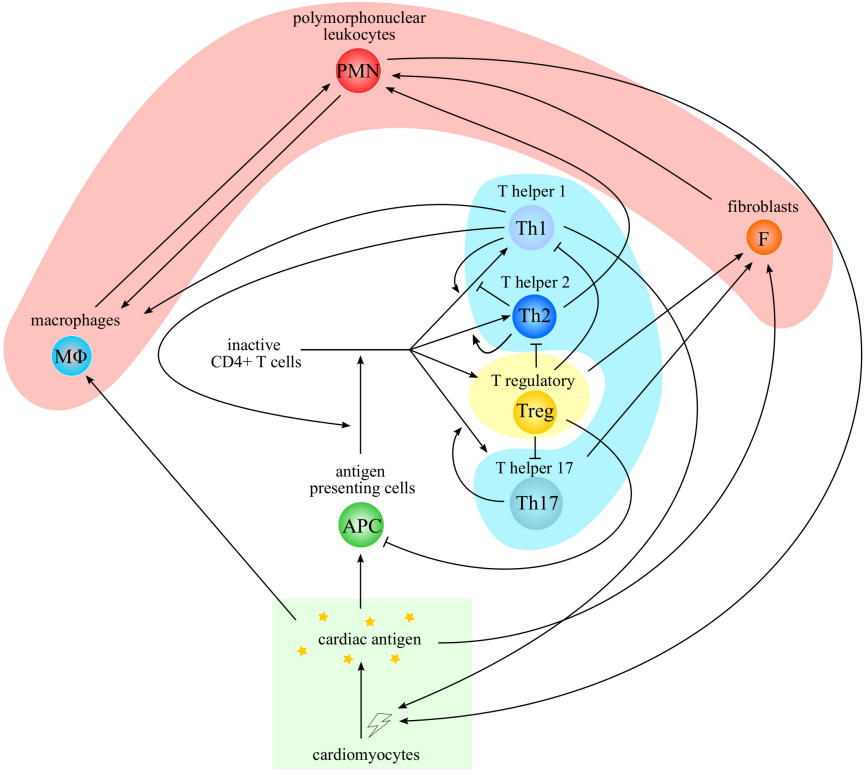


Figure 3.1: An overview of the immunology underpinning the development and progression of autoimmune myocarditis. Arrowheads indicate stimulation, barheads indicate inhibition. Shaded areas represent the lumped variables in the mathematical model.

tive importance of the different pathways is not obvious. To aid in establishing which cell types and interactions present in Fig. 3.1 are of primary importance, in this chapter we develop and analyse a mathematical model of autoimmune myocarditis. Our aim in developing the model is to strike a balance between simplicity, so as to allow analytical methods to be applied in its study, and biological realism, by incorporating cell types determined to be essential by experimental research and ensuring that model behaviour preserves disease characteristics observed in patients.

Having established the biological basis on which to build the model, we next consider what behaviour we expect from the model mathematically. The two potential outcomes of autoimmune myocarditis, resolution and chronic inflammation, are reminiscent of the "three Es" of immunoediting, where the interactions between a tumour

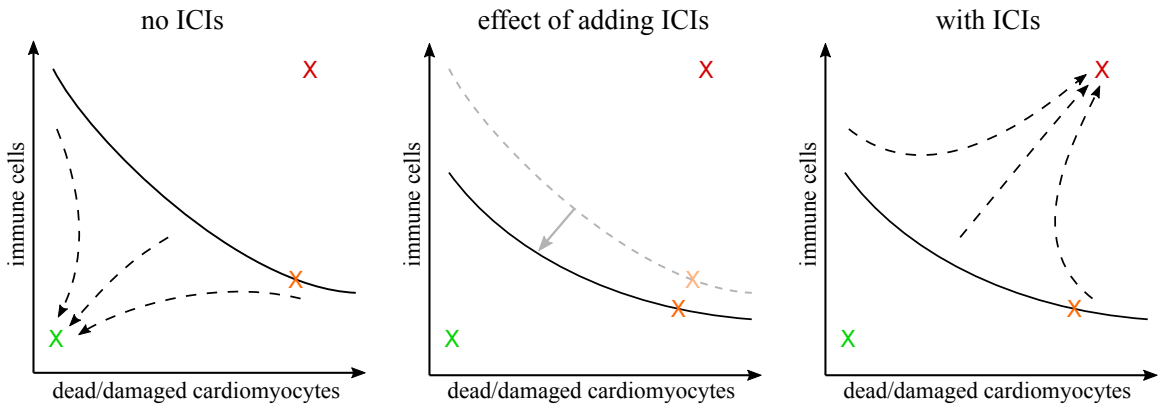


Figure 3.2: Illustration of the steady state and stability concepts related to the hypothesized mechanism of autoimmune myocarditis and the role of ICIs therein. The green cross indicates a healthy stable steady state with low cardiac damage and low numbers of immune cells. The red cross indicates a diseased stable steady state with high cardiac damage and high numbers of immune cells. The orange cross indicates a saddle point, and the black curve indicates the separatrix. The dashed arrows show trajectories, starting at the same initial condition in both the left and right panels. By adding ICIs (middle panel), the separatrix is moved closer to the healthy steady state, lowering the threshold for an immune response to develop into myocarditis. In the middle panel, the ICI–free position of the separatrix and saddle point are indicated by the grey curve and light orange cross, respectively.

and the immune system can lead to tumour elimination, equilibrium or escape [41]. This concept has been captured in mathematical models using systems of ODEs that describe the dynamic interactions between tumour and immune system components and admit multiple steady state solutions. In these systems, each steady state represents one of the possible outcomes of tumour–immune interactions, *i.e.* one of the three Es [38].

It is reasonable to expect a mathematical model of autoimmune myocarditis to have at least two non–negative steady states which are stable to small perturbations, one representing a healthy state, *i.e.* resolution (represented by the green cross in Fig. 3.2), and one representing a diseased state, *i.e.* chronic autoimmune myocarditis, with high numbers of immune cells and dead/damaged cardiomyocytes (the red cross in Fig. 3.2). A third non–negative steady state, a saddle point, lies on the separatrix which separates the basins of attraction of the two stable steady states (the orange

cross in Fig. 3.2). We propose that the separatrix between the basins of attraction of the healthy and diseased steady states can be seen as a tolerance threshold in the immune system: any challenge to the immune system that causes an immune response that stays below this threshold will be tolerated and will not lead to disease, *i.e.* the system stays in the basin of attraction of the healthy steady state (left panel in Fig. 3.2). However, if the challenge to the immune system is sufficiently large to cause an immune response that crosses the threshold, the system can no longer return to the healthy steady state and instead autoimmune myocarditis develops.

The immune system is constantly responding to (perceived) challenges from foreign organisms, causing immune cell numbers to vary over time. It is rare, however, for someone to develop autoimmune myocarditis spontaneously, so the body must be able to control these fluctuations in immune components most of the time. As the body of a healthy individual is able to tolerate small perturbations to the immune system, the separatrix symbolizing the immune threshold for development of autoimmune myocarditis should not lie too close to the healthy steady state.

Within this framework, we hypothesize that the impact of ICIs on the tolerance threshold reflects the increased risk of developing myocarditis when a cancer patient is treated with ICIs [50, 95]. We hypothesize that the presence of ICIs moves the separatrix closer to the healthy steady state, and an immune challenge that could be controlled before treatment with ICIs, now leads to disease. This is shown in the right panel in Fig. 3.2 where the trajectories that approached the healthy steady state before (dashed arrows in the left panel), now lie above the lowered tolerance threshold and thus approach the diseased steady state instead (dashed arrows in the right panel). In this chapter, we test whether a model based on the immunological pathways that are marked as essential in the literature (see Fig. 3.1) can exhibit the behaviour shown in Fig. 3.2, and whether the hypothesized movement of the separatrix in response to ICIs is an emergent behaviour of such a model.

The steady state hypothesis in Fig. 3.2 is shown in two dimensions for ease of presentation, whereas the mathematical model developed in this chapter has four state variables. The general idea of steady state behaviour hypothesized above translates from two to four dimensions, where we can still have two stable steady states and one steady state providing a separatrix. This separatrix will be a multidimensional surface rather than a line but can still function as an immune threshold for autoimmune myocarditis. However, it is possible in four dimensions for the steady state providing the separatrix to show a type of linear stability for which the idea that only sufficiently large perturbations will cause autoimmune myocarditis breaks down. For example, this would be the case if the unstable steady state has a conjugate pair of complex eigenvalues, a real positive eigenvalue and a real negative eigenvalue. We note that in the parameter regimes explored in the remainder of this work, we do not come across this type of linear stability, and the model behaves in four dimensions as expected based on the two-dimensional representation in Fig. 3.2.

3.1.1 Chapter outline

We propose a mathematical model of autoimmune myocarditis based on the known biological pathways underlying the disease, and explore whether such a model can produce the behaviours postulated in Fig. 3.2. The cell types included in this model represent a minimal set required to describe the mechanisms underlying the development and progression of autoimmune myocarditis.

In Section 3.2 we present the model, the steady states and the parameterisation. We start by further simplifying the biology in Fig. 3.1 to obtain a graph of four interacting cell types, which we translate into a system of four ODEs. We derive a set of analytic expressions which together define the nonnegative steady states of the system. To parameterize the model, we select parameter sets that satisfy a set of criteria based on the biological basis and the hypothesized behaviour of the

model. In Section 3.3 we show that the presence of ICIs in the system increases the likelihood of a patient developing autoimmune myocarditis. Through numerical simulations and bifurcation analysis, we show that the effect of ICIs on the system is as predicted above. This is established for nivolumab and ipilimumab monotherapy, as well as combination therapy. In Section 3.4 we explore how the administration of ICIs affects the risk of autoimmune myocarditis for different patients by investigating model behaviour in different parameter regimes. We show that small changes in model parameters can lead to qualitatively and quantitatively different outcomes when ICIs are introduced. Lastly, in Section 3.5 we explore the identifiability of model parameters. We assess the structural and practical identifiability, exploring the limits of what is computationally feasible for the model developed in this chapter. We conclude in Section 3.6 by highlighting the novel contributions to the field, and suggesting potential avenues for future work.

3.2 The mathematical model

Here, we introduce the mathematical model of autoimmune myocarditis and the role of ICIs. Our guiding aim is to capture the salient immunological mechanisms and preserve the characteristics of the disease observed in patients while keeping the model tractable. The relative simplicity of ODE models combined with their ability to capture complex, nonlinear immunological mechanisms, makes them a suitable starting point for modelling this disease.

Fig. 3.1 shows the complexity of the key immunological pathways. However, including all these cell types and interactions would lead to a hugely complex model and preclude a number of analytic approaches that provide insight into the behaviour of the system, such as the derivation of expressions for the steady states and their linear stability. To further simplify the system so as to allow for these types of analysis,

immune cell types that have similar effects on the cycle of autoimmunity are grouped together in the model. This applies for example to Th1, Th2 and Th17 cells, shaded in blue in Fig. 3.1, all of which are activated through similar mechanisms and involved in stimulating innate immune cells and/or damaging cardiomyocytes. Similarly, the innate immune cell types, *i.e.* macrophages, fibroblasts and PMNs shaded in red, are all activated by T cells or directly by damage to the myocardium and ultimately lead to activated PMNs attacking cardiomyocytes, and thus we also group these cell types together. This inevitably leads to the neglect of certain immunological details, such as the interactions that happen within the grouped populations of immune cells. However, our goal is to build a minimal model that includes the immunological pathways required to capture the development of autoimmune myocarditis.

As discussed in Chapter 2, nivolumab is an antibody against PD–1, which is found on the surface of T cells, particularly Th cells [126]. Anergy, a state of hypoactivity, is induced in T cells when PD–1 binds one of its ligands, PD–L1 or PD–L2, which are found on the surfaces of peripheral cells and dendritic cells, respectively. We would thus expect the presence of nivolumab in our model to inhibit deactivation of pathogenic CD4+ T cells as PD–1 can no longer bind PD–L1, so anergy cannot be induced in active T cells by contact with peripheral cells. Furthermore, we expect the presence of nivolumab to increase activation of T cells as PD–1 can no longer bind PD–L2 on dendritic cells, which would inhibit the T cell activation process. Ipilimumab, on the other hand, is an antibody targeting CTLA4, which is present on the cell surface of activated T cells, particularly regulatory T (Treg) cells. Binding of CTLA4 to its target, CD80/CD86 found on dendritic cells, interferes with the activation of T cells [113, 124]. The presence of ipilimumab, which blocks CTLA4, in the model is thus expected to increase activation of T cells. Furthermore, it has been proposed that ipilimumab causes the depletion of Treg cells all together, so we would also expect ipilimumab to increase the deactivation rate of Treg cells in the model

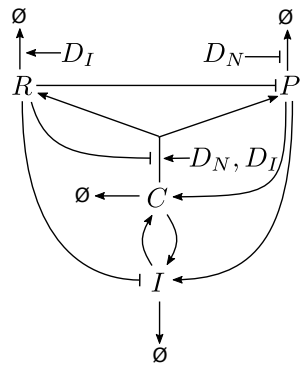


Figure 3.3: Graphical representation of the mathematical model of autoimmune myocarditis, indicating the effects of nivolumab (D_N) and ipilimumab (D_I) in the network of interactions. The cell types included are dead/damaged cardiomyocytes (C), innate immune cells (I), pathogenic CD4+ T cells (P) and regulatory T cells (R). Arrow–heads indicate stimulation, bar–heads indicate inhibition, Ø indicate cell death and/or clearance.

[145, 157].

Our model consists of a system of four coupled nonlinear ODEs, each describing one class of cell types: the number of damaged/dead cardiomyocytes (C), an indicator of the extent of cardiac damage and disease progression; the number of innate immune cells (I); the number of pathogenic CD4+ T cells (P); and the number of regulatory T cells (R). Fig. 3.3 shows a graphical representation of the model. The variable t denotes the time in units of days.

For the remainder of this work we will use the terms "T cells" and "CD4+ T cells" interchangeably, as we define the population of pathogenic T cells in our model to specifically mean CD4+ T cells. We note that CD8+ T cells have been found to infiltrate cardiac muscle tissue during inflammation in monkeys and humans suffering from autoimmune myocarditis [68, 108]. However, studies with mice and observational data obtained from humans show that CD4+ T cells are the primary drivers of autoimmune myocarditis in s [15, 93]. In an effort to keep the mathematical model as simple as possible, we therefore choose to focus here on the pathway driven by CD4+ T cells and leave the immunological mechanisms including CD8+ T cells for future work.

3.2.1 The mathematical model

The dynamics of dead/damaged cardiomyocyte numbers is governed by the ODE

$$\frac{\mathrm{d}C}{\mathrm{d}t} = a_1 \frac{P^2}{R_1 + R} + a_2 \frac{I}{R_2 + R} - d_1 C, \qquad C(0) = C_0, \tag{3.1}$$

where a_1, a_2, R_1, R_2 and d_1 are positive parameters, and C_0 is the nonnegative initial number of damaged cardiomyocytes. The first term on the right–hand side (RHS) of Eq. (3.1) describes the nonlinear effects of damage to the myocardium due to the presence of pathogenic CD4+ T cells. The nonlinearity in P in this term balances the high complexity of the immunological processes driving cardiomyocyte dynamics with the desire to keep the model as simple as possible. Nonlinearity in the model is necessary in order for it to show the steady state behaviour required of our model of autoimmune myocarditis. We acknowledge that all processes included in this model are likely highly nonlinear in nature, however, squaring P in this term is one of the simplest ways this nonlinearity can be included. The denominator $R_1 + R$ represents the inhibition of this process by Treg cells, where R_1 is the number of regulatory immune cells required for 50% inhibition of pathogenic CD4+ T cell activity. The second term on the RHS of Eq. (3.1) describes myocardial damage due to the presence of innate immune cells [15, 67, 104, 156]. This process is inhibited by Treg cells, hence the denominator, where R_2 is the number of regulatory immune cells required for 50% inhibition of innate immune cells. Lastly, the third term on the RHS of Eq. (3.1) describes the rapid clearance of dead cardiomyocytes by *e.g.* macrophages [159, 167].

The number of innate immune cells evolves over time according to the ODE

$$\frac{\mathrm{d}I}{\mathrm{d}t} = a_4 \frac{P}{R_1 + R} + a_5 C - d_3 I, \qquad I(0) = I_0, \tag{3.2}$$

where a_4, a_5 and d_3 are positive parameters, and I_0 is the nonnegative initial num-

ber of innate immune cells. The first term on the RHS of Eq. (3.2) describes the stimulation of innate immune cell activation by pathogenic CD4+ T cells [10, 67]. Again, pathogenic CD4+ T cells are inhibited by Treg cells. The second term on the RHS of Eq. (3.2) describes the recruitment of innate immune cells by dead/damaged cardiomyocytes through the innate immune pathway [75, 118, 133]. The last term describes the death or deactivation of innate immune cells [53, 111].

The dynamics of pathogenic CD4+ T cells is governed by the ODE

$$\frac{\mathrm{d}P}{\mathrm{d}t} = a_6\phi\frac{C}{R_3+R} - d_4P, \qquad P(0) = P_0, \tag{3.3}$$

where a_6, d_4, ϕ and R_3 are positive parameters, and P_0 is the nonnegative initial number of pathogenic T cells. The first term on the RHS of Eq. (3.3) describes the activation of inactive T cells by dead/damaged cardiomyocytes and their subsequent clonal expansion [11, 133]. Because this process is stimulated by both nivolumab and ipilimumab, the value of the rate parameter a_6 increases in the presence of these drugs. A fraction $0 < \phi < 1$ of activated T cells then adopt a pathogenic phenotype [20, 133, 154]. The processes of activation and differentiation are inhibited by Treg cells, hence the denominator $R_3 + R$ where R_3 is the number of regulatory immune cells required for 50% inhibition of pathogenic CD4+ T cell activation [33, 104, 131, 133]. The second term on the RHS of Eq. (3.3) denotes the deactivation or death of pathogenic CD4+ T cells [15, 17, 54, 112, 126, 170]. As discussed above, nivolumab prevents T cell inhibition in the myocardium, thus the value of the rate parameter d_4 decreases as a consequence of the presence of nivolumab.

The number of Treg cells evolves according to the ODE

$$\frac{\mathrm{d}R}{\mathrm{d}t} = a_6(1-\phi)\frac{C}{R_3+R} - d_5R, \qquad R(0) = R_0, \tag{3.4}$$

where a_6, d_5, ϕ and R_3 are positive parameters, and R_0 is the nonnegative initial

number of Treg cells. The activation and differentiation of Treg cells occurs according to the same processes as described for pathogenic CD4+ T cells, the only difference being that a fraction $(1 - \phi)$ of the clonally expanded T cells adopt a regulatory phenotype. This is reflected in the first term on the RHS of Eq. (3.4). The second term on the RHS of Eq. (3.4) describes the death or deactivation of Treg cells. Ipilimumab depletes Treg cells, thus the value of the rate parameter d_5 increases as a consequence of the presence of ipilimumab.

3.2.2 Steady states and linear stability

The model described in Section 3.2.1 was built based on the immunology underlying the progression of autoimmune myocarditis. To assess whether this model can capture the hypothesized behaviour discussed in Section 3.1 and shown in Fig. 3.2, we now seek an analytical expression for the steady states of the system. A model of autoimmune myocarditis should admit three non–negative steady states: one stable steady state with low counts of inflammatory cells and dead/damaged cardiomyocytes; one stable steady state with high counts of inflammatory cells and dead/damaged cardiomyocytes; and one saddle point with intermediate cell counts.

The origin is a steady state of the system (3.1)–(3.4). To determine the non–zero steady states of the model, we set the RHSs of Eq. (3.1)–(3.4) equal to zero, use algebraic manipulation to eliminate C, I and R, and establish a polynomial equation satisfied by the non–zero steady states in terms of P. The non–zero steady state

numbers of pathogenic CD4+ T cells, which we denote P_*, satisfy the cubic equation

$$\begin{aligned} 0 \; &= d_3 K_R^3 P_*^3 \\ &+ \left(d_3(R_1+R_2+R_3)K_R^2 - \frac{a_2 a_5}{d_1} K_R^2 - \frac{a_1 K_P}{d_1} d_3 K_R \right) P_*^2 \\ &+ \Bigg[d_3(R_1R_2+R_1R_3+R_2R_3)K_R - \frac{a_2 a_5}{d_1}(R_1+R_3)K_R \\ &\qquad - \frac{a_1 K_P}{d_1}\left(R_2 d_3 - \frac{a_2 a_5}{d_1} \right) - \frac{a_1 a_2 a_5 K_P}{d_1^2} \Bigg] P_* \\ &+ d_3 R_1 R_2 R_3 - \frac{a_2 a_5}{d_1} R_1 R_3 - \frac{a_2 a_4 K_P}{d_1}, \end{aligned} \tag{3.5}$$

where $K_R = d_4(1-\phi)/(d_5\phi)$ and $K_P = a_6\phi/d_4$. The steady state cell counts of the remaining variables are given in terms of P_* by

$$R_* = K_R P_*, \tag{3.6}$$

$$I_* = \frac{\left(a_4 + \frac{a_1 a_5}{d_1} P_* \right)\left(\frac{P_*}{K_R P_* + R_1} \right)}{d_3 - \frac{a_2 a_5}{d_1(K_R P_* + R_2)}}, \tag{3.7}$$

and

$$C_* = \frac{a_1 P_*^2}{d_1(R_1+R_*)} + \frac{a_2 I_*}{d_1(R_2+R_*)}. \tag{3.8}$$

The linear stability of the non–negative steady states with respect to small perturbations is determined by the eigenvalues λ_i, $i = 1, \ldots, 4$, computed by solving

$$|J - \lambda_i I| = 0, \tag{3.9}$$

where I is the four–by–four identity matrix and J is the the Jacobian matrix:

$$J = \begin{pmatrix} -d_1 & \frac{a_2}{R_2+R_*} & \frac{2a_1P_*}{R_1+R_*} & -\frac{a_1P_*^2}{(R_1+R_*)^2} - \frac{a_2I_*}{(R_2+R_*)^2} \\ a_5 & -d_3 & \frac{a_4}{R_1+R_*} & -\frac{a_4P_*}{(R_1+R_*)^2} \\ \frac{a_6\phi}{R_3+R_*} & 0 & -d_4 & \frac{-a_6\phi C_*}{(R_3+R_*)^2} \\ \frac{a_6(1-\phi)}{R_3+R_*} & 0 & 0 & \frac{-a_6(1-\phi)C_*}{(R_3+R_*)^2} - d_5 \end{pmatrix}. \tag{3.10}$$

If $\mathrm{Re}(\lambda_i) < 0$ for all i, the steady state is linearly stable. If $\mathrm{Re}(\lambda_i) > 0$ for all i, the steady state is linearly unstable. If a subset of $\mathrm{Re}(\lambda_i)$ are positive and the remaining $\mathrm{Re}(\lambda_i)$ are negative, the steady state is a saddle point.

3.2.3 Parameter selection

Our overarching aim is to develop a model capable of recapitulating the increased risk of developing autoimmune myocarditis upon application of ICIs. The mathematical model has 13 parameters (see Table 3.1 for definitions of the parameters). The grouping of different types of immune cells into classes, and the consolidation of many steps in the immunological process into single terms in the model, means the model is simple enough to be amenable to some analytical approaches, but it makes it difficult to estimate most parameter values directly from the literature. As a result, we seek parameter values that give rise to the hypothesized number and types of steady states required for the model behaviour illustrated in Fig. 3.2, in an ICI–free context. The values of parameters d_4 and d_5 can be obtained from literature [121]. For the remaining 11 parameters, we find suitable ranges based on a combination of insights gained from the literature and constraints that give rise to the behaviour depicted in Fig. 3.2. Recall, we propose that a model of autoimmune myocarditis should have two non–negative, stable steady states and one saddle point, whose position in phase space is influenced by the presence of ICIs. In this section, we explore whether, and for what parameter values, the model can display these characteristics.

Parameter	Definition	Value	Units
a_1	Rate at which pathogenic CD4+ T cells damage/kill cardiomyocytes	80	day^{-1}
a_2	Rate at which innate immune cells damage/kill cardiomyocytes	6000	cells day^{-1}
a_4	Rate at which pathogenic CD4+ T cells activate innate immune cells	500,000	cells day^{-1}
a_5	Rate at which damaged/dead cardiomyocytes activate innate immune cells	35	day^{-1}
a_6	Rate at which damaged/dead cardiomyocytes activate pathogenic CD4+ T cells	750	cells day^{-1}
d_1	Clearance rate of damaged/dead cardiomyocytes	75	day^{-1}
d_3	Innate immune cell death rate	7	day^{-1}
d_4	Pathogenic CD4+ T cell death rate, from [121]	0.1	day^{-1}
d_5	Regulatory T cell death rate, from [121]	0.1	day^{-1}
R_1	Number of regulatory immune cells for 50% inhibition of pathogenic CD4+ T cell activity	1100	cells
R_2	Number of regulatory immune cells for 50% inhibition of innate immune cells	40000	cells
R_3	Number of regulatory immune cells for 50% inhibition of pathogenic CD4+ T cell activation	2000	cells
ϕ	Fraction activated CD4+ T cells that take on a pathogenic phenotype	0.7	unitless

Table 3.1: Definitions, values and units of the 13 model parameters.

We seek parameter sets that satisfy the following criteria (summarized in Table 3.2):

1. The parameter set must produce three nonnegative steady states. Two must be stable to small perturbations, and one must be a saddle point so that the behaviour of the model captures the ideas in Fig. 3.2. The origin corresponds to the "healthy" state, with no inflammatory cells. This steady state must be stable. At the origin where $[C_*,\ I_*,\ P_*,\ R_*] = [0,\ 0,\ 0,\ 0]$, the Jacobian matrix is

$$J(0,0,0,0) = \begin{pmatrix} -d_1 & \frac{a_2}{R_2} & 0 & 0 \\ a_5 & -d_3 & \frac{a_4}{R_1} & 0 \\ \frac{a_6\phi}{R_3} & 0 & -d_4 & 0 \\ \frac{a_6(1-\phi)}{R_3} & 0 & 0 & -d_5 \end{pmatrix}, \tag{3.11}$$

   from which we derive the following condition for stability which ensures that the real parts of all the eigenvalues are negative:

$$d_1 d_3 d_4 R_1 R_2 R_3 > a_2 a_5 d_4 R_1 R_3 + a_2 a_4 a_6 \phi. \tag{3.12}$$

2. The saddle point cell counts of damaged cardiomyocytes, innate immune cells and pathogenic CD4+ T cells should lie between 10^2 and 10^4 cells, to ensure that a moderately sized perturbation in CD4+ T cells is needed to cause autoimmune myocarditis to develop. It is unknown how many immune cells or dead/damaged cardiomyocytes are needed to trigger autoimmune myocarditis, so we fix the counts of damaged cardiomyocytes, innate immune cells and pathogenic CD4+ T cells within these ranges to ensure that model behaviour is consistent with the expected behaviour of the immunological system in terms of requiring a significant perturbation to cause autoimmune myocarditis to develop.

3. The maximum number of damaged cardiomyocytes at any time point must be

below 10^6. The heart consists of approximately 2–3 $\times$ 10^9 cardiomyocytes in total. To be biologically realistic, the model should not predict the entire heart to be destroyed within a matter of days [148].

4. The number of T cells should peak around seven days after an immune challenge [86]. We use the initial conditions $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 10^4,\ 0]$ cells to simulate the model when it has been perturbed away from the healthy steady state by a strong immune challenge. A perturbation of this size will cause myocarditis to develop for all parameter sets that satisfy Criterion 2 as it will push the system beyond the separatrix, and can thus be reliably used to assess the time it takes a T cell response to develop for all parameter sets.

5. All terms in the model must have a significant impact on model outcomes. The model was constructed to include those interactions which, based on the existing literature, are believed to be important for the development of autoimmune myocarditis, so all these interactions should have a significant impact on the dynamics.

We obtain estimates for d_4 and d_5 from Robertson–Tessi *et al.* [121], setting $d_4 = d_5 = 0.1$ day $^{-1}$. We implement a comprehensive parameter sweep to establish reasonable values for the 11 remaining parameters. All algorithms, here and in the following sections, are implemented and all simulations run in MATLAB R2020b (published by Mathworks, Natick, MA). Unless noted otherwise, we use the ode23s solver with default tolerances. A stiff solver is required because of the rapid changes in cell counts over time, and we note that the ODE solutions converge for the default tolerances.

To find parameter sets that satisfy Criteria 1 and 2, we solve Eq. (3.5) for each possible parameter set. By Descartes Rule of Signs, the number of sign changes in the coefficients of the polynomial in Eq. (3.5) gives the maximum number of positive

Requirement	Reference
1. Three nonnegative steady states, two stable and one saddle.	Fig. 3.2
2. Number of damaged/dead cardiomyocytes, pathogenic CD4+ T cells and Treg cells at the saddle point must be between 10^2 and 10^4.	[102, 132]
3. Number of damaged/dead cardiomyocytes should never exceed 10^6.	[148]
4. Numbers of T cells peak around seven days when the initial condition is $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 10^4,\ 0]$ cells.	[86]
5. All terms must have significant impact on model output.	Chapter 2
6. $d_4 = d_5 = 0.1$ day^{-1}.	[121]

Table 3.2: Requirements for establishing parameter values.

solutions for P_*. Equation (3.5) can be solved using the roots() function in MATLAB to obtain the nonzero steady state cell counts of P_*. After solving for the steady state cell counts of the other variables, we then select those parameter sets which give rise to steady states with suitable positive cell counts for all variables and the desired linear stability.

To further narrow down the number of feasible parameter sets to those satisfying Criteria 3 and 4, we simulate the model for all the remaining parameter sets using the initial condition $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 10^4,\ 0]$ cells. The healthy steady state is at the origin. Because active, cardiac–specific pathogenic CD4+ T cells are a known (partial) trigger of autoimmune myocarditis, setting an initial condition of $P_0 = 10^4$ is employed as a means to perturb the system away from a healthy state [52, 80]. From these simulations, we extract the maximum number of damaged/dead cardiomyocytes in the system at any point and discard those parameter sets for which this cell count is too high. Next, we select parameter sets for which the number of T cells in the system peaks as close to seven days as possible. The remaining parameter sets thus satisfy Criteria 3 and 4. When simulating the model with each of the remaining parameter sets, we find that the simulation results are qualitatively similar. We then select one set of parameter values from the many that satisfy the criteria listed above

to make progress on modelling autoimmune myocarditis. The parameter values in this "default parameter" set are listed in Table 3.1.

To confirm that the default parameter set satisfies Criterion 5, we apply a parameter sensitivity analysis on an ICI–free model using the implementation of Latin Hypercube Sampling Partial Rank Correlation Coefficent (LHS–PRCC) made available by Marino *et al.* [87]. This metric evaluates how much a particular output of the model is affected by an increase or decrease in a certain parameter value, and whether or not this impact is significant compared to a dummy parameter. We choose this method because the computational costs are feasible, as compared to other methods, such as the extended Fourier Amplitude Sensitivity Test, for such high–dimensional parameter spaces.

We take the initial conditions for the parameter sensitivity analysis to be $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 10^4,\ 0]$ cells, so the system is perturbed away from the healthy steady state and towards autoimmune myocarditis, and the effect of the parameters on the development of the disease can be assessed. The impact that the parameters have on each of the variables is measured after 60 days, a relevant timescale for the development and pathology of autoimmune myocarditis, as discussed in Chapter 2 Section 2.3. We vary each of the parameter values over two orders of magnitude around their default values. The only exception to this is ϕ which is constrained between zero and one. A new parameter set is generated for each run, where the parameter values are independently varied around their default value. Approximately 6,000 runs are required for estimates of the correlation between the parameter values and model outcome to converge for all parameters. Results in Fig. 3.4 show that the values of the variables C, I, P and R at the final time point are sensitive to changes in the value of any parameter in the model. This parameter set thus satisfies Criterion 5.

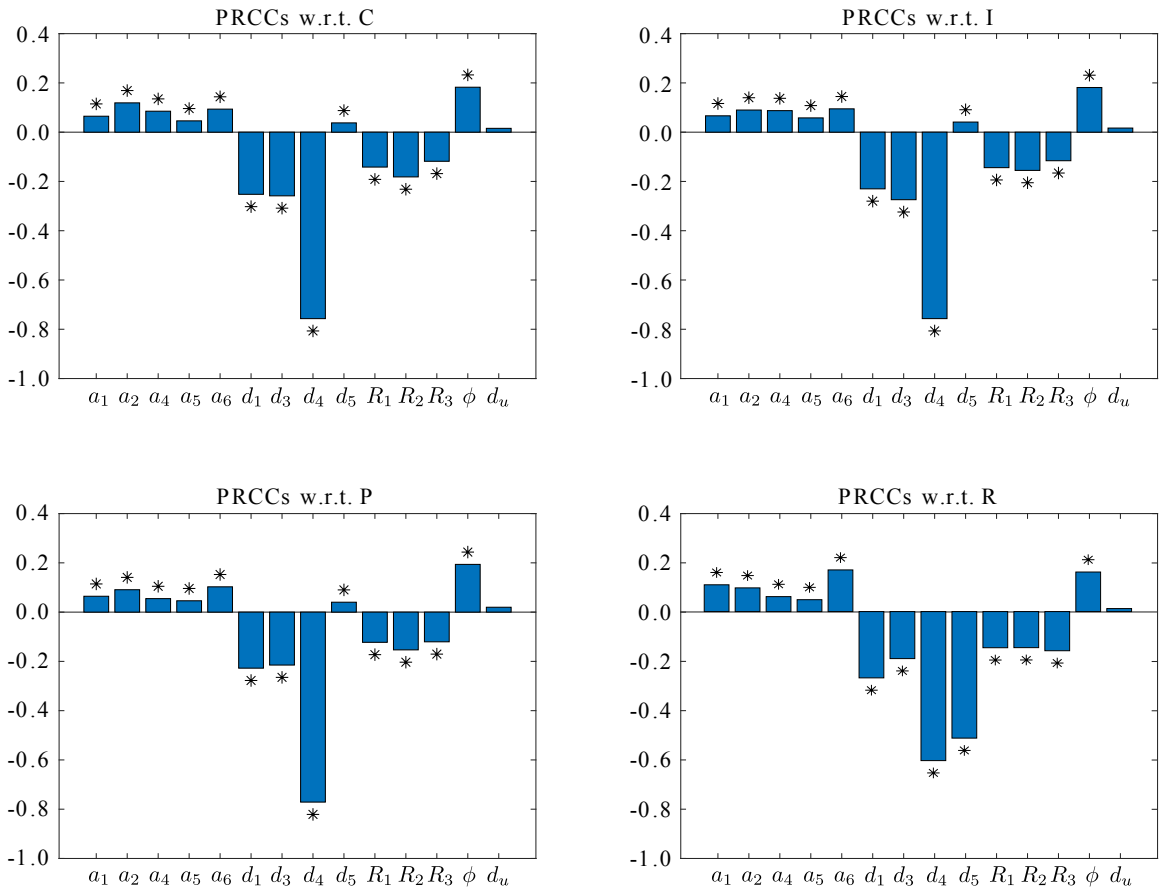


Figure 3.4: Parameter sensitivity analysis with baseline parameter values as given in Table 3.1. In each plot, d_u denotes the dummy parameter to which the impact of the model parameters is compared to ensure that any sensitivity of the model is not an artifact of the method. Results obtained for initial conditions $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 10^4,\ 0]$ cells, at $t = 60$ days and for 6,000 runs. Asterisk indicates a PRCC value that is significantly different from the dummy parameter, $p \leq 0.01$, and the parameter thus has a significant impact on the output variable.

3.2.4 Model simulation with default parameter set

We expect the model behaviour to change significantly depending on the magnitude of the perturbation away from the healthy steady state. For small perturbations, the system should be able to recover to the healthy steady state, but for perturbations above a certain threshold the system should tend to the diseased steady state. To explore this hypothesis in more detail, we simulate the model with the default parameter set, setting all initial conditions to zero, except for P_0 which is varied

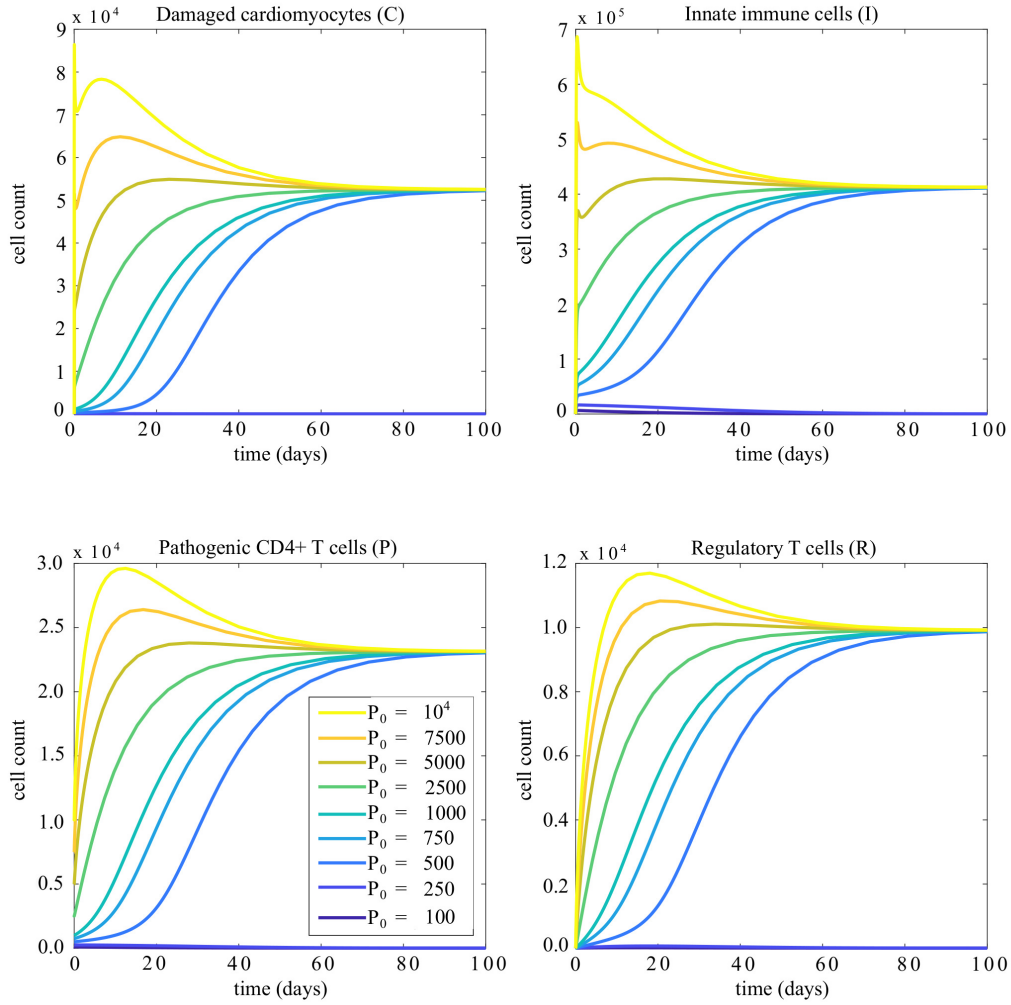


Figure 3.5: Time traces of the model variables for different initial conditions with default parameter values. The initial conditions are $[C_0,\ I_0,\ R_0] = [0,\ 0,\ 0]$ cells, with P_0 varied as noted.

to simulate different immune perturbations away from the healthy steady state at the origin. The time traces in Fig. 3.5 show that between 250 and 500 pathogenic CD4+ T cells are minimally required for the system to approach the diseased steady state, representing the fact that autoimmune myocarditis has developed. Indeed we compute by Eq. (3.5) that the cell count of pathogenic T cells at the saddle point is 339. These cell counts fall within the range of 100–10,000 T cells estimated based on existing literature [102, 132]. When autoimmune myocarditis develops, the T cell response takes approximately seven days to launch, while the innate immune response

develops almost instantaneously. These behavioural characteristics confirm that the model, simulated using the default parameter set, is capable of reproducing experimental observations [86, 133]. Thus, with this particular parameter set, the minimal set of immune cell types and pathways included in this model suffices to recapitulate the disease characteristics seen in patients, and to produce the behaviour required from a mathematical model of autoimmune myocarditis, as postulated in Fig. 3.2, in an ICI–free context.

For the default parameter set, the saddle point lies at approximately $[C,\ I,\ P,\ R] = [138,\ 20000,\ 339,\ 145]$ cells. Fig. 3.5 shows that once the system reaches cell counts above this saddle point, the immune response cannot be resolved and the development of autoimmune myocarditis is inevitable.

3.3 The effects of ICIs

In the previous section we established a mathematical model that is rooted in the experimental literature, and can replicate observations of the development of autoimmune myocarditis. We now turn to consider the effects that ICIs have on the system by exploring changes in the number and nature of the steady states as nivolumab and ipilimumab are administered.

Bifurcation structures provide insight into how the number of steady states and their linear stability change as a function of one or more model parameters. In the ICI–free model, the two stable steady states represent the two possible outcomes of autoimmune myocarditis: resolution and chronic inflammation. The changing values of a_6, d_4 and d_5 as ICIs are administered affect the number and linear stability of the steady states, and thus how the system evolves after an immune challenge. We assume in this chapter that the values of the parameters a_6, d_4 and d_5 remain constant over time, but can vary between simulations in order to represent different levels of

ICIs in the system. In Chapter 4 and Chapter 5 we relax this assumption and let ICI levels in the system, and with it the values of a_6, d_4 and d_5, vary over time.

3.3.1 Nivolumab

The left panel in Fig. 3.6 shows how the number of steady states changes when nivolumab is administered to the system, changing the values of a_6 and d_4. Nivolumab increases the rate of activation of T cells and decreases the rate of deactivation/death of pathogenic CD4+ T cells. Thus, as the dose of nivolumab is increased, the value of a_6 increases and the value of d_4 decreases, moving the system from the default parameter set (starred in the left panel in Fig. 3.6) towards the bottom right–hand corner. If enough nivolumab is applied, the system crosses a bifurcation which engenders a switch from three steady states to only two steady states.

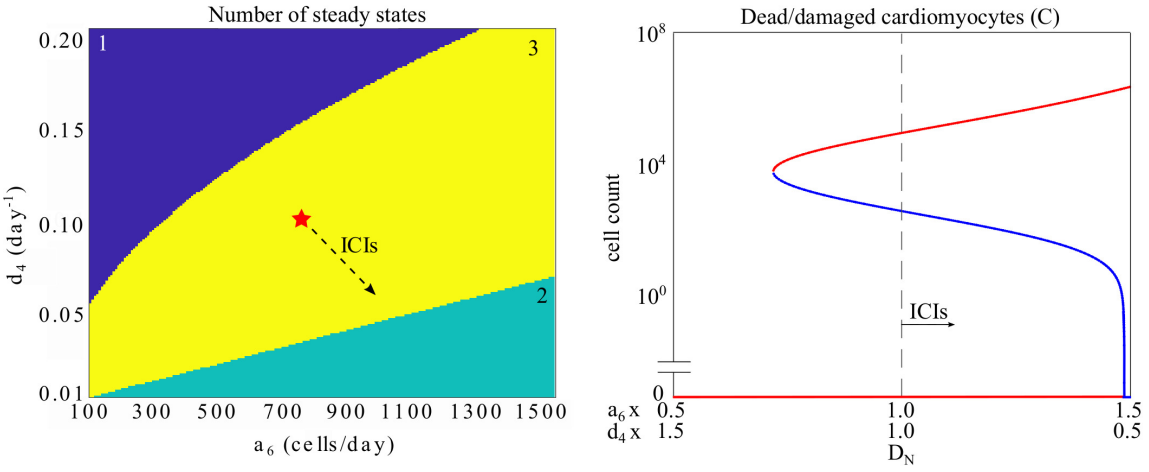


Figure 3.6: Left: Bifurcation diagram for varying values of a_6 and d_4. The yellow shading indicates regions of parameter space where the model has three non–negative steady states. The green shading indicates regions with two steady states, and the purple shading indicates regions with one steady state. The red star indicates the default parameter values of a_6 and d_4. The dashed arrow indicates the general direction the system moves in across parameter space as an increasing dose of nivolumab is applied. Right: A transcritical bifurcation occurs as the value of a_6 is increased from its default value of 750 cells day^{-1} and the value of d_4 is decreased from its default value of 0.1 day^{-1}. Red indicates a stable steady state, blue indicates a saddle point. The grey dashed line indicates the default values of a_6 and d_4. The arrow indicates the direction the system moves in when nivolumab is applied. Default values are used for all parameters unless otherwise noted.

To explore the type of bifurcation that occurs as nivolumab is applied, the values of a_6 and d_4 are varied together and the steady states and their linear stability are tracked. It is assumed here that the effect of nivolumab on these parameters is identical in terms of the percent change in parameter value as a function of nivolumab dose. The right panel in Fig. 3.6 shows what happens to the stable steady states (red) and saddle point (blue) as the values of a_6 and d_4 are varied between 50% and 150% of their default values. The vertical dashed line in the right panel in Fig. 3.6 indicates an ICI–free system, where a_6 and d_4 are at their default values; here there are three steady states, two stable steady states and one saddle point, as established in previous sections. As nivolumab is applied in increasing amounts the cell count at the saddle point decreases, and the saddle point moves towards the origin. This means that fewer active immune cells are required to push the system past the separatrix, thus applying nivolumab lowers the threshold for the development of autoimmune myocarditis. As a_6 nears 1.5 times its default value and d_4 nears 0.5 times its default value, the saddle point collides with the origin in a transcritical bifurcation, leaving the origin unstable. Past this bifurcation point, the system only has two steady states: a stable steady state with high immune cell counts and an unstable steady state at the origin. This means that any perturbation to the system away from the healthy steady state leads to the development of autoimmune myocarditis. Applying nivolumab to the system thus initially increases the likelihood of developing autoimmune myocarditis by lowering the threshold for immune perturbations to push the system across the separatrix until, for a sufficiently high dose, any perturbation will lead to disease as the healthy steady state at the origin becomes unstable. Additionally, as more nivolumab is applied, the cell counts at the diseased steady state increases which suggests more severe disease when autoimmune myocarditis develops.

3.3.2 Ipilimumab

As the dose of ipilimumab is increased, both the activation rate of T cells, a_6, and the death rate of Treg cells, d_5, increase. The left panel in Fig. 3.7 demonstrates that the system moves away from the default parameter set (starred) towards the top–right corner as more ipilimumab is applied. As with nivolumab, the system has three steady states when the default parameter values of a_6 and d_5 are used, and crosses a bifurcation as an increasing dose of ipilimumab is applied, leaving the system with two steady states. In contrast to nivolumab, only one parameter, a_6, controls the occurrence of this bifurcation. For large values of a_6, the number of steady states changes from three to two, but increasing only d_5 from its default value does not change the number of steady states.

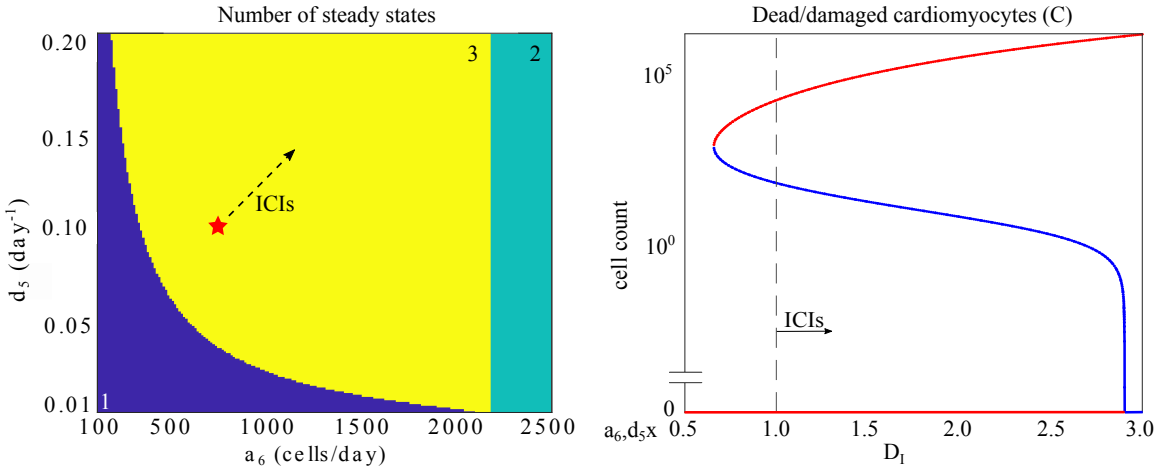


Figure 3.7: Left: Bifurcation diagram for varying values of a_6 and d_5 as all other parameters are kept at default values. The yellow shading indicates regions of parameter space where the model has three non–negative steady states. The green shading indicates regions with two steady states, and the purple shading indicates regions with one steady state. The red star indicates the default values of a_6 and d_5. The dashed arrow indicates the general direction the system moves in as an increasing dose of ipilimumab is applied. Right: A transcritical bifurcation occurs as the value of a_6 is increased from 750 cells day^{-1} and the value of d_5 is increased from 0.1 day^{-1}. Default values are used for all other parameters. Red indicates a stable steady state, blue indicates a saddle point. The grey dashed line indicates the default values of a_6 and d_5. The arrow indicates the direction the system moves in when ipilimumab is applied.

To explore the type of bifurcation that occurs as the dose of ipilimumab is increased, a_6 and d_5 are increased simultaneously while the number of steady states and their linear stability are tracked (see right panel in Fig. 3.7). Again, it is assumed here that the effect of ipilimumab on these parameters is identical in terms of the percentage change in parameter value as a function of ipilimumab dose. The grey dashed line indicates where a_6 and d_5 attain their default values. As the dose of ipilimumab is increased the system moves to the right, away from the dashed line. Similar to nivolumab, increasing doses of ipilimumab initially lowers the threshold for autoimmune myocarditis to develop before the system crosses a transcritical bifurcation and the origin becomes unstable, making the development of the disease almost inevitable. To reach this bifurcation, enough ipilimumab must be added to increase a_6 to almost three times its default value. Furthermore, as with nivolumab, we observe an increase in cell counts at the diseased steady state as the dose of ipilimumab is increased, suggesting more severe autoimmune myocarditis develops as a consequence of higher doses of ICIs.

In summary, applying ICIs in increasing doses can drastically change the number and linear stability of the steady states of the model. The presence of increasing levels of nivolumab or ipilimumab leads to a lower threshold for the development of autoimmune myocarditis as the saddle point moves towards the healthy steady state. For sufficiently high levels of ICIs, a bifurcation occurs, leading to an unstable origin and a system in which even the smallest perturbation away from the healthy steady state will lead to disease.

3.4 Patient–specific responses

Not all patients receiving ICI treatment develop autoimmune myocarditis. The risk of developing autoimmune myocarditis is represented in the model by how close the

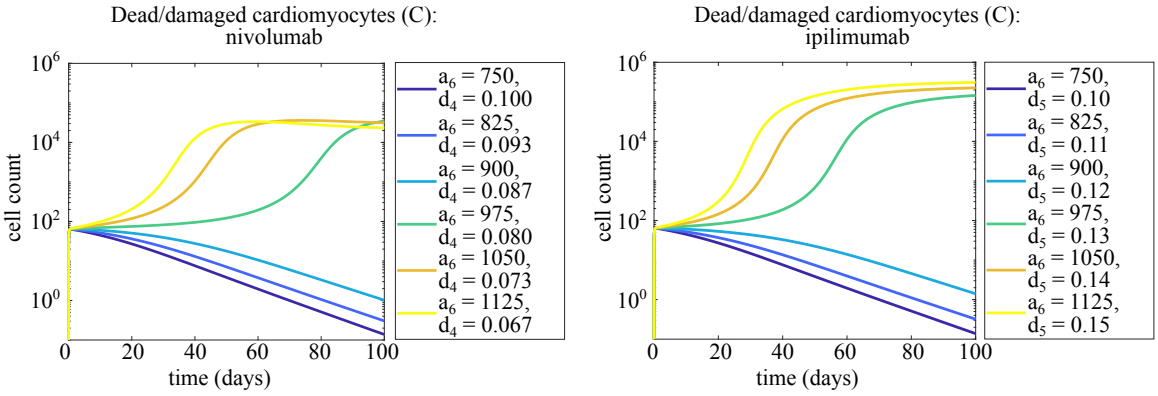


Figure 3.8: Left: Dynamics of dead/damaged cardiomyocytes in response to increasing levels of nivolumab, represented by changes in the values of a_6 and d_4. The default parameter set is used, unless otherwise stated. Initial conditions $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 200,\ 0]$ cells. Right: Dynamics of dead/damaged cardiomyocytes in response to increasing levels ipilimumab, represented by changes in the values of a_6 and d_5. The default parameter set is used, unless otherwise stated. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 200,\ 0]$ cells.

separatrix is to the healthy steady state. The closer the separatrix is to the healthy steady state, the higher the risk of autoimmune myocarditis. The position of the separatrix is dictated by the parameter values, and how much it shifts as a consequence of ICI treatment depends on the doses of ICIs administered and how the level of ICIs impacts the values of a_6, d_4 and d_5. Ultimately, the aim of this research is to predict the levels of ICIs a patient can receive without developing autoimmune myocarditis, or how different patients might respond to the same treatment protocol. Although the parameterized model presented here is not sufficient for this goal, we can explore what effect changes in the values of a_6, d_4 and d_5 have on model behaviour.

Fig. 3.8 shows model predictions for how an individual might respond to an immune perturbation given the administration of specific doses of nivolumab and ipilimumab. The initial condition $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 200,\ 0]$ cells represents an immune perturbation that in an ICI–free system brings the systems close to the separatrix, but does not cause autoimmune myocarditis to develop. However, a relatively small change in the position of the separatrix due to the dosing of ICIs could change the perturbation outcome. For a patient who experiences fluctuations in immune cell

counts of this magnitude, it is thus important to carefully consider the amount of ICIs administered, to ensure they do not develop autoimmune myocarditis. The left panel in Fig. 3.8 shows how the model evolves after an influx of pathogenic CD4+ T cells as the value of a_6 is increased step–wise up to 1.5 times its default value, and d_4 is simultaneously decreased to 0.67 times its default value. For small doses of nivolumab (blue time traces), the model predicts that the patient can resolve the immune response and autoimmune myocarditis does not develop. For larger doses of nivolumab, however, the model predicts that the perturbation pushes the dynamics across the separatrix and disease develops (green, yellow and orange time traces). The right panel in Fig. 3.8 shows similar results for ipilimumab, where doses of ipilimumab are represented by step–wise increases of a_6 and d_5 up to 1.5 times their respective default values.

Another important question to ask is how different patients might respond to the same dose of ICIs. The rate at which different immune processes occur is likely to vary between individuals, thus exploring variations in parameter values and the effect this has on outcome in ICI–free and ICI–dosed systems is important. Parameter a_1, the rate with which pathogenic CD4+ T cells damage or kill cardiomyocytes, is chosen to illustrate the impact of individual variation, and its value is varied around its default value of 80 day^{-1}. The four panels in Fig. 3.9 show how a slight increase or decrease in the value of a_1 can have a large effect on the outcome after immune perturbation when different types of ICI therapy are applied. The initial conditions $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 200,\ 0]$ cells are used to simulate a small immune perturbation which, in an ICI–free environment, pushes the system close to the separatrix but not across it. The top–left panel in Fig. 3.9 shows that in an ICI–free environment the outcome is always resolution of inflammation, regardless of small changes to the value of a_1. For nivolumab and ipilimumab monotherapy, ICI doses are simulated by a 10% increase or decrease in the appropriate parameter values. The incidence rate of

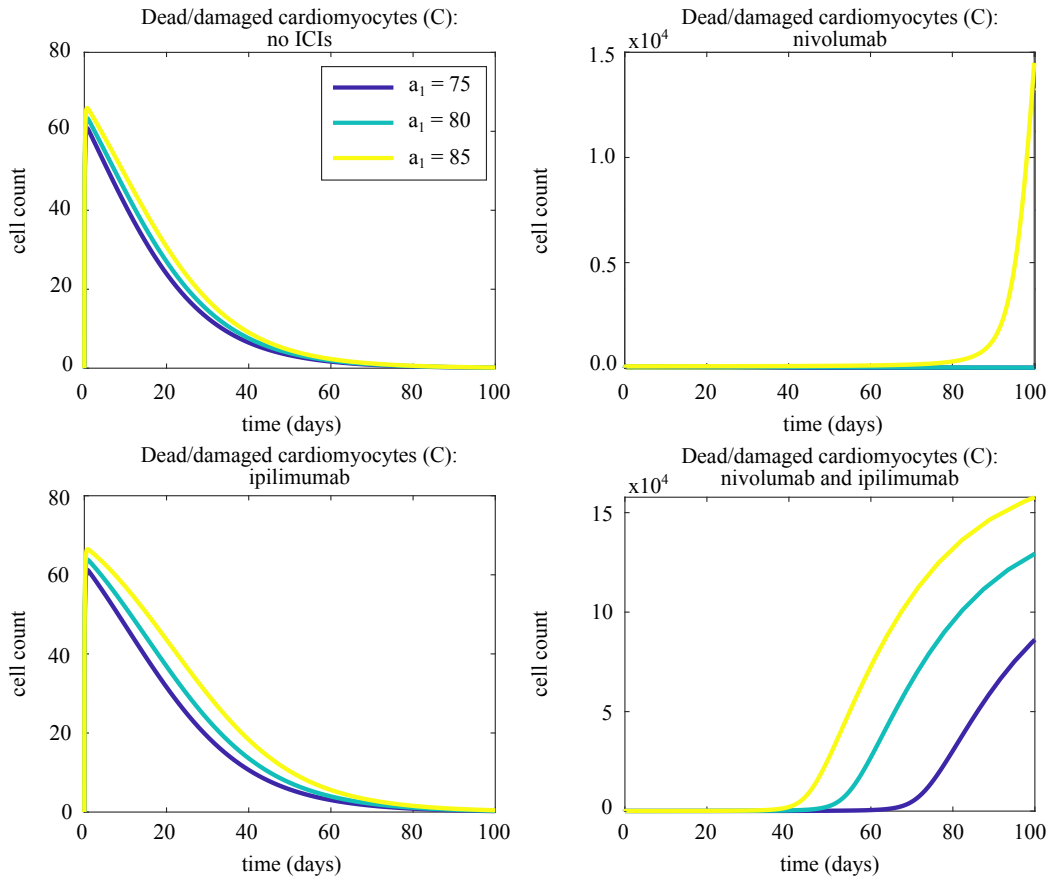


Figure 3.9: Varying outcomes for different values of a_1 to a dose of nivolumab (top right), ipilimumab (bottom left), or both (bottom right), as indicated by the number of damaged/dead cardiomyocytes. ICI doses are modelled by a 10% increase in the value of a_6 and a 10% decrease in the value of d_4 (nivolumab monotherapy), a 10% increase in the values of a_6 and d_5 (ipilimumab monotherapy), or a 21% increase in a_6, a 10% decrease in d_4 and a 10% increase in d_5 (combination therapy). The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 200,\ 0]$ cells, and the default parameter values are used, unless otherwise noted.

autoimmune myocarditis as a side–effect of combination therapy (0.27%) compared to monotherapy (0.06–0.09%) suggests a nonlinear effect on the immunological mechanisms underlying the disease when both ICIs are administered [5, 95]. Although these incidence rates are somewhat uncertain due to the rarity of autoimmune myocarditis, we use $a_6 = 1.1 \times 1.1\times$ default value, *i.e.* a 21% increase in value, for the simulation of combination therapy to account for this nonlinear effect when both nivolumab and ipilimumab are applied. These increases in the values of a_6, d_4 and d_5 ensure that,

although the threshold for developing autoimmune myocarditis is lower than in an ICI–free environment, the system has not crossed the transcritical bifurcation that occurs for high doses of ICIs, so the perturbation away from the healthy steady state introduced by the initial conditions is not guaranteed to lead to the development of disease. The bottom left panel in Fig. 3.9 shows that the model predicts that administering ipilimumab does not change the outcome of inflammation and the system still resolves inflammation for all values of a_1. When administering nivolumab (top right panel), however, this changes for $a_1 = 85$ day^{-1} the model predicts that autoimmune myocarditis will develop because the saddle point for $a_1 = 85$ day^{-1}, $a_6 = 1.21 \times 750$ cells day^{-1} and $d_4 = 0.9 \times 0.1$ day^{-1} lies at $[C,\ I,\ P,\ R] = [50,\ 1.1 \times 10^4,\ 170,\ 59]$ cells, which the system is above at the final time point of the simulation. When administering combination therapy (bottom right panel), all values of a_1 explored here lead to autoimmune myocarditis. A small change in a parameter can thus lead to different outcomes when ICIs are applied, showing that individual patient characteristics are important to consider when it comes to designing treatment regimen.

3.5 Parameter identifiability analysis

Although no experimental data is currently available from which we can estimate model parameters, the potential exists for such data to become available in the future. When it does, it will be essential to ensure that the model parameters are practically identifiable given the quantity, quality and type of experimental data obtained. Identifiability analysis establishes the requirements for a data set to allow parameter value estimation. These requirements include the specific output variables and time points that must be observed, together with the level of measurement noise that can be tolerated. Performing an identifiability analysis prior to carrying out experiments allows for more efficient experiment design, and more informed data col-

lection. In this section, we explore the structural and practical identifiability of the model presented in this chapter, noting the challenges of applying these methods to models with a high ratio of parameters to output variables.

3.5.1 Structural identifiability

A model is structurally identifiable if, given perfect, noise–free and time–continuous data for a set of output variables, all parameters in the model can be uniquely identified. The number of dependent variables that must be observed for local structural identifiability can be determined using the scaling method developed by Castro and de Boer [25]. This method uses the scaling invariance of the model equations to determine whether the model parameters are identifiable or not. In brief, for a given set of observed variables, the method can be used to assess which parameters are structurally identifiable by determining which parameter values can be changed without affecting the observed output variables of the model. Parameters which can be scaled without affecting model output are not structurally identifiable.

For a model with n dependent variables, $x_1, \ldots, x_n$, and m parameters, $p_1, \ldots, p_m$, the system of ODEs is written in the form

$$\frac{\mathrm{d}x_i}{\mathrm{d}t} = f_{i,1}(x_1, \ldots, x_n; p_1, \ldots, p_m) + \ldots + f_{i,k_i}(x_1, \ldots, x_n; p_1, \ldots, p_m), \tag{3.13}$$

where $i \in \{1, \ldots, n\}$ denotes the dependent variable, and the function $f_{i,j}$, $j = 1, \ldots, k_i$, is the jth term on the RHS of Eq. (3.13) describing the evolution of the ith variable over time. The ODEs must be written in such a way that the functions $f_{i,j}$ for a given i and $j \in \{1, \ldots, k_i\}$ are linearly independent.

To explore the structural identifiability of the parameters in the model, we scale all parameters p_l, $l = 1, \ldots, m$, and variables x_i, $i = 1, \ldots, n$, in the model by the scaling factors u_{p_l} and u_{x_i}, respectively. For variables that we observe, we set $u_{x_i} = 1$.

so that x_i must be unchanged by scaling. We conclude that a given parameter p_l is structurally identifiable for a given set of observed variables if $u_{p_l} = 1$ is the unique value of u_{p_l} for which the observed output of the scaled system of ODEs is the same as the observed output of the unscaled system of ODEs. If this is not the case, p_l can be scaled and, by scaling another parameter or unobserved variable, the output remains unchanged. Hence there is not a unique value of p_l for which the observed output is obtained and the parameter p_l is not structurally identifiable for this specific observed output. Because the terms $f_{i,j}$ in the ODEs of the model are linearly independent, we can assess on a term–by–term basis whether the process of scaling the model parameters change the model output by setting the unscaled and scaled versions of each term equal to each other according to

$$f_{i,j}(x_1, x_2; \lambda_1, \lambda_2, \lambda_3, \lambda_4) = \frac{1}{u_{x_i}} f_{i,j}(u_{x_1}x_1, u_{x_2}x_2; u_{\lambda_1}\lambda_1, u_{\lambda_2}\lambda_2, u_{\lambda_3}\lambda_3, u_{\lambda_4}\lambda_4). \quad (3.14)$$

In the next section, we demonstrate how this method works in practice with a simple example.

3.5.1.1 An example

We reproduce an example from the paper by Castro and de Boer here to demonstrate how their method for determining structural identifiability works [25]. We consider the system of ODEs

$$\frac{\mathrm{d}x_1}{\mathrm{d}t} = \lambda_1 x_1^2 + \lambda_2 x_1 x_2, \quad (3.15)$$

$$\frac{\mathrm{d}x_2}{\mathrm{d}t} = \lambda_3 x_1^2 + \lambda_4 x_1 x_2, \quad (3.16)$$

for $\lambda_i > 0$, $i = 1, \ldots, 4$, and the initial conditions $x_1(0) = 0$ and $x_2(0) = 0$. This system has $n = 2$ variables, $m = 4$ parameters and each ODE consists of two terms, so $k_1 = 2$ and $k_2 = 2$. We assume that we observe the variable x_1. We now re–scale

all model parameters and unobserved state variables, so x_2 becomes $u_{x_2}x_2$, and λ_i becomes $u_{\lambda_i}\lambda_i$ for $i = 1, \ldots, 4$. Because the terms on the RHS of Eq. (3.15) and Eq. (3.16) are linearly independent, we can set the unscaled version of each term equal to its scaled version to obtain the four relations

$$\lambda_1 x_1^2 = u_{\lambda_1}\lambda_1 x_1^2, \tag{3.17}$$

$$\lambda_2 x_1 x_2 = u_{\lambda_2}u_{x_2}\lambda_2 x_1 x_2, \tag{3.18}$$

$$\lambda_3 x_1^2 = \frac{u_{\lambda_3}}{u_{x_2}}\lambda_3 x_1^2, \tag{3.19}$$

$$\lambda_4 x_1 x_2 = u_{\lambda_4}\lambda_4 x_1 x_2. \tag{3.20}$$

The unique solution to Eq. (3.17) is $u_{\lambda_1} = 1$. This means that if the value of the parameter λ_1 is changed, this changes the observed output of the system. As a result, we conclude that the parameter λ_1 is structurally identifiable if the state variable x_1 is observed. The unique solution to Eq. (3.20) is $u_{\lambda_4} = 1$, and we similarly conclude that the parameter λ_4 is structurally identifiable if the variable x_1 is observed. Equations (3.18) and (3.19) hold as long as $u_{\lambda_2}u_{x_2} = 1$ and $u_{\lambda_3}/u_{x_2} = 1$, respectively. The parameters λ_2 and λ_3 are thus not structurally identifiable if only x_1 is observed, because a change in the value of either of these parameters could lead to a change in the unobserved variable x_2 yet leave the observed output, x_1, unchanged. When only x_1 is observed the values of λ_2 and λ_3 which produce that output are non–unique. We note that if x_2 is also observed, so $u_{x_2} = 1$, then $u_{\lambda_2} = 1$ is the unique solution to Eq. (3.18) and $u_{\lambda_3} = 1$ is the unique solution to Eq. (3.19). Thus, all parameters in the system are structurally identifiable if both state variables are observed. In the next section, we apply this method to our model of autoimmune myocarditis.

3.5.1.2 Structural identifiability of the model of autoimmune myocarditis

Our aim is to determine which variables we must observe for all model parameters to be structurally identifiable. In our model of autoimmune myocarditis, defined by Eq. (3.1)–(3.4), we have four variables which we might observe, 13 parameters, and a total of ten linearly independent terms across all four ODEs. We scale all variables x_i, $i = 1, \ldots, 4$, by u_{x_i} and all parameters p_l, $l = 1, \ldots, 13$, by u_{p_l} and obtain the following equalities according to Eq. (3.14):

$$\frac{a_1 P^2}{R_1 + R} = \frac{u_{a_1} u_P^2}{u_R u_C} \; \frac{a_1 P^2}{\frac{u_{R_1}}{u_R} R_1 + R}; \tag{3.21}$$

$$\frac{a_2 I}{R_2 + R} = \frac{u_{a_2} u_I}{u_R u_C} \; \frac{a_2 I}{\frac{u_{R_2}}{u_R} R_2 + R}; \tag{3.22}$$

$$d_1 C = u_{d_1} d_1 C; \tag{3.23}$$

$$\frac{a_4 P}{R_1 + R} = \frac{u_{a_4} u_P}{u_R u_I} \; \frac{a_4 P}{\frac{u_{R_1}}{u_R} R_1 + R}; \tag{3.24}$$

$$a_5 C = \frac{u_{a_5} u_C}{u_I} \; a_5 C; \tag{3.25}$$

$$d_3 I = u_{d_3} d_3 I; \tag{3.26}$$

$$\phi \frac{a_6 C}{R_3 + R} = \frac{u_{a_6} u_C}{u_R u_P} \; u_\phi \phi \frac{a_6 C}{\frac{u_{R_3}}{u_R} R_3 + R}; \tag{3.27}$$

$$d_4 P = u_{d_4} d_4 P; \tag{3.28}$$

$$(1 - \phi) \frac{a_6 C}{R_3 + R} = \frac{u_{a_6} u_C}{u_R^2} \; (1 - u_\phi \phi) \frac{a_6 C}{\frac{u_{R_3}}{u_R} R_3 + R}; \tag{3.29}$$

$$d_5 R = u_{d_5} d_5 R. \tag{3.30}$$

We want to determine for which variables we must have $u_{x_i} = 1$ in order for $u_{p_l} = 1$ to be the unique solution to the Eq. (3.21)–(3.30) for all $l = 1, \ldots, 13$. This will enable us to determine which dependent variables must be observed in order for all model parameters to be structurally identifiable.

From Eqs. (3.23), (3.26), (3.28), and (3.30), we can conclude that the deactivation/clearance rates d_1, d_3, d_4 and d_5 are structurally identifiable regardless of which variables we observe, because $u_{d_1} = 1$, $u_{d_3} = 1$, $u_{d_4} = 1$, and $u_{d_5} = 1$ are the unique solutions to these four equations, respectively. Furthermore, from Eq. (3.29) we observe that $u_\phi = 1$ is required for the equality to hold, thus the parameter ϕ is also structurally identifiable regardless of which variables we observe.

From Eq. (3.21)–(3.30), we derive the following set of conditions for the parameter scalings:

$$\frac{u_{a_1} u_P^2}{u_C u_R} = 1; \tag{3.31}$$

$$\frac{u_{a_2} u_I}{u_C u_R} = 1; \tag{3.32}$$

$$\frac{u_{a_4} u_P}{u_I u_R} = 1; \tag{3.33}$$

$$\frac{u_{a_5} u_C}{u_I} = 1; \tag{3.34}$$

$$\frac{u_{a_6} u_C}{u_P u_R} = 1; \tag{3.35}$$

$$\frac{u_{a_6} u_C}{u_R^2} = 1. \tag{3.36}$$

From Eq. (3.31), we find that we must observe the variables C, P and R for the parameter a_1 to be structurally identifiable. From Eq. (3.32), the variables C, I and R must be observed for the parameter a_2 to be structurally identifiable. Thus, in order for the two parameters a_1 and a_2 alone to be structurally identifiable, we must observe all four state variables. Equations (3.33)–(3.36) show that if all four variables are observed, then the parameters a_4, a_5, and a_6 are also structurally identifiable.

This leaves us with the parameters R_1, R_2 and R_3. We take Eq. (3.21) as an example. We have already shown that $(u_{a_1} u_P^2)/(u_C u_R) = 1$, so for Eq. (3.21) to hold, we must have $u_{R_1}/u_R = 1$. From Eqs. (3.22), (3.27) and (3.30), we have $u_{R_2}/u_R = 1$ and $u_{R_3}/u_R = 1$, so we have $u_{R_1} = u_{R_2} = u_{R_3} = u_R$. We have already shown

that we must observe R to achieve structurally identifiability of other parameters in the model, so $u_R = 1$ and therefore the parameters R_1, R_2 and R_3 are structurally identifiable when all state variables are observed.

In conclusion, we must observe all four state variables for the model parameters to be structurally identifiable. Global structural identifiability analysis performed using DAISY software confirms that our model is only globally structurally identifiable if all four variables are measured as outputs [13].

3.5.2 Practical identifiability

Given the results of the structural identifiability analysis, we explore practical identifiability analysis knowing that a minimum constraint on the data is that all four variables must be observed to estimate parameter values. Several methods exist to explore the practical identifiability of models, such as Markov Chain Monte Carlo (MCMC), the Fisher information matrix, and maximum likelihood profiling. Here, we choose to apply maximum likelihood profiling as it is computationally less expensive than MCMC and handles nonlinear models more reliably than the Fisher information matrix [161].

We compute the maximum loglikelihood profiles of subsets of the 13 parameters in the ICI–free model using the methods explained in Simpson *et al.* [140] which we detail below. Firstly, a synthetic data set is obtained by simulating the model with default parameter values for 60 days, obtaining an "observation" of all four variables on days 0, 1, 2, 5, 10, 30 and 60. The observations are concentrated in the first few days to capture the fast early–time dynamics of the system. To each data point we add measurement noise, which is sampled from a normal distribution with a mean of zero and a standard deviation of 5% of the cell count at that data point.[1] In order to investigate identifiability for scenarios where autoimmune myocarditis develops, we

[1]Using 1% noise instead does not qualitatively alter the results.

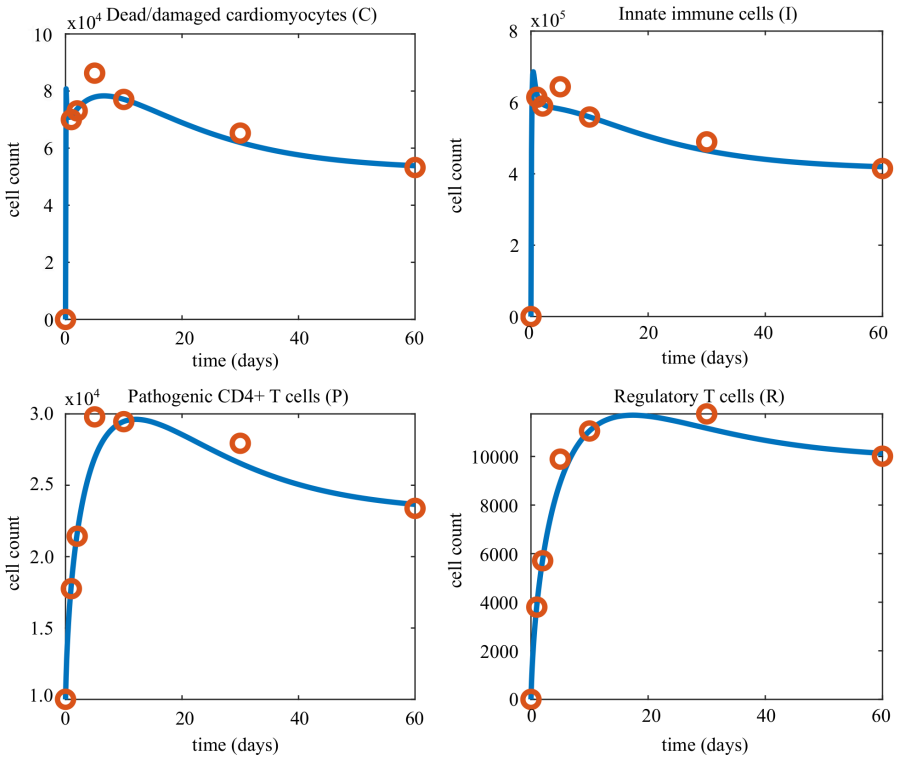


Figure 3.10: A representative simulation (blue line) and resulting synthetic data set (red circles). Observations from the simulation (blue line) are made on days 0, 1, 2, 5, 10, 30 and 60. To each data point, normally distributed noise with a standard deviation of 5% of the cell count of that observation is added.

take as initial conditions $[C_0,\ I_0,\ P_0,\ R_0] = [0,\ 0,\ 10^4,\ 0]$. A representative synthetic data set is shown in Fig. 3.10.

Having obtained a synthetic data set, we now proceed to compute the maximum loglikelihood profiles. The large number of parameters and relatively small number of variables in the model makes a comprehensive practical identifiability analysis computationally infeasible as optimization over the full 13–dimensional parameter space is computationally prohibitive. The parameter set must therefore be split into subsets to investigate which subsets of model parameters are identifiable. Such analysis can potentially inform a sequence of experiments in which parameters outside a practically identifiable subset are determined independently before the model is fit to data to estimate the values of the practically identifiable parameters. We select a subset S of model parameters for which we wish to compute maximum loglikelihood

profiles. All parameters not in S are kept at their default value throughout the analysis. Maximum loglikelihood profiles for the parameters in S are computed as follows:

Step 1. Select a parameter of interest from the subset S.

Step 2. Select a list of values $[v_1, \ldots, v_k]$, ordered smallest to largest, for which the maximum loglikelihood profile for the parameter of interest is to be calculated.

Step 3. Iterate over the ordered list of values selected in **Step 2**. For q in $[1, \ldots, k]$:

Step 3a. Set the parameter of interest equal to v_q.

Step 3b. If $q = 1$, set the initial guesses for the optimal values of the remaining parameters in S equal to their respective default values. If $q \neq 1$, set the initial guesses for values of the remaining parameters to their respective optimal values obtained when the parameter of interest is equal to v_{q-1}.

Step 3c. While keeping the parameter of interest constant at the value from **Step 3a**, optimize the values of the other parameters in S such that the loglikelihood is maximized, starting the initial guesses set in **Step 3b**. The loglikelihood is given by

$$L^2 = \left(\sum_{j=1}^{4} \left(\sum_{i=1}^{7} -0.5 \log(2\pi) - \log(SD_{i,j}) - \frac{(y_{i,j} - y_{i,j}^d)^2}{2\ SD_{i,j}^2} \right) \right)^2, \tag{3.37}$$

where $y_{i,j}$ is the model output of the variable j at time point i for a given parameter set, $y_{i,j}^d$ is the corresponding synthetic data point, and $SD_{i,j}$ is the standard deviation for a specific data point, computed as $SD_{i,j} = 0.05 y_{i,j}^d$. Our model has four variables, hence $j = 1, \ldots, 4$, and we observe each of those variables at the seven time points.

Step 4. Plot the the maximum loglikelihoods computed in **Step 3** as a function of the value of the parameter of interest.

Step 5. Repeat from **Step 1** until a maximum loglikelihood profile is obtained for every parameter in S.

The range of the parameter of interest in **Step 2** is set to be between 0.1 and five times its default parameter value (see Table 3.1). For the parameter values to be optimized in **Step 3c**, the lower bound is 10^{-14} all parameters, and the upper bound is positive infinity for all parameters except for ϕ, which cannot exceed one. The optimization in **Step 3c** is performed using a global search algorithm, with the MATLAB optimizer fmincon and algorithm spq. We specify the FunctionTolerance and the XTolerance to be 10^{-10}.

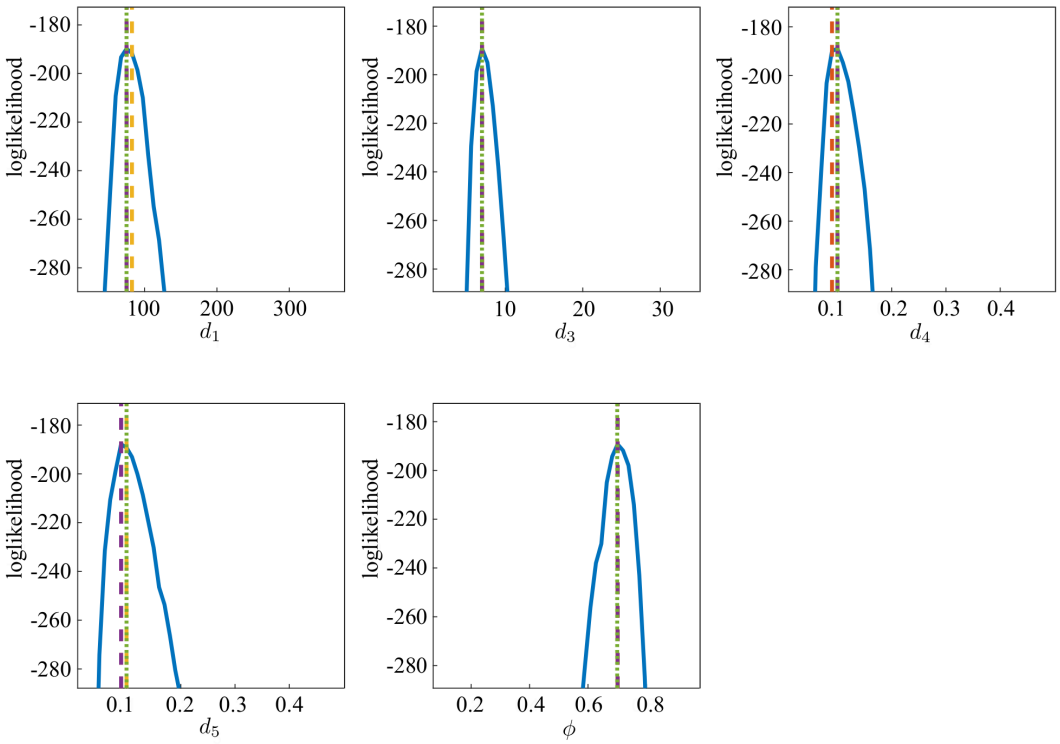


Figure 3.11: Maximum loglikelihood profiles (blue) of the subset of parameters d_1, d_3, d_4, d_5, ϕ in the ICI–free model. The green dotted line indicates the true value of the parameter. The purple dashed line indicates the value of the parameter for which the maximum loglikelihood value is attained. The red and yellow dashed lines indicate the 95% confidence interval around the loglikelihood maximum. Green, purple, yellow and red lines may overlap for narrow confidence intervals. For values of d_3, d_4, d_5 and ϕ for which no maximum loglikehood is plotted, the maximum loglikelihood was lower than the range to which the y–axis is limited here.

The first subset of model parameters that we investigate consists of the death rates d_1, d_3, d_4, d_5, ϕ. Fig. 3.11 shows the profile loglikelihoods for this subset. These parameters form a practically identifiable subset for the given synthetic data set, as indicated by the sharply peaked profiles which are maximized for the default parameter value, indicated with a green dotted vertical line. If all other model parameters, *i.e.* the activation rates and the inhibition parameters can be independently estimated, then Fig. 3.11 demonstrates that the death rates and ϕ can be confidently estimated from data of similar type, quantity and quality as the synthetic data generated here.

The subset of model parameters consisting of the activation rates a_1, a_2, a_4, a_5 is an example of a subset of parameters which is not a practically identifiable subset given the noise level and specific algorithm settings used here, as is shown in Fig. 3.12. The profile loglikelihood of parameter a_2 is completely flat, meaning that within the range explored here, we cannot say anything useful about its value. The loglikelihood profile for parameter a_1 is more sharply peaked, indicating more confidence in the estimated parameter value. However, the maximum loglikelihood is close but not equal to the value of a_1 used to generate the synthetic data. For parameters a_4 and a_5 we see a large plateau of the profile loglikelihood, meaning that we can establish a range within which the parameter value is likely to lie, as indicated by the high loglikelihood, but we cannot confidently identify a single value for these parameters.

Instead of grouping parameters by similarity in biological meaning, the subsets can also be formed by grouping parameters by the ODE they appear in. As is to be expected, this leads to practically unidentifiable subsets as parameter values in these subsets can be adjusted to compensate for the varying value of the parameter of interest. An example of this is shown in Fig. 3.13, which shows the flat profile loglikelihoods for the parameter subset a_1, a_2, d_1, all of which appear in Eq. (3.1), describing the dynamics of damaged/dead cardiomyocytes. The flatness of these profile loglikelihoods indicates that this is a practically unidentifiable subset of parameters for the chosen data set.

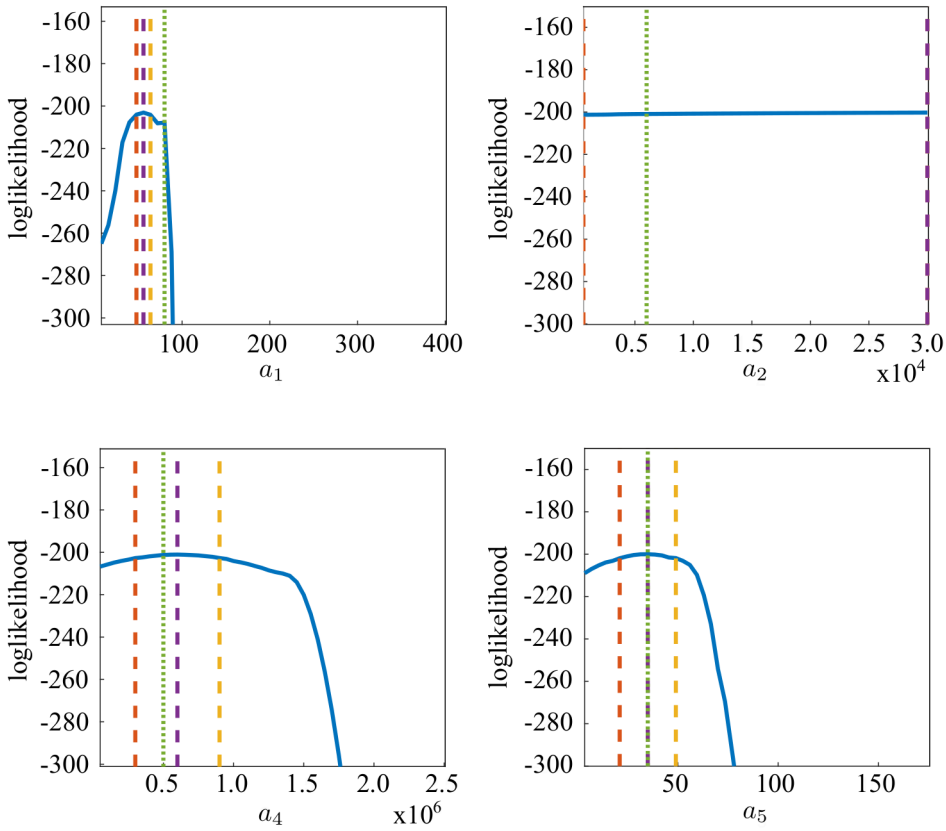


Figure 3.12: Maximum loglikelihood profiles (blue) for the parameter subset a_1, a_2, a_4, a_5. The green dotted line indicates the true value of the parameter. The purple dashed line indicates the value of the parameter for which the maximum loglikelihood value is attained. The red and yellow dashed lines indicate the 95% confidence interval around the loglikelihood maximum. Green, purple, yellow and red lines may overlap for narrow confidence intervals, or fall outside the range of parameter values plotted here. For values of a_1, a_2, a_4 and a_5 for which no maximum loglikehood is plotted, the maximum loglikelihood was lower than the range to which the y–axis is limited here.

We have thus shown that for the specific type of data set assumed here, certain subsets of the model parameters are practically identifiable, suggesting that a number of parameter values will need to be estimated independently before the remaining ones can be determined by fitting the model to this type of data set.

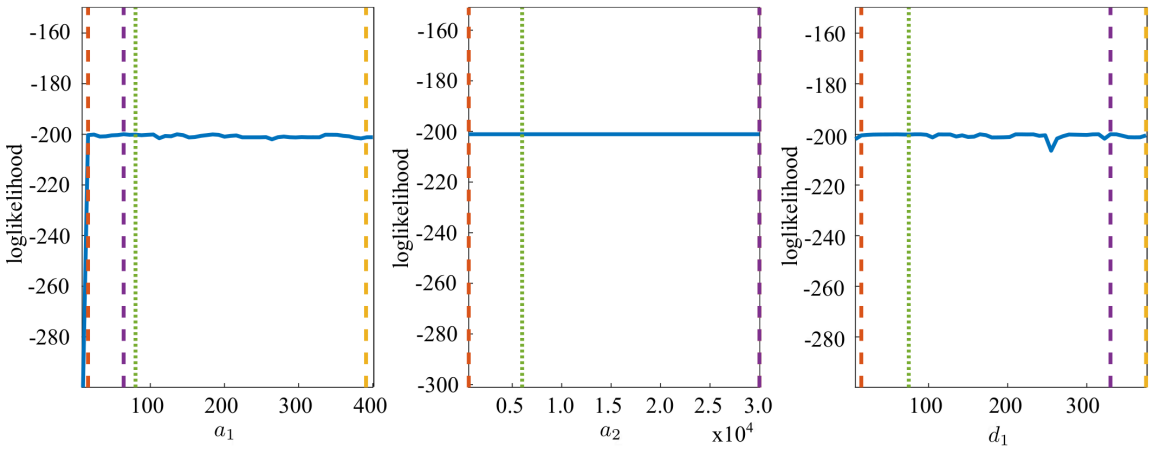


Figure 3.13: Maximum loglikelihood profiles (blue) for the parameter subset a_1, a_2, d_1. The green dotted line indicates the true value of the parameter. The purple dashed line indicates the value of the parameter for which the maximum loglikelihood value is attained. The red and yellow dashed lines indicate the 95% confidence interval around the loglikelihood maximum. Green, purple, yellow and red lines may overlap for narrow confidence intervals, or fall outside the range of parameter values plotted here.

3.6 Discussion

The main novel contribution of this chapter is the development and analysis of the first mathematical model of autoimmune myocarditis and the impact of ICIs on its development and progression. Not only is this the first model to capture the minimal set of immunological interactions required to reproduce the characteristics of the disease observed in patients, it is, to our knowledge, the first explicit model of a side–effect of cancer therapy. Although other models have considered the effect of certain therapies on a single variable representing "normal" cells [81, 82], the highly complex network of interactions between the immune system, healthy tissue and the treatment warrants a more detailed exploration if it is to be completely understood.

The model of autoimmune myocarditis presented here describes how the numbers of damaged/dead cardiomyocytes, innate immune cells, pathogenic CD4+ T cells and regulatory T cells evolve over time, eventually leading to either resolution or chronic

inflammation. The cell types and interactions in this model were included because they have been previously identified through experimental research as playing a role in the immunological mechanisms underlying this disease. The fact that the model can reproduce observations in patients with autoimmune myocarditis supports the hypothesis that the cell types included in this model form a set of essential immune components for the development of the disease.

We have shown that the addition of ICIs makes it more likely that autoimmune myocarditis develops, *i.e.* a smaller perturbation away from the healthy state is required to trigger the disease. Furthermore, the number of immune cells in the diseased state is higher when ICIs are present, suggesting more severe disease. We have also shown that small changes in parameter values, representing slight differences between individual patients, can influence whether a particular individual experiences resolution or chronic inflammation. For a certain ICI dose or immunological perturbation, the model predicts that some patients may be pushed towards chronic inflammation while others return to a healthy state.

Structural identifiability analysis of the model showed that the parameters of the model are only structurally identifiable if all four state variables are observed. Through practical parameter identifiability analysis, we have shown that, for the chosen type of data set, there are subsets of model parameters which are practically identifiable. Exploring the practical identifiability of the model further by varying the subsets of parameters, noise levels, and number and timing of measurements will provide more insight into the required quantity, quality and type of data that must be obtained experimentally to estimate different subsets of model parameters.

Our model predicts that for sufficiently large doses of either nivolumab or ipilimumab, the development of autoimmune myocarditis is inevitable for even the smallest perturbation away from the healthy steady state. This is a prediction that could potentially be tested using an *in vitro* assay where the dosage of ICIs and the size

of the perturbation, *i.e.* the number of active pathogenic CD4+ T cells added to the system, can be carefully controlled. Results of such an experiment could be used to validate the mechanisms included in the model as the minimum required set for reproducing the characteristics of autoimmune myocarditis as observed in patients.

The parameter values relevant for this model are highly uncertain. No data with which to calibrate the model is currently available, and the grouping of many cell types and interactions into single variables and terms makes it difficult to estimate parameter values from existing literature. As the magnitude of the perturbation away from the healthy steady state significantly impacts the timescale on which the cell counts evolve, and the timescale on which pathogenic T cells peak is one of the criteria (Criterion 4 in Table 3.2) around which the model is parameterized, the parameter values resulting from the method we have employed in Section 3.2.3 are clearly dependent on the size of the perturbation. We set the perturbation to be 10^4 pathogenic T cells as this perturbation will cause autoimmune myocarditis to develop for all parameter sets that satisfy Criterion 2 in Table 3.2 and can thus reliably be used to assess the time it takes the T cell response to develop for all parameter sets. Any change to the assumed magnitude of the perturbation may require changes to the parameter values to ensure that Criterion 4 is still met. However, we highlight that the qualitative behaviours of the system are not affected.

Should patient–specific data with which to calibrate the model become available, the calibrated model could be used to estimate the maximum tolerated dose of ICIs. Fitting the model parameters to patient–specific data will enable determination of where the separatrix lies, how far away it is from the healthy steady state, and thus the size of the immune challenge to the healthy steady state a patient can tolerate. Assessing how the separatrix shifts as ICIs are administered will aid in determining the maximum dose of ICIs a patient can receive before the risk of autoimmune myocarditis becomes unacceptable.

With this model, we have taken the first steps in consolidating the current knowledge of autoimmune myocarditis and the impact of ICIs on its development and progression. Though much work remains to be done, we have started on the path towards a mathematical model of autoimmune myocarditis and the role of ICIs in driving the disease, that can be used to determine which cell types, cytokines and other components are critical in the development and progression of this severe side effect. With these essential immune components identified, work can start towards reducing the risk of severe cardiotoxic side effects of ICIs in patients. Furthermore, the balancing of the positive, tumour–treating and negative, cardiotoxic effects of ICIs suggests that there will be ideal treatment regimens for ICIs where the tumour is maximally treated while myocarditis is avoided. In the next chapter, we use optimal control theory to explore what these ideal treatment schedules might look like for different types of ICI therapy.

Chapter 4

Optimal control of autoimmune myocarditis

4.1 Introduction

In Chapter 3, we introduced the first mathematical model of autoimmune myocarditis, and showed that it can reproduce the dynamics of the immune system seen in patients with this disease. Ultimately, the goal is to use the understanding gained from this model to work towards the prevention of autoimmune myocarditis as a side–effect of ICI treatment. One way to do this is to optimize treatment schedules to balance the benefit of ICIs in treating the tumour with the risk of developing autoimmune myocarditis: ICIs stimulate the immune system, which is beneficial when directed at the tumour, but detrimental when the immune reaction targets another tissue or organ, such as the heart. Cardiac antigen–specific T cells are known to circulate in the body of healthy individuals. Furthermore, some tumour–specific T cells activated in the anti–tumour immune response are known to cross–react against cardiac self–antigens [20, 52, 80, 117]. An anti–tumour immune response can thus lead to an increase in cardiac–specific pathogenic T cells circulating throughout the body, and

ICI treatment, by stimulating the immune response, can cause a further increase in the number of these T cells. A higher number of cardiac–specific pathogenic T cells circulating throughout the body means a higher risk of such cells ending up in the heart and potentially causing autoimmune myocarditis to develop.

As a first step towards designing safer ICI treatment schedules, the question we want to answer in this chapter is: can we design a dosing schedule for ICIs such that patients experiencing perturbations in their cardiac immune cell counts during treatment do not develop autoimmune myocarditis as a consequence of ICIs? The goal is to prevent autoimmune myocarditis from developing while the patient continues treatment with ICIs, so that no other interventions, such as anti–inflammatory drugs, are required. We turn to optimal control theory to explore what such treatment schedules might look like.

Optimal control theory uses control variables to optimize the output of dynamical systems by maximizing or minimizing a payoff function [73]. The theory has been used to find optimal regimens for treating leukemia, minimizing autophagy and eradicating gliobastoma cells, among other contexts [88, 119, 120, 129, 136, 138]. In these studies, the control variable is often the concentration of a drug in the patient's system, or the rate at which the drug is administered. The two most common types of controls used in designing optimal treatments are continuous and bang–bang controls. Although these types of controls are relatively straightforward to implement, they lack biological realism. With continuous controls, the control variable is continuous in time and its value at every time point is determined through the application of control theoretic approaches. Having this level of control over drug concentrations or influx rates in the body would be ideal for carefully balancing drug benefit with the risk of side–effects but it is not typically realistic in a clinical setting. Bang–bang controls, on the other hand, allow for the control to be either on or off at each time point by allowing the control variable to switch between two distinct values, usually zero and some positive

number. If the control variable is taken to be the drug concentration in the patient's body, this approach is also unrealistic as drug levels cannot be that tightly controlled. However, if the control variable is instead the influx rate of drug into the patient, this set–up is closer to how treatment is administered in the clinic, with the patient either receiving the drug, or not receiving the drug, at each time point. ICI treatments are generally administered over a period of 30 to 90 minutes, however, so a realistic treatment schedule should not have the control variable be "on" for long periods of time. This causes issues with the use of bang–bang controls because solutions where treatments are "on" for long periods of time are permitted [43, 44, 136].

To obtain realistic treatment schedules, we select an approach in which instead of optimizing the drug concentration or influx rate over time, we optimize the time points at which a drug dose of a set magnitude is administered to the patient instantaneously [23]. This avoids the assumption that a clinician can control the concentration of drug in the patient at all times, and the possibility that the optimal dosing schedules require long periods where the treatment is being administered. Whereas continuous and bang–bang control are widely used, this approach has, to our knowledge, only been used previously in a handful of biomedical contexts [23, 24, 119].

The method developed by Castiglione and Piccoli [23] optimizes dosing time points with high precision. However, as discussed in Chapter 3, our model of autoimmune myocarditis contains much uncertainty, particularly in the parameter values. As a result, we should be careful in placing too much emphasis on the fine details of the optimized dosing schedules and focus instead on the general trends that are predicted. The level of precision provided by Castiglione and Piccoli's method will become relevant as the model of autoimmune myocarditis is developed further and more accurate parameter values are obtained. With the work in this chapter, we anticipate this future need for more precision and develop a framework for schedule optimization in which future models can easily be incorporated.

In Section 4.2, we adapt the mathematical model presented in Chapter 3 for use within an optimal control framework. We previously included the effects of ICIs implicitly by varying model parameter values. Now we include these effects explicitly by defining how nivolumab and ipilimumab impact the immunological processes that drive the development of autoimmune myocarditis, and defining two ODEs that describe the dynamics of ICIs in the body over time. This allows us to model both the pharmacokinetics, *i.e.* how drugs are transformed and transported in the body, and the pharmacodynamics, *i.e.* the physiological effects of the drugs, of ICIs in patients. In Section 4.3, we outline the optimal control method used to solve the problem described above. In Section 4.4, we solve the optimal control problem for nivolumab monotherapy, ipilimumab monotherapy and combination therapy. We finish with a discussion of the results and future avenues of research in Section 4.5.

4.2 The model

The mathematical model of autoimmune myocarditis presented in Chapter 3 describes the evolution of dead/damaged cardiomyocytes (C), innate immune cells (I), pathogenic CD4+ T cells (P) and regulatory T cells (R). It implicitly captures how nivolumab and ipilimumab affect the immunological processes underlying autoimmune myocarditis by varying the parameters a_6, d_4 and d_5, which describe the rates of T cell activation, pathogenic CD4+ T cell deactivation, and regulatory T cell deactivation, respectively, based on the presence of nivolumab and ipilimumab. As nivolumab stimulates T cell activation and inhibits deactivation of pathogenic T cells, the presence of nivolumab was simulated by increasing and decreasing the values of a_6 and d_4, respectively. Ipilimumab also stimulates T cell activation and furthermore stimulates deactivation of regulatory T cells, thus the presence of ipilimumab was simulated by increasing the values of a_6 and d_5.

To track the effects of ICIs on the system as the drugs are administered and cleared, we replace the parameters a_6, d_4 and d_5 by the functions $f_{a_6}(D_N, D_I)$, $f_{d_4}(D_N)$ and $f_{d_5}(D_I)$ to explicitly define how the rates of T cell activation, pathogenic CD4+ T cell deactivation, and regulatory T cell deactivation depend on the presence of ICIs. We define the variable D_N as the mass of nivolumab in the body in grams, and V_N as the dose size of nivolumab in grams. We similarly define the variable D_I as the mass of ipilimumab in the body in grams, and V_I as the dose size of ipilimumab in grams. The changes in D_N and D_I when a dose of ICI is administered are assumed to be instantaneous as the time it takes to administer a dose of nivolumab or ipilimumab (30–90 minutes) is short compared to the length of the treatment window [43, 44, 96].

Nivolumab and ipilimumab each affect multiple separate immunological processes. It is unclear how the amount of the ICIs in the body affects the rate of these different processes. We introduce the parameters γ_i, $i = 1, \ldots, 4$, with units grams^{-1} which govern the pharmacodynamics in the model by translating the mass of ICIs in the body to fractional increases or decreases in the rate parameters of the various immunological processes affected by their presence. We define

$$f_{a_6}(D_N, D_I) = a_6(1 + \gamma_1 D_N)(1 + \gamma_2 D_I), \tag{4.1}$$

$$f_{d_4}(D_N) = d_4/(1 + \gamma_3 D_N), \tag{4.2}$$

$$f_{d_5}(D_I) = d_5(1 + \gamma_4 D_I), \tag{4.3}$$

where a_6, d_4 and d_5 remain the ICI–free rates of T cell activation, pathogenic CD4+ T cell deactivation, and regulatory T cell deactivation, respectively, as defined in Chapter 3. The values of γ_1 and γ_3 represent the impact of nivolumab on T cell activation and pathogenic CD4+ T cell deactivation, respectively, and the values of γ_2 and γ_4 represent the impact of ipilimumab on T cell activation and regulatory T cell deactivation, respectively. While the rarity of autoimmune myocarditis as a side–effect

of ICIs makes estimates of incidence rates uncertain, the incidence rate of autoimmune myocarditis as a side–effect of combination therapy (0.27%) compared to monotherapy (0.06–0.09%) suggests a nonlinear effect on the immunological mechanisms underlying the disease when both ICIs are administered [5, 95]. To reflect this, we multiply the drug terms where these ICIs affect the same process, *i.e.* in the activation of T cells in $f_{a_6}(D_N, D_I)$.

For nivolumab monotherapy, the standard–of–care treatment schedule is one dose every two weeks until disease progression or unacceptable side–effects, for a maximum of 24 months [43]. For ipilimumab monotherapy, the standard–of–care treatment schedule is one dose every three weeks. The duration of treatment with ipilimumab monotherapy depends on the type of cancer, but treatment consists of at least four doses unless the disease progresses or unacceptable side–effects occur [44]. For combination therapy, the standard–of–care treatment schedule is a simultaneous dose of nivolumab and ipilimumab every three weeks for four doses, after which the patient continues on nivolumab monotherapy. The dose of nivolumab is the same for combination therapy and monotherapy, but the dose of ipilimumab in combination therapy is, on average, approximately one tenth of its value in monotherapy [43, 44]. We assume doses are administered instantaneously at time points specified by the treatment schedule.

The number of heart–specific T cells that circulate through the body fluctuates, both naturally and as part of the anti–tumour immune response [20, 52, 80, 117]. These fluctuations may trigger an immune response that leads to autoimmune myocarditis, depending on the magnitude of the fluctuation and the level of ICIs in the patient's system. To simulate this, we inject a set number of pathogenic CD4+ T cells at a given time point.

The model of autoimmune myocarditis with explicit consideration of the effects of nivolumab and ipilimumab and subject to a perturbation in the number of pathogenic

T cells is

$$\frac{\mathrm{d}C}{\mathrm{d}t} = a_1\frac{P^2}{R_1+R} + a_2\frac{I}{R_2+R} - d_1C, \tag{4.4}$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = a_4\frac{P}{R_1+R} + a_5C - d_3I, \tag{4.5}$$

$$\frac{\mathrm{d}P}{\mathrm{d}t} = a_6(1+\gamma_1 D_N)(1+\gamma_2 D_I)\phi\frac{C}{R+R_3} - \frac{d_4}{1+\gamma_3 D_N}P, \text{ with} \tag{4.6}$$

$$P(t=T) = P(t=T^-) + P_{\text{perturb}}$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = a_6(1+\gamma_1 D_N)(1+\gamma_2 D_I)(1-\phi)\frac{C}{R+R_3} - d_5(1+\gamma_4 D_I)R, \tag{4.7}$$

where all parameters are positive and T is the time at which an immune perturbation consisting of P_{perturb} pathogenic T cells is introduced. We assume that the immune perturbation represented by P_{perturb} occurs due to the natural presence of cardiac-antigen specific T cells in otherwise healthy humans, hence it is independent of the level of ICIs in the system [52, 80]. As we now consider the effects of nivolumab and ipilimumab explicitly, we define two ODEs describing the pharmacokinetics of these ICIs

$$\frac{\mathrm{d}D_N}{\mathrm{d}t} = -d_6D_N, \text{ with } D_N(t=t_{i_N,N}) = D_N(t=t_{i_N,N}^-) + V_N, \tag{4.8}$$

$$\frac{\mathrm{d}D_I}{\mathrm{d}t} = -d_7D_I, \text{ with } D_I(t=t_{i_I,I}) = D_I(t=t_{i_I,I}^-) + V_I, \tag{4.9}$$

where $t_{i_N,N}$, $i_N = 1,\ldots,n_N$, and $t_{i_I,I}$, $i_I = 1,\ldots,n_I$, are the time points at which a dose of nivolumab or ipilimumab, respectively, is applied to the system. The parameters d_6 and d_7 are the decay rates of nivolumab and ipilimumab, respectively. The system is closed by the initial conditions $[C(0),\ I(0),\ P(0),\ R(0),\ D_N(0),\ D_I(0)] = [C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}]$.

Parameter values estimated in Chapter 3 are used here. We set $d_6 = \log(2)/25$ days^{-1} as the half–life of nivolumab in the body is 25 days [27]. We set $d_7 = \log(2)/14.7$ days^{-1} as the half–life of ipilimumab in the body is 14.7 days [27]. For

simplicity, we set $\gamma_i = 1$ grams^{-1} for $i = 1, \ldots, 4$. In Appendix A we vary the values of γ_3 and γ_4 to explore the impact of changing the relative effects of ipilimumab and nivolumab on the different immunological processes. The case where the values of all γ_i, $i = 1, \ldots, 4$, are varied simultaneously is covered by the exploration of different values of V_N and V_I, also included in Appendix A. Varying γ_4 alone does not qualitatively change the model behaviour, but varying γ_3 or the values of V_N and V_I does, potentially leading to different outcomes after a perturbation. The full set of parameter values used in this chapter is included in Table 4.1.

In this work, we conclude that myocarditis develops as a consequence of the applied ICI therapy, if, at the final time point of the simulation, all the cell counts in the system are above the cell counts at the ICI–free saddle point. The saddle point has the highest cell counts when the system is ICI–free, therefore if all the cell counts in the system are higher than at the ICI–free saddle point, the system is guaranteed to be in the basin of attraction of the diseased steady state and can never recover to the healthy steady state. This is an overly strict criterion for development of autoimmune myocarditis as autoimmune myocarditis will develop as long as the cell counts in the system lie above the separatrix. However, for ease of implementation this definition is sufficient for our purposes.

The position of the ICI–free saddle point can be computed by solving the polynomial presented in Chapter 3, and reproduced here, for the number of pathogenic

Param.	Definition	Value	Units
a_1	Rate at which pathogenic CD4+ T cells damage/kill cardiomyocytes	80	day^{-1}
a_2	Rate at which innate immune cells damage/kill cardiomyocytes	6000	cells day^{-1}
a_4	Rate at which pathogenic CD4+ T cells activate innate immune cells	500,000	cells day^{-1}
a_5	Rate at which damaged/dead cardiomyocytes activate innate immune cells	35	day^{-1}
a_6	Rate at which damaged/dead cardiomyocytes activate pathogenic CD4+ T cells	750	cells day^{-1}
d_1	Clearance rate of damaged/dead cardiomyocytes	75	day^{-1}
d_3	Innate immune cell death rate	7	day^{-1}
d_4	Pathogenic CD4+ T cell death rate	0.1	day^{-1}
d_5	Regulatory T cell death rate	0.1	day^{-1}
d_6	Degradation rate of nivolumab	$\log(2)/25$	day^{-1}
d_7	Degradation rate of ipilimumab	$\log(2)/14.7$	day^{-1}
R_1	Number of regulatory immune cells for 50% inhibition of pathogenic CD4+ T cell activity	1100	cells
R_2	Number of regulatory immune cells for 50% inhibition of innate immune cells	40000	cells
R_3	Number of regulatory immune cells for 50% inhibition of pathogenic CD4+ T cell activation	2000	cells
ϕ	Fraction activated CD4+ T cells that take on a pathogenic phenotype	0.7	unitless
γ_1	Effect of the mass of nivolumab on T cell activation	1	grams^{-1}
γ_2	Effect of the mass of ipilimumab on T cell activation	1	grams^{-1}
γ_3	Effect of the mass of nivolumab on pathogenic T cell death	1	grams^{-1}
γ_4	Effect of the mass of ipilimumab on regulatory T cell death	1	grams^{-1}

Table 4.1: Definitions, values and units of model parameters.

CD4+ T cells at the steady state, P_*,

$$\begin{aligned}
0 \; &= d_3 K_R^3 P_*^3 \\
&+ \left(d_3(R_1 + R_2 + R_3) K_R^2 - \frac{a_2 a_5}{d_1} K_R^2 - \frac{a_1 K_P}{d_1} d_3 K_R \right) P_*^2 \\
&+ \Bigg[d_3(R_1 R_2 + R_1 R_3 + R_2 R_3) K_R - \frac{a_2 a_5}{d_1}(R_1 + R_3) K_R \\
&\qquad - \frac{a_1 K_P}{d_1} \left(R_2 d_3 - \frac{a_2 a_5}{d_1} \right) - \frac{a_1 a_2 a_5 K_P}{d_1^2} \Bigg] P_* \\
&+ d_3 R_1 R_2 R_3 - \frac{a_2 a_5}{d_1} R_1 R_3 - \frac{a_2 a_4 K_P}{d_1},
\end{aligned} \tag{4.10}$$

where $K_R = d_4(1-\phi)/(d_5\phi)$ and $K_P = a_6\phi/d_4$. The ICI–free steady state cell counts of the remaining variables are then given in terms of P_* by

$$R_* = K_R P_*, \tag{4.11}$$

$$I_* = \frac{\left(a_4 + \frac{a_1 a_5}{d_1} P_* \right) \left(\frac{P_*}{K_R P_* + R_1} \right)}{d_3 - \frac{a_2 a_5}{d_1 (K_R P_* + R_2)}}, \tag{4.12}$$

and

$$C_* = \frac{a_1 P_*^2}{d_1(R_1 + R_*)} + \frac{a_2 I_*}{d_1(R_2 + R_*)}. \tag{4.13}$$

In Section 4.2.1, we show that, given default parameter values and standard–of–care treatment schedules, autoimmune myocarditis develops for certain values of V_N and V_I.

4.2.1 Standard–of–care treatment

Because our aim in this chapter is to set up and demonstrate a pipeline to design ICI treatment schedules that minimize the risk of autoimmune myocarditis, we select the

values of V_N and V_I such that the model predicts that autoimmune myocarditis develops when a standard–of–care treatment schedule is used for nivolumab or ipilimumab monotherapy. For nivolumab the same dose is used in monotherapy and combination therapy. For ipilimumab the dose used in combination therapy is approximately 10% of the dose used in monotherapy [43, 44]. A full exploration of the effects of varying the values of V_N and V_I is included in Appendix A, where we show that for small values of V_N and V_I the standard–of–care dosing schedule does not cause autoimmune myocarditis to develop, but for intermediate and large values of V_N and V_I it does.

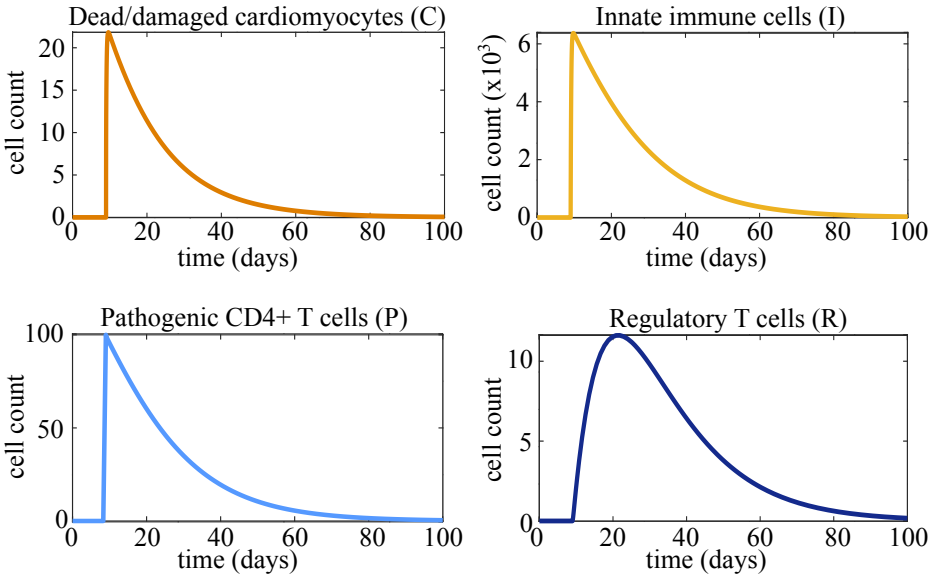


Figure 4.1: Time traces of the model variables in an ICI–free environment. An immune perturbation of 100 pathogenic T cells occurs on day 9. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$, and default parameter values are used (see Table 4.1).

4.2.1.1 Nivolumab monotherapy

For nivolumab monotherapy, the standard–of–care treatment schedule is one dose every two weeks [43]. We assume V_N to be 0.25 grams. We assume doses are administered instantaneously at time points specified by the treatment schedule. We also assume that the heart is free of inflammation and ICIs before the treatment starts, thus $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$. The first dose of nivolumab is

administered on day 2, and then subsequent doses are given every two weeks, with the last dose on day 58, hence five doses are administered in total. On day 9, we introduce an immune perturbation consisting of 100 pathogenic CD4+ T cells (see inset panel in bottom left graph in Fig. 4.2). We note that all cell counts remain zero prior to the perturbation on day 9 due to the inflammation–free initial conditions. The immune perturbation at day 9 simulates the fluctuating number of cardiac–specific T cells that circulate through the body, both naturally and as part of the anti–tumour immune response [20, 52, 80, 117]. We set day 90 to be the final time point of the simulation. A perturbation of this magnitude does not cause autoimmune myocarditis in an ICI–free environment (see Fig. 4.1). However, as Fig. 4.2 shows, when $V_N = 0.25$ grams and nivolumab is administered every two weeks, autoimmune myocarditis develops, as indicated by the high number of inflammatory cells at day 90. In an ICI–free environment, the saddle point lies at $[C,\ I,\ P,\ R] = [138,\ 20000,\ 339,\ 145]$ cells so by our definition of myocarditis presented in Section 4.2, the disease has developed.

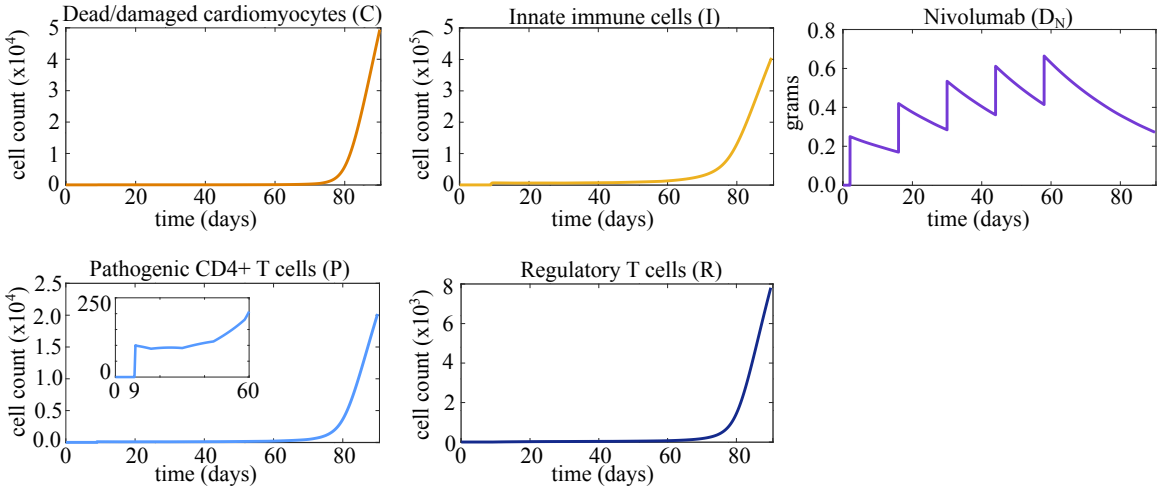


Figure 4.2: When nivolumab is administered every two weeks with $V_N = 0.25$ grams, autoimmune myocarditis develops. Doses are administered on days 2, 16, 30, 44 and 58. An immune perturbation of 100 pathogenic T cells occurs on day 9. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$, and default parameter values are used (see Table 4.1).

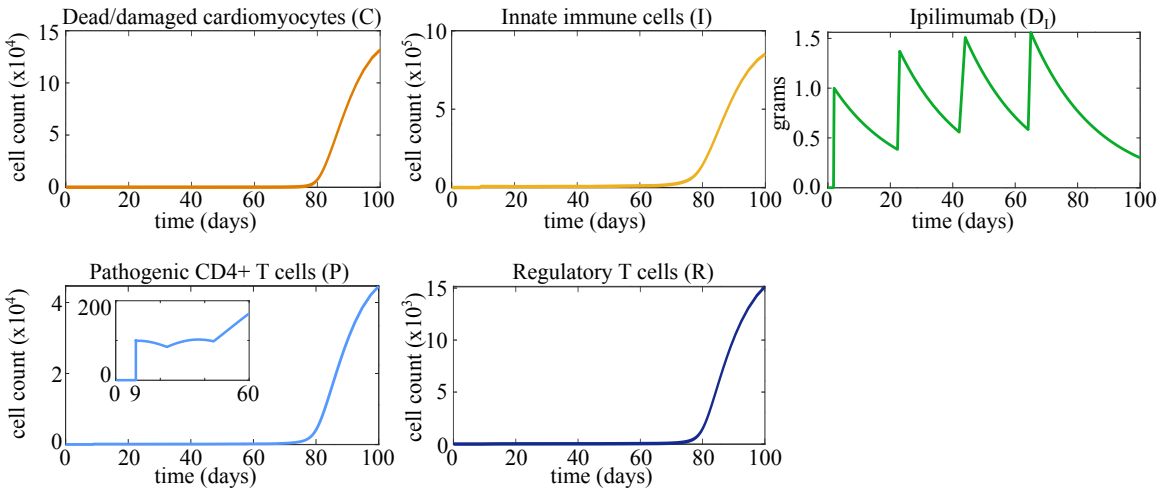


Figure 4.3: When ipilimumab is administered every three weeks with $V_I = 1.0$ gram, autoimmune myocarditis develops. Doses are administered on days 2, 23, 44, 65. An immune perturbation of 100 pathogenic T cells occurs on day 9. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$, and default parameter values are used (see Table 4.1).

4.2.1.2 Ipilimumab monotherapy

For ipilimumab monotherapy, the standard–of–care treatment schedule is one dose every three weeks, with at least four doses administered in total [44]. We assume $V_I = 1.0$ gram. Again, we assume that doses are administered instantaneously at the times specified by the treatment schedule and that the heart is healthy and free of inflammation and ICIs at day 0, captured by the initial conditions $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$. The first dose is administered on day 2, and then every three weeks with the last dose on day 65. On day 9, we introduce an immune perturbation consisting of 100 pathogenic CD4+ T cells (see inset panel in bottom left graph in Fig. 4.3). We set the final time point of the simulation to be day 100. Fig. 4.3 shows that in this scenario, autoimmune myocarditis develops as the cell counts at the final time point of the simulation are well above the ICI–free saddle point at $[C,\ I,\ P,\ R] = [138,\ 20000,\ 339,\ 145]$ cells.

4.2.1.3 Combination therapy

For combination therapy, the standard–of–care treatment schedule is a simultaneous dose of nivolumab and ipilimumab every three weeks for four doses, after which the patient continues on nivolumab monotherapy of which we only simulate the first dose. The dose of nivolumab is the same for combination therapy and monotherapy, but the dose of ipilimumab in combination therapy is, on average, approximately one tenth of what it is in monotherapy [43, 44]. We thus set $V_N = 0.25$ grams and $V_I = 0.1$ grams. As for the monotherapies, we assume the heart is healthy and ICI–free at the start of the simulation, so the initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$. On days 2, 23, 44 and 65, doses of nivolumab and ipilimumab are administered simultaneously. On day 86, a dose of nivolumab only is administered. We simulate an immune perturbation by introducing 100 pathogenic CD4+ T cells on day 9 (see inset panel in bottom left graph in Fig. 4.4). We set the final time point of the simulation to be day 200. Fig. 4.4 shows that in this scenario, autoimmune myocarditis develops as the cell counts at the final time point of the simulation are well above the ICI–free saddle point at $[C,\ I,\ P,\ R] = [138,\ 20000,\ 339,\ 145]$ cells.

For all three types of ICI treatment, these simulations indicate that the patient develops autoimmune myocarditis as a side–effect when standard–of–care treatment is applied with the values of V_N and V_I as stated above. We now turn to optimal control theory to optimize the dosing schedules such that the same number of ICI doses of the same size can be applied without the patient developing autoimmune myocarditis.

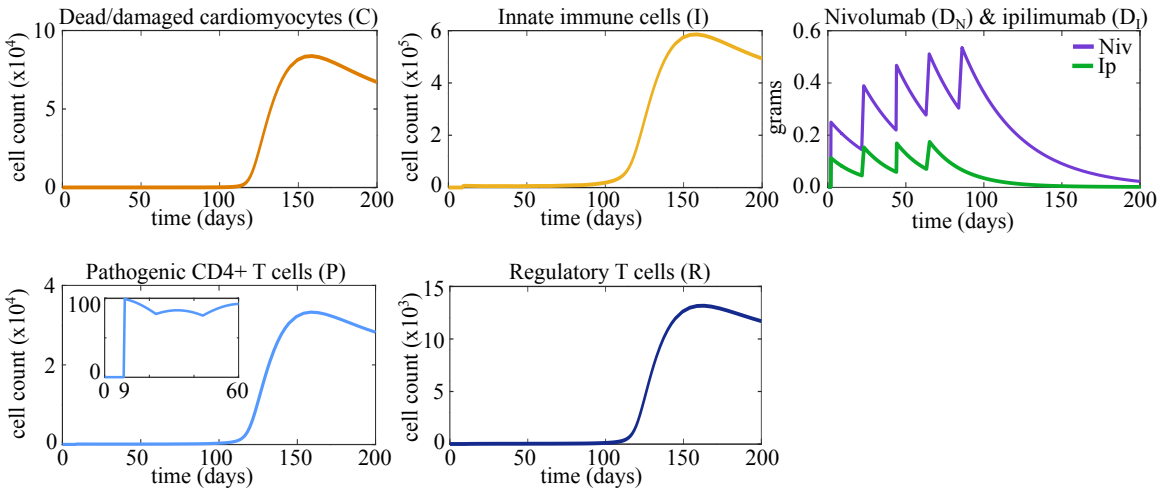


Figure 4.4: When combination therapy is administered with $V_N = 0.25$ grams and $V_I = 0.1$ grams on a standard–of–care treatment schedule, autoimmune myocarditis develops. Doses of nivolumab are administered on days 2, 23, 44, 65 and 86. Doses of ipilimumab are administered on days 2, 23, 44 and 65. An immune perturbation of 100 pathogenic T cells occurs on day 9. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$, and default parameter values are used (see Table 4.1).

4.3 Optimal control method for the timing of dosing schedules

In this section, we set out the algorithm developed by Castiglione and Piccoli that we use to optimize the timing of the ICI doses such that autoimmune myocarditis does not develop [23, 116]. This method relies on finding the optimal schedule by minimizing a payoff function through gradient descent, and applies to a system of ODEs of the form

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \mathbf{f}(\mathbf{x}, t) + \mathbf{g}(\mathbf{x}, t)\mathbf{u}(t), \qquad \mathbf{x}(0) = \mathbf{x_0}, \tag{4.14}$$

where $\mathbf{x}$ are the state variables and $\mathbf{u}$ are the control variables. For the purpose of optimizing ICI treatment schedules to prevent autoimmune myocarditis, we define the function $\mathbf{f}(\mathbf{x}, t)$ as the RHS of Eq. (4.4)–(4.9), the vector $\mathbf{x} = [C, I, P, R, D_N, D_I]$ and

the control $\mathbf{u}(t) = \left[0, 0, 0, 0, \sum_{i_N=1}^{n_N} V_N\ \chi[t_{i_N,N}, t_{i_N,N} + \eta], \sum_{i_I=1}^{n_I} V_I\ \chi[t_{i_I,I}, t_{i_I,I} + \eta]\right]$ where n_N is the number of nivolumab doses in the schedule, n_I is the number of ipilimumab doses in the schedule, χ is an indicator function which is equal to one when $t_{i_\alpha,\alpha} \leq t \leq t_{i_\alpha,\alpha} + \eta$ where $\alpha \in \{N, I\}$ and zero everywhere else, and η is the time span over which an ICI dose is administered, which is considered to be negligibly small [23]. Since we consider η to be negligibly small, we will assume instantaneous dosing at the time points specified in the dosing schedule for the remainder of this work. Lastly, we set $\mathbf{g}(\mathbf{x}, t) = 1$. The payoff function takes the form

$$J = \Psi(\mathbf{x}(t_f, \mathbf{u}(t))) + \int_0^{t_f} L(\mathbf{x}(t, \mathbf{u}(t)), \mathbf{u}(t))\mathrm{d}t, \tag{4.15}$$

where t_f is the final time point of the simulation. The payoff function can capture both how the state and control variables evolve over time in the function $L(\mathbf{x}(t, \mathbf{u}(t)), \mathbf{u}(t))$, as well as the state of the system at the final time point of the simulation via the function $\Psi(\mathbf{x}(t_f, \mathbf{u}(t)))$. The dosing schedule $S = [t_{1,N}, \ldots, t_{n_N,N}, t_{1,I}, \ldots, t_{n_I,I}]$ specifies the time points at which doses of ICIs are administered. If a time point appears in S multiple times, this means that multiple doses of ICIs are applied at the same time point. In the case of monotherapy, the same ICI is applied at all time points. For combination therapy, each time point is linked to a specific type of ICI.

Below we set out the algorithm for optimizing dosing schedules that consist of n_N doses of nivolumab and n_I doses of ipilimumab. The algorithm proceeds as follows:

1. Set the ICI dose sizes V_N and V_I, the initial dosing schedule $S_0 = [t_{1,N}, \ldots, t_{n_N,N}, t_{1,I}, \ldots, t_{n_I,I}]$ for $n = n_N + n_I$ total doses to be administered, the length of the treatment window $[t_{\min}, t_{\max}]$, the final time point of the simulation t_f, the gradient descent step size parameter h, the initial conditions $\mathbf{x_0}$, and the convergence threshold for the payoff c_p, which determines the maximum absolute difference in payoffs between iterations before the algorithm is

considered to have converged.

2. Solve the ODEs in Eq. (4.4)–(4.9) using the current schedule $S = [t_{1,N}, \ldots, t_{n_N,N}, t_{1,I}, \ldots, t_{n_I,I}]$ (S_0 if the first iteration) to obtain the time traces $\mathbf{x}(t)$ of the state variables.

3. For each $t_{i_\alpha,\alpha}$ in S where $\alpha \in \{N, I\}$, $i_N = [1, \ldots, n_N]$, and $i_I = [1, \ldots, n_I]$, solve the variation equations over the time span $[t_{i_\alpha,\alpha},\ t_f]$, where the variation equations are defined as

$$\begin{aligned} \frac{\mathrm{d}\mathbf{v}_{i_\alpha}}{\mathrm{d}t} &= \nabla\mathbf{f}(\mathbf{x}(t))\ \mathbf{v}_{i_\alpha}(t), \\ \mathbf{v}_{i_\alpha}(t^-_{i_\alpha,\alpha}) &= \mathbf{f}(\mathbf{x}(t^-_{i_\alpha,\alpha})) - \mathbf{f}(\mathbf{x}(t^-_{i_\alpha,\alpha}) + \mathbf{e}_V), \\ \frac{\mathrm{d}\nu_{i_\alpha}}{\mathrm{d}t} &= \nabla L(\mathbf{x}(t), \mathbf{u}(t)) \cdot \mathbf{v}_{i_\alpha}(t), \\ \nu_{i_\alpha}(t^-_{i_\alpha,\alpha}) &= 0, \end{aligned} \tag{4.16}$$

where the matrix $\nabla\mathbf{f}$ and the vector ∇L denote the Jacobians of $\mathbf{f}$ and L, respectively, and $\mathbf{x}(t^-_{i_\alpha,\alpha})$ denotes the state of the system right before the dose is given. Specifically, we define $t^-_{i_\alpha,\alpha} = t_{i_\alpha,\alpha} - 10^{-10}$. The vector $\mathbf{e}_V$ denotes the ICI dose that is given at time $t_{i_\alpha,\alpha}$, *i.e.* $\mathbf{e}_V = [0,\ 0,\ 0,\ 0,\ V_N,\ 0]$ if $\alpha = N$, and $\mathbf{e}_V = [0,\ 0,\ 0,\ 0,\ 0,\ V_I]$ if $\alpha = I$.

4. For each $t_{i_\alpha,\alpha}$ in S, where $\alpha \in \{N, I\}$, $i_N = [1, \ldots, n_N]$, and $i_I = [1, \ldots, n_I]$, compute the direction of steepest descent as

$$\left.\frac{\mathrm{d}\phi}{\mathrm{d}t}\right|_{i_\alpha,\alpha} = \nabla\Psi(t, \mathbf{x}(\mathbf{t_f}, \mathbf{u})) \cdot \mathbf{v}_{i_\alpha}(t_f) + \nu_{i_\alpha}(t_f), \tag{4.17}$$

where the vector $\nabla\Psi$ denotes the Jacobian of Ψ.

5. Update the treatment schedule S as

$$S_{\text{new}} = S - h\frac{\mathrm{d}\phi/\mathrm{d}t}{||\mathrm{d}\phi/\mathrm{d}t||}, \tag{4.18}$$

where

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = \left[\left.\frac{\mathrm{d}\phi}{\mathrm{d}t}\right|_{1,N}, \ldots, \left.\frac{\mathrm{d}\phi}{\mathrm{d}t}\right|_{n_N,N}, \left.\frac{\mathrm{d}\phi}{\mathrm{d}t}\right|_{1,I}, \ldots, \left.\frac{\mathrm{d}\phi}{\mathrm{d}t}\right|_{n_I,I}\right] \tag{4.19}$$

and h is the step size parameter.

6. Compute the payoff of the current schedule. If the absolute difference between the current payoff and the payoff of the previous iteration is less than c_p, terminate the algorithm. If not, return to Step 2 with $S = S_{\text{new}}$.

In Step 5, $\mathrm{d}\phi/\mathrm{d}t$ is divided by its norm to improve convergence: the values in the vector $\mathrm{d}\phi/\mathrm{d}t$ can vary widely, so dividing by the norm controls the size of the steps taken in the gradient descent algorithm. Other options for optimizing the gradient descent algorithm in Step 5 can be found in [30], however we found the approach in Eq. (4.18) sufficient from a computational perspective.

To ensure that all ICI doses are administered within the desired treatment window, we adapt the algorithm from Castiglione and Piccoli [23] by implementing projected gradient descent in Step 5. This entails projecting any time points in S_{new} that fall outside of the window $[t_{\min}, t_{\max}]$ to the nearest time point within the window. Thus, if any time point in S_{new} is before the first day, $t_{\min}$, or after the last day, $t_{\max}$, of the treatment window, we set those doses to be administered on the first or on the last day, respectively. Within the treatment window, dosing times can vary continuously.

4.4 Optimized ICI treatment

In this section, we optimize the treatment schedules for nivolumab monotherapy, ipilimumab monotherapy and combination therapy for a cancer patient with a healthy,

ICI–free heart pre–treatment. The question we seek to answer is: if this patient experiences a perturbation in the number of cardiac–specific pathogenic CD4+ T cells while on standard–of–care treatment with ICIs, can we adjust the scheduling of the subsequent doses of ICIs such that the patient does not develop autoimmune myocarditis?

The payoff landscape explored by the gradient descent algorithm is large and high dimensional. It is thus difficult to ascertain whether the algorithm has converged to the global minimum (which represents the optimal dosing schedule). However, since our primary aim is to identify a schedule which prevents autoimmune myocarditis, which may be achieved by a schedule at or close to a local minimum, a sufficiently optimized, though not optimal, schedule will thus suffice to achieve our objective.

For all types of treatment, the default parameter values established in Chapter 3 are used, and we set the step size parameter $h = 0.5$. The effect of varying h for each of the ICI therapies is explored in Appendix A, where we show that varying the value of h does not have a qualitative impact on the results of the optimizations included in this chapter. We set the convergence threshold for the payoff to be $c_p = 0.01$. The payoff function to be minimized for each case is

$$J = C(t_f). \tag{4.20}$$

This payoff function captures the need to keep the number of dead/damaged cardiomyocytes low in order to prevent autoimmune myocarditis. Only the number of dead/damaged cardiomyocytes at the final time point of the simulation is considered, because our main concern is whether myocarditis develops or not and the count of dead/damaged cardiomyocytes at the final time point is sufficient to indicate this. A high count of dead/damaged cardiomyocytes indicates that the system is above the separatrix, and autoimmune myocarditis therefore either has already developed or

will develop in the future. In Appendix A.1, we show that the results do not change significantly when we instead integrate the number of dead/damaged cardiomyocytes over time so that the payoff function becomes

$$J = \int_0^{t_f} C(t)\mathrm{d}t, \tag{4.21}$$

i.e. we count the total number of dead/damaged cardiomyocytes.

4.4.1 Nivolumab monotherapy

We now turn to optimizing the treatment schedule for nivolumab monotherapy using the algorithm presented in Section 4.3. We again use $V_N = 0.25$ grams, as in Section 4.2.1.1. We assume the heart is healthy and ICI–free on day 0, so the initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$. The first dose of nivolumab is administered on day 2. On day 9 an immune perturbation consisting of 100 pathogenic T cells is introduced. Because detection of autoimmune myocarditis, and subsequent adjustment of the dosing schedule, is only possible after the immune perturbation has taken place and its consequences have been detected, the first dose of nivolumab is fixed to occur on day 2 and is not part of the optimization. We optimize the timing of the subsequent four doses, which are initialized at the standard–of–care schedule $S_0 = [16,\ 30,\ 44,\ 58]$. We choose to consider the first five doses of the treatment course because the median time to onset of myocarditis symptoms for patients receiving ICI treatment is on the order of a month after the first dose of ICIs is administered and considering the first five doses of nivolumab monotherapy covers this time frame [108].

The four doses which are optimized can be administered from day $t_{\min} = 10$ to day $t_{\max} = 60$. The lower limit on the treatment window comes from the fact that it is only necessary to change the dosing schedule if an immune perturbation occurs and

causes immune cell counts to rise to such an extent that it might cause autoimmune myocarditis. The choice to optimize the dosing schedule thus relies on the detection of an immune perturbation. With the immune perturbation set to occur on day 9, the earliest we could possibly respond to the detection of high levels of immune cell counts or the patient showing symptoms, is on day 10. This assumes that onset of symptoms is very rapid or the patient happens to be tested for high levels of immune cells right after the immune perturbation, which may not be true. As such, in Appendix A.1.7 we show that if the perturbation occurs much later in treatment, or the consequences of the perturbation are detected later during treatment, we can still identify optimized dosing schedules that do not cause autoimmune myocarditis. We set the final day of the treatment window to be day 60 so that the optimized doses are administered approximately within the same time span as the standard–of–care schedule. We set the final time point of the simulation to be day 160 ($t_f = 160$) to ensure that if a dose of nivolumab administered on day 60 causes myocarditis to develop, this is captured in the payoff function. In this chapter, the immune perturbation thus always occurs between the first and second dose, and the timing of every dose after the first one can be adjusted.

Fig. 4.5 shows how the dosing schedule and payoff change over the iterations of the control algorithm as the nivolumab monotherapy schedule is optimized. Fig. 4.6 shows the time traces produced by the model when the optimized dosing schedule is applied. After 69 iterations, convergence is achieved and the final dosing schedule is $S = [42.1,\ 44.9,\ 51.6,\ 59.1]$.

The left–hand panel in Fig. 4.5 shows that the optimization of the dosing schedule primarily delays the administration of nivolumab after the immune perturbation occurs. After the first dose of nivolumab on day 2 and the perturbation to the number of pathogenic T cells, P, this delay gives the system time to recover, and the numbers of inflammatory cells, I, P and R, to drop before the next dose is applied. As discussed

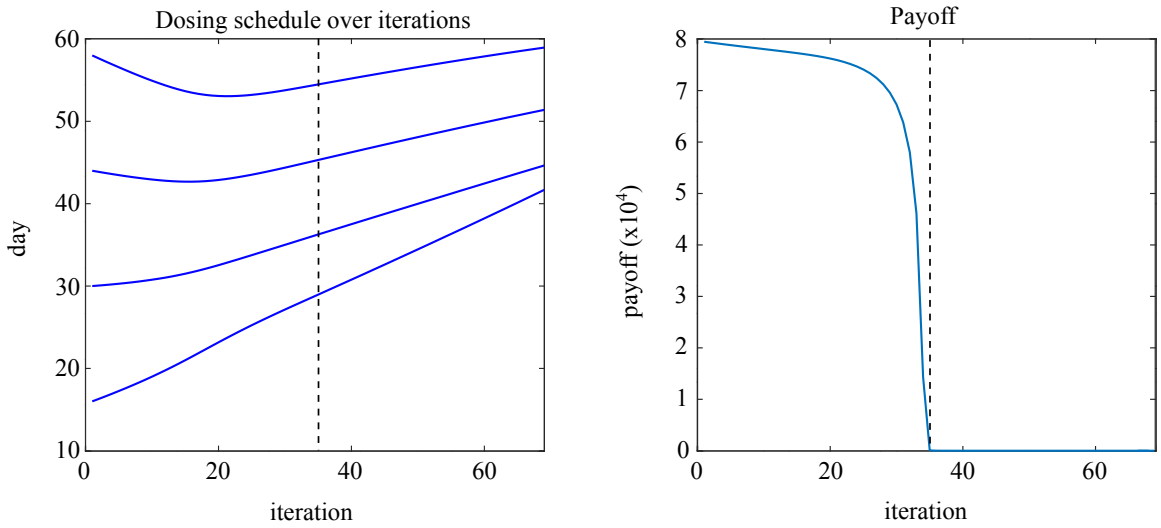


Figure 4.5: The dose schedule for nivolumab monotherapy and payoff evolve over iterations of the optimization algorithm. The dashed line indicates the iteration at which a myocarditis–free schedule is first found. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$, with default parameter values used (see Table 4.1).

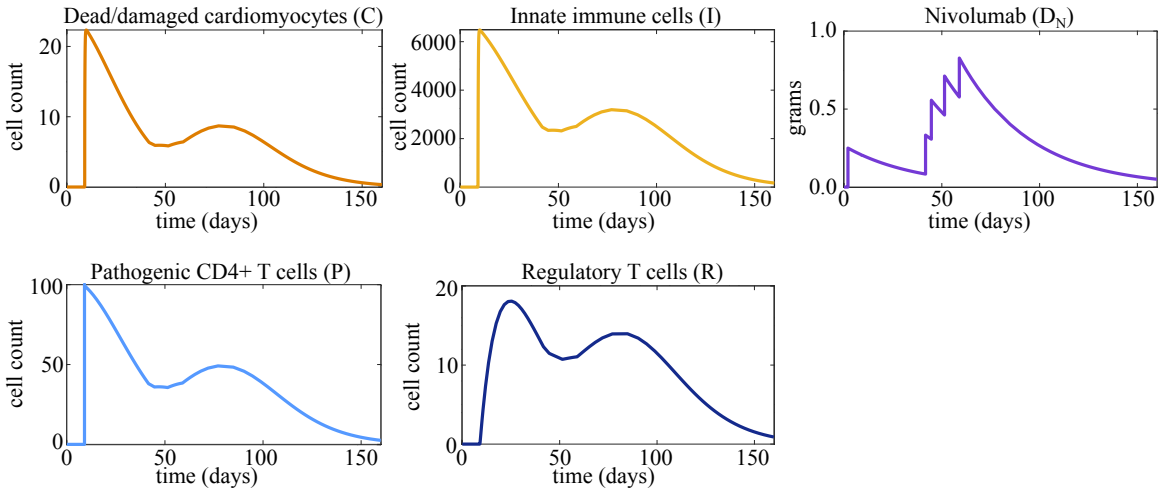


Figure 4.6: Time traces of state variables when the optimized dosing schedule for nivolumab monotherapy is applied. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$, with default parameter values used (see Table 4.1).

in Chapter 3, the presence of ICIs in the system brings the separatrix closer to the origin. If the system crosses the separatrix, autoimmune myocarditis develops and it is difficult for the system to return to the healthy steady state at the origin. Thus, with immune cell counts high after the immune perturbation, another dose of nivolumab administered too soon could push the separatrix between the healthy steady state at the origin and the current state of the system, thus causing immune cell counts to rise further as autoimmune myocarditis develops. With a delay before the next dose of nivolumab, the immune cell counts can decrease further, thus decreasing the risk of pushing the separatrix below the current state of the system. This principle is illustrated in Fig. 4.7 where in the left–hand panel the dosing schedule causes the saddle point (dashed line) to be lowered too early, causing the system (represented by the pathogenic CD4+ T cells (light blue)) to cross the separatrix and autoimmune myocarditis to develop. In the right–hand panel, the dosing schedule is adjusted such that myocarditis does not develop. Here we see that the third dose of nivolumab does cause the saddle point to drop below the current count of pathogenic CD4+ T cells, but at that time the levels of inflammation are much lower compared to when the system crosses the separatrix in the left–hand panel. Because levels of inflammation are lower, the decay of nivolumab is sufficiently rapid and the increase in pathogenic CD4+ T cells sufficiently slow so that the cell counts at the saddle point rise faster than the cell counts present in the system, causing a return to the healthy steady state at the origin as the separatrix crosses over the system around $t = 80$ days.

In the optimized dosing schedule, all four doses of nivolumab are administered between day 40 and day 60, the last day of the treatment window, as shown in Fig. 4.6. This rapid succession of doses does not cause autoimmune myocarditis, because levels of inflammation are sufficiently low. It is questionable, however, whether administering four doses of nivolumab so quickly after each other is wise in terms of tumour treatment and the risk of other potential side–effects. The right–hand panel in Fig. 4.5

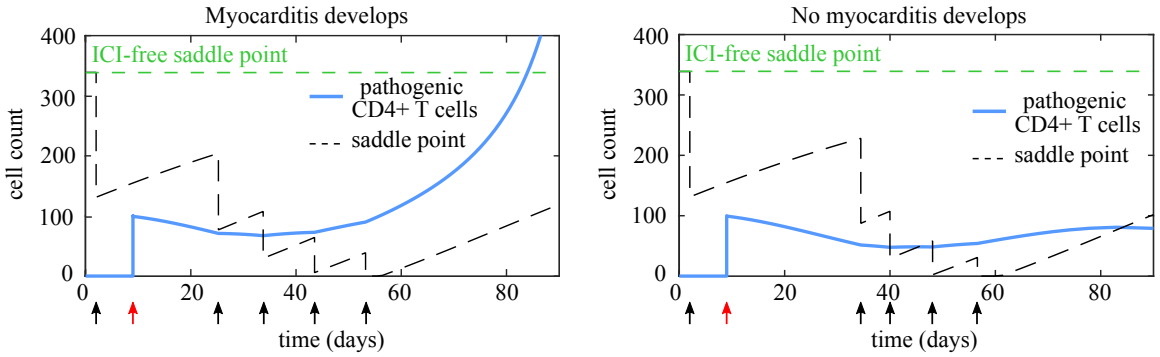


Figure 4.7: Time traces of pathogenic CD4+ T cells for nivolumab dosing schedules which do (left, obtained after 25 iterations) and do not (right, obtained after 50 iterations) cause autoimmune myocarditis to develop. Note that at the final time point of the simulation ($t_f = 160$, not shown here) in the left–hand panel the immune cell counts exceed those at the ICI–free saddle point (dashed green line), hence we conclude that autoimmune myocarditis will develop in this case. The dashed line indicates the number of pathogenic CD4+ T cells at the saddle point. The black arrows indicate time points at which a dose of nivolumab is applied. The red arrows indicate the time point at which the immune perturbation occurs. All default parameter settings are used (see Table 4.1).

shows the evolution of the payoff over the iterations. During the first 35 iterations, the payoff is significantly reduced, after which the decrease in payoff with each iteration is small and the payoff profile becomes almost flat. Most importantly, for all schedules with iteration numbers above 34 the patient remains free of autoimmune myocarditis.

Table 4.2 shows that if more iterations are allowed the dosing schedule evolves such that eventually all doses are administered on the final day of the treatment window, day 60. Note that all schedules in Table 4.2 are better than the standard–of–care schedule as indicated by the lower payoff value. The results in this table show that there is a wide variety of schedules which do not cause autoimmune myocarditis to develop. A clinician would thus have a wide range of treatment schedules to choose from, all of which prevent myocarditis. This flexibility in schedules would allow them to take in account other factors, such as other side–effects, efficacy in tumour reduction, patient variability, etc. in treatment decisions. To illustrate this further,

time traces of the state variables when the schedules of iterations 50, 100, and 150 are applied are included in Appendix A.1.

The default size of the immune perturbation is 100 pathogenic CD4+ T cells. Through simulations (not shown here) we find three different regimes when the number of cells in the perturbation is varied. For 83 pathogenic CD4+ T cells or fewer, autoimmune myocarditis does not develop under the standard–of–care schedule, which is consistent with our hypothesis stated in Chapter 3 that the body should be able to control small perturbations in the immune system. The threshold for the size of the immune perturbation that can be controlled is lower when standard–of–care ICI treatment is administered compared to an ICI–free environment as the presence of ICIs lowers the threshold for triggering autoimmune myocarditis. If the immune perturbation consists of 84–150 pathogenic CD4+ T cells, autoimmune myocarditis develops when a standard–of–care treatment schedule is applied but an optimized dosing schedule for which autoimmune myocarditis is prevented can be identified by the algorithm presented in this chapter. If more than 150 pathogenic CD4+ T cells are introduced, no dosing schedule can be identified (within the constraints of the default parameter settings) which prevents autoimmune myocarditis from developing. In this scenario, the stimulation of the immune system is so strong that myocarditis develops regardless of when subsequent doses are administered.

In Appendix A.1 we show that the results of the optimization do not change qualitatively as the value of the step parameter h is varied. Increasing the length of the treatment window does affect the optimized dosing schedule identified by the algorithm, but it does not lead to a significant decrease in the optimized payoff. We also show that as the value of V_N is varied we see three regimes: (1) for small values of V_N the standard–of–care dosing schedule does not cause autoimmune myocarditis to develop; (2) for intermediate values of V_N, including the default value used in this work, the standard–of–care schedule does cause autoimmune myocarditis but opti-

Iteration	Schedule	Payoff
10	19.0, 30.8, 42.9, 54.9	7.8076×10^4
20	23.2, 32.5, 42.9, 53.1	7.6233×10^4
50	34.5, 40.0, 48.1, 56.6	0.7491
69	42.1, 44.9, 51.6, 59.1	0.3392
100	53.0, 53.2, 55.9, 60.0	0.1964
150	60.0, 60.0, 60.0, 60.0	0.1561

Table 4.2: The schedule of nivolumab doses and the payoff vary over the iterations. After iteration 69, the convergence threshold is reached. After 35 iterations, a schedule is first achieved for which autoimmune myocarditis does not develop.

mized dosing schedules that prevent the disease can be found; and (3) for high values of V_N no dosing schedules can be found that do not cause autoimmune myocarditis with the default settings used here. Lastly, we show that if the immune perturbation occurs later on in treatment, or the consequences of the immune perturbation are detected later and we optimize doses four through seven instead of doses two through five, an optimized schedule which prevents autoimmune myocarditis can still be identified.

4.4.2 Ipilimumab monotherapy

We approach the optimization of the ipilimumab monotherapy similarly to the nivolumab monotherapy in the previous section. We set $V_I = 1.0$ gram, as in Section 4.2.1.2. The heart is assumed to be healthy and ICI–free on day 0, so the initial conditions for the simulations are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$. The first dose of ipilimumab is administered on day 2, and an immune perturbation consisting of 100 pathogenic CD4 + T cells is introduced on day 9. Only the ICI doses administered in standard–of–care after the immune perturbation on day 9 are optimized, so $S_0 = [23,\ 44,\ 65]$. To accommodate the dosing schedule of ipilimumab, we set the

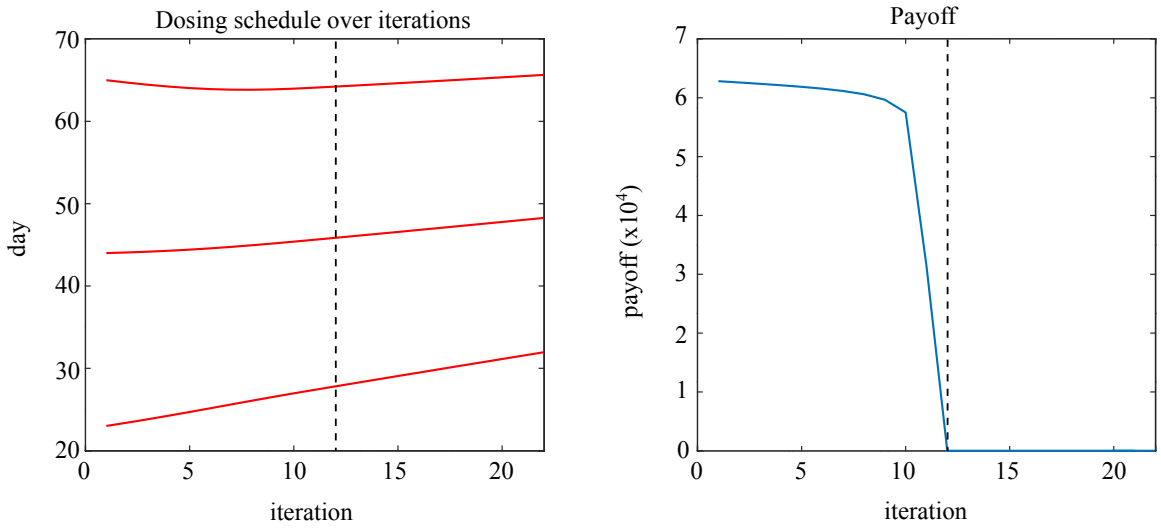


Figure 4.8: Evolution of the dosing schedule schedule and payoff over the iterations as the ipilimumab monotherapy treatment is optimized. The dashed line indicates the iteration at which a myocarditis–free schedule is first found. Default parameter values are used (see Table 4.1).

treatment window from day $t_{\min} = 10$ through day $t_{\max} = 70$ and set the final time point of the simulation to be day $t_f = 170$.

Fig. 4.8 shows how the dosing schedule and payoff evolve over the iterations as we optimize the schedule of ipilimumab monotherapy. The optimized schedule of treatment is $S = [32.4,\ 48.5,\ 65.8]$. Fig. 4.9 shows the time traces of the state variables when the optimized dosing schedule is applied. Similarly to the nivolumab results, the main difference between this optimized schedule and the standard–of–care schedule is the delay of the ipilimumab doses after the immune perturbation occurs. This allows the number of immune cells in the heart to decrease before the next dose of ipilimumab lowers the separatrix, thus decreasing the risk of developing myocarditis. When this delay is included, the model predicts that the remaining doses can be administered in quicker succession, *i.e.* with less than three weeks between doses, without causing autoimmune myocarditis to develop.

The right–hand panel in Fig. 4.8 shows that the payoff function decreases steeply over the first 12 iterations, after which the profile flattens out. Iteration 12 is the

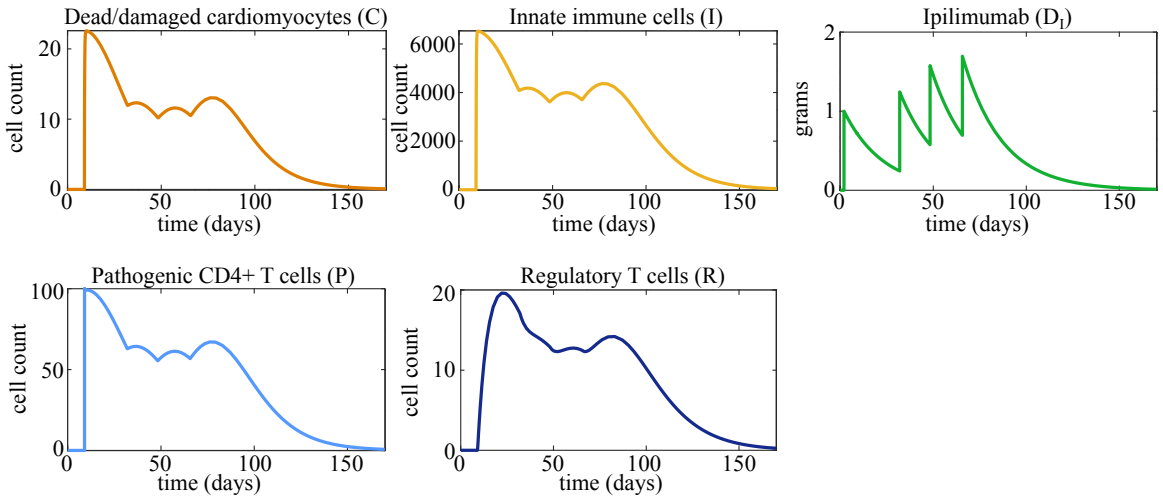


Figure 4.9: Time traces of state variables when the optimized dosing schedule for ipilimumab monotherapy is applied. Default parameter values are used (see Table 4.1).

Iteration	Schedule	Payoff
10	27.0, 45.4, 64.0	1.9032×10^5
20	31.1, 47.8, 65.3	1.3430×10^5
22	32.4, 48.5, 65.8	0.0879
50	43.0, 55.5, 70.0	0.0293
200	70.0, 70.0, 70.0	0.0181

Table 4.3: The schedule of ipilimumab doses and the payoff vary over the iterations. After iteration 22, the convergence threshold is met. At iteration 12, a schedule is first achieved for which autoimmune myocarditis does not develop.

first iteration at which a myocarditis–free schedule is found, and all schedules found in iterations thereafter do not cause autoimmune myocarditis. Table 4.3 shows the schedules and payoffs after a range of iterations both before and after the convergence threshold is met, showing the wide variety of treatment schedules that attain the objective of preventing autoimmune myocarditis. For ipilimumab monotherapy there is thus also room to consider other factors, such as the effect on the tumour and other possible side–effects when designing the optimal treatment plan.

An exploration of the effects of varying the values of the parameters V_I, h, and the length of the treatment window, is included in Appendix A.2. As for nivolumab

monotherapy, we show that the results of the optimization are qualitatively robust to the changes in the step parameter h, that increasing the length of the treatment window affects the optimized dosing schedule identified by the algorithm but that this does not lead to a significantly lower payoff, and that we see the same three possible outcome scenarios when V_I is varied as we saw when V_N is varied.

4.4.3 Combination therapy

We now consider combination therapy with both nivolumab and ipilimumab. We approach the optimization of this type of treatment similarly to the previous sections. The dose of nivolumab is the same in monotherapy and combination therapy so we set $V_N = 0.25$ grams. The dose of ipilimumab is approximately ten times lower than in monotherapy so we set $V_I = 0.1$ grams [43, 44]. The standard–of–care treatment starts with four combined doses of both ICIs which are administered every three weeks starting on day 2, and then continues with one further dose of nivolumab monotherapy. Not including the dose on day 2, which is not optimized, the initial schedule, S_0, thus consists of an initial schedule of doses of nivolumab, $S_{0,N}$, which are administered on days $S_{0,N} = [23, 44, 65, 86]$, as in standard–of–care treatment, and an initial schedule of doses of ipilimumab, $S_{0,I}$, which are administered on days $S_{0,I} = [23, 44, 65]$, as in standard–of–care treatment. The optimization algorithm can vary each dose individually, *i.e.* the doses of ipilimumab and nivolumab that are administered on the same day in the initial schedule need not be administered on the same day in the optimized schedule. We do not optimize the doses on day 2 as these fall before the immune perturbation of 100 pathogenic T cells on day 9. We set the treatment window from day $t_{\min} = 10$ through day $t_{\max} = 90$ and set the final time point of the simulation to be 100 days after the end of the treatment window on day $t_f = 190$. The initial conditions for the simulations are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0]$, representing the healthy, ICI–free heart before treatment starts.

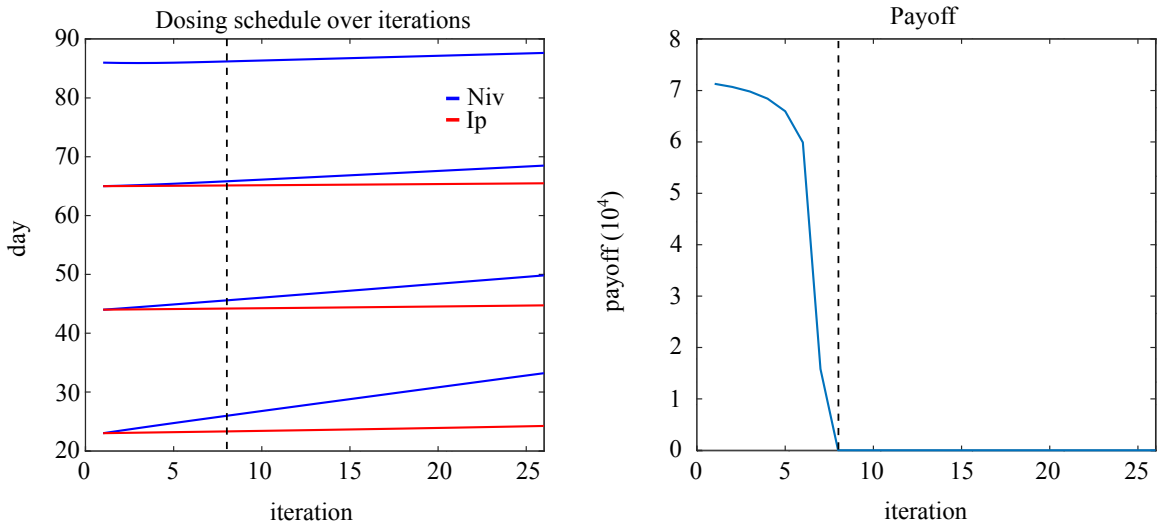


Figure 4.10: Evolution of the dosing schedule and payoff over the iterations as we optimize combination therapy treatment. In the left–hand plot, doses of nivolumab are shown in blue and doses of ipilimumab are shown in red. The dashed line indicates the iteration at which a myocarditis–free schedule is first found. Default parameter values are used (see Table 4.1).

Fig. 4.10 shows how the dosing schedule and payoff evolve over the iterations as we optimize the dosing schedule of combination therapy. The optimized schedule is $S_N = [33.6,\ 50.1,\ 68.6,\ 87.7]$ for nivolumab, and $S_I = [24.3,\ 44.8,\ 65.5]$ for ipilimumab. Fig. 4.11 shows the time traces of the state variables when the optimized dosing schedule is applied. The pattern of evolution of the dosing schedules is qualitatively similar to what we observed for the monotherapies. The left–hand panel in Fig. 4.10 shows that all ICI doses in the schedule are delayed in the optimized schedule compared to the standard–of–care schedule. This delay is significantly larger for nivolumab doses compared to the ipilimumab doses which are initialized on the same day.

Similarly to what we saw with the other types of ICI therapy, the first eight iterations of the optimal control algorithm make significant progress in reducing the payoff and then the payoff profile flattens out (see right–hand panel in Fig. 4.10). Iteration 8 is the first iteration at which a myocarditis–free schedule is found. Table 4.4 shows the schedules and payoffs after a range of iterations, both before and

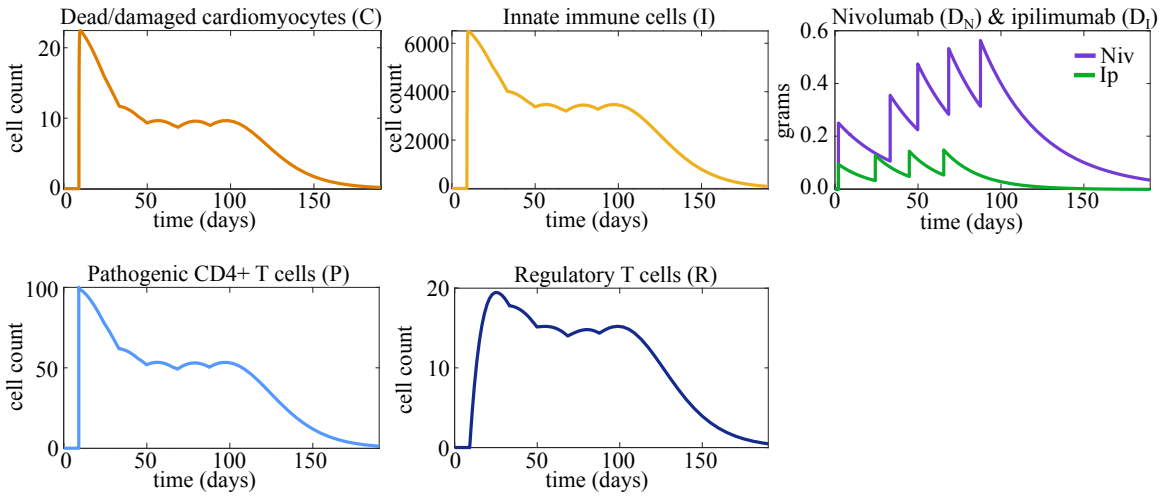


Figure 4.11: Time traces of state variables when the optimized dosing schedule for combination therapy is applied. Default parameter values are used (see Table 4.1).

after convergence is achieved. There is a wide variety of dosing schedules available that reduce the risk of myocarditis while still administering the same number of doses of ICIs as the standard–of–care treatment for combination therapy. For combination therapy, similar to the monotherapies, there is thus room to consider other factors besides prevention of autoimmune myocarditis when designing treatment for a specific patient.

An exploration of the effects of varying the values of V_N, V_I, h, and the length of the treatment window is included in Appendix A.3. As before, we can show that the results of the optimization are qualitatively robust to changes to the step parameter h, and that increasing the length of the treatment window qualitatively changes the optimized dosing schedule but that this does not lead to a significantly lower payoff. As V_N and V_I are varied simultaneously, the three scenarios again occur where either standard–of–care treatment does not cause autoimmune myocarditis, standard–of–care treatment does cause autoimmune myocarditis but an optimized schedule does not, or no schedule can be identified where autoimmune myocarditis does not occur.

Iteration	Schedule D_N	Schedule D_I	Payoff
5	24.7, 44.9, 65.4, 86.0	23.2, 44.1, 65.0	6.5973×10^4
10	26.8, 46.0, 66.1, 86.4	23.4, 44.2, 65.2	1.2874
20	30.8, 48.4, 67.6, 87.2	23.9, 44.5, 65.4	0.2463
26	33.6, 50.1, 68.6, 87.7	24.3, 44.8, 65.5	0.1688
50	42.7, 55.5, 72.2, 89.5	25.9, 45.6, 66.1	0.0811
100	63.5, 66.2, 79.0, 90.0	30.3, 46.6, 67.6	0.0449
150	80.6, 80.7, 82.5, 89.0	34.7, 47.9, 66.5	0.0326

Table 4.4: The schedule of nivolumab doses, the schedule of ipilimumab, and the payoff vary over the number of iterations. After iteration 26, the convergence threshold is met. At iteration 8, a schedule which does not cause autoimmune myocarditis is first achieved.

4.5 Discussion

In this chapter we have shown that by applying optimal control theory to our model of autoimmune myocarditis, we can improve treatment with ICIs such that autoimmune myocarditis does not develop. The method we have used optimizes the timing of ICI doses, rather than directly controlling the ICI levels in the patient, which makes for a more biologically and clinically realistic scenario.

For all types of ICI therapy, monotherapy with nivolumab or ipilimumab or combination therapy with both, we showed that autoimmune myocarditis does not develop when ICIs are administered over a large range of dosing schedules whereby the patient receives the same number of ICI doses as the standard–of–care treatment regime. Compared to the standard–of–care dosing schedule, the common trend among all the myocarditis–free schedules is a delay in dosing with ICIs after the immune perturbation occurs to allow the number of immune cells in the heart to decrease. As shown in Tables 4.2–4.4, allowing more iterations after the convergence threshold is met leads to optimized dosing schedules where all doses are administered on the last day of the treatment window. A dosing schedule where all doses are administered as late as

possible is an intuitive optimal solution to the problem presented in this chapter as it allows the most time for immune cell counts to decrease before dosing again, thus minimizing the risk of autoimmune myocarditis.

The variety of dosing schedules that prevent autoimmune myocarditis allows the clinician who is designing the treatment for a specific patient to take into account many other factors, such as the effect of the magnitude and timing of ICI doses on the tumour, and the risk of other detrimental side–effects that may come with large doses of ICIs. However, an effective clinical protocol based on this modelling and optimization approach would require the clinician to carefully monitor the patients levels of pathogenic CD4+ T cells as the timing, and magnitude of the immune perturbation, and how quickly the perturbation is detected form an important part of the modelling approach.

We chose the values of V_N and V_I such that the model predicts the development of autoimmune myocarditis when standard–of–care treatment for monotherapies is applied, and thus there is a problem to solve with optimal control theory. The value of V_N that we found in Section 4.2.1, 0.25 grams, is close to the real dose of nivolumab in monotherapy that is administered every two weeks, 0.24 grams, as well as the dose of nivolumab used in some combination therapy schedules [43, 44]. The values of V_I are also the correct order of magnitude as doses of 10 mg/kg for monotherapy and 1 mg/kg for combination therapy are recommended in various studies [44, 162].

In future work, it will be important to further consider if the choices of V_N, V_I, and γ_i, $i = 1, \ldots, 4$, are biologically realistic. There are three possible scenarios when a dose of ICIs is applied: (1) no autoimmune myocarditis develops with the standard–of–care treatment schedule; (2) autoimmune myocarditis develops with the standard–of–care treatment schedule, but using optimal control theory we can identify treatment schedules for which this dose does not cause the disease; or (3) autoimmune myocarditis develops with the standard–of–care treatment schedule, and this cannot

be prevented. The exploration of the effect of the parameters V_N and V_I on the model in the Appendix A shows that these scenarios are all possible in our model, and that the relative values of the γ_i, $i = 1, \ldots, 4$, can have a significant impact on whether or not autoimmune myocarditis develops. The question then remains: for a specific patient, where do the boundaries between these different outcome regions lie, and what are the true values of the γ_i? The answers to these questions are essential when we consider how to avoid inducing autoimmune myocarditis in a patient when they are treated with ICIs. A more in–depth exploration of patient–specific values for model parameters and biologically realistic ICI doses in the context of this model will enable us to shed light on these issues.

In addition to the patient–specific parameters that must be accurately determined, other questions arise such as the optimal length of the treatment window and the optimal number of doses to be administered. The framework presented in this chapter, and extensions thereof, can be used in the future to further explore these questions [24]. In this chapter, the payoff function only included the number of dead/damaged cardiomyocytes at the final time point of the simulation because the number and size of the ICI doses was not varied. If in the future the number and size of the doses is varied as well, the payoff function will need to be adjusted to reflect the benefits of administering more or less ICIs to the patient.

The control algorithm identifies a wide variety of dosing schedules which do not cause autoimmune myocarditis to develop. Considering other possible side–effects and the effects of ICIs on the tumour could further narrow down what an optimal dosing schedule for ICIs may look like. Off–target effects of ICIs occur in other organs besides the heart, such as the lungs, liver and bones, and these can be just as dangerous to patients as autoimmune myocarditis [117]. It is likely inadvisable to administer four doses of ICIs at once to a patient due to toxicity, but also because letting the tumour grow uninhibited for nearly 60 days is likely to be very detrimental

to the patient. A logical next step would thus be to extend the model of autoimmune myocarditis and place it within a larger framework that considers the systemic effects of ICIs throughout the body. This framework could consist of different connected compartments, each modelling the effect of ICIs on a specific organ or tissue. In the next chapter, we start by adding a simple compartment representing the tumour and the effects of ICIs on its growth.

Chapter 5

A heart–tumour model of autoimmune myocarditis as a side–effect of immune checkpoint inhibitors

5.1 Introduction

In optimizing the treatment strategies for nivolumab and ipilimumab therapies in Chapter 4, we focused solely on preventing the development of autoimmune myocarditis. We did not consider how changing the dosing schedule of nivolumab and/or ipilimumab might impact tumour growth. To design treatment schedules that are clinically relevant, we must consider both the positive effects of the ICI treatment on the tumour as well as the negative effects of ICIs on the heart. In this chapter, we extend our model of autoimmune myocarditis by including a simple tumour compartment and use optimal control approaches to identify optimized dosing schedules that both prevent autoimmune myocarditis and delay tumour growth.

When it comes to mathematical modelling of tumour growth, the literature is vast, and many tumour growth laws have been proposed over the years [98, 166]. As different types of tumours follow different growth curves under different experimental conditions, no choice of growth law is universally applicable and thus a variety have been used in modelling cancer [81, 121, 134, 168]. Little is known about growth patterns of tumours after detection, and even less is known about how tumours develop before they are clinically detectable. The hypothesis that the immune system detects and eliminates pre–cancerous lesions and tiny tumours, proposed by Burnet and Thomas, was initially rejected due to a lack of experimental evidence, but has since become more popular [18, 19, 146]. Evidence in support of this hypothesis includes the fact that immunodeficiency, whether primary or acquired, has been associated with an increased risk of cancer, and that cancer cells carry mutations in genes which encode proteins that are specifically recognized by the adaptive immune system [32, 40].

The elimination of small tumours by the immune system, and the growth of those tumours which manage to escape it, suggests that tumour growth could be described by a mathematical model exhibiting similar steady state behaviour to our model of autoimmune myocarditis, as discussed in Chapter 3. In the case of tumour growth, the stable steady state at the origin represents a healthy, cancer–free state. A second stable steady state represents a tumour that has escaped the immune system and grown to carrying capacity. Lastly, the saddle point represents the threshold size of the tumour that the immune system can control and/or eliminate. Cancer treatment, for example with ICIs, inhibits tumour growth, so we would expect the presence of ICIs to move the saddle point away from the origin.

In this chapter, we develop a simple model of tumour growth that exhibits the steady state behaviour described above, and connect it to our model of autoimmune myocarditis. We apply the optimal control algorithm from Chapter 4 to the two–

compartment model to design treatment strategies that delay tumour growth while preventing autoimmune myocarditis from developing. In Section 5.2, we present the tumour–only model, set the parameter values and use optimal control theory to show that standard–of–care ICI treatment is optimal for treating the tumour. In Section 5.3, we couple the tumour model to the model of autoimmune myocarditis, find its steady states, and set the parameter values. In Section 5.4, we use optimal control theory to design treatment schedules for nivolumab monotherapy, ipilimumab monotherapy and combination therapy that delay tumour growth while preventing autoimmune myocarditis. In Section 5.5, we discuss the results and point out avenues of research that future work may explore.

5.2 Tumour model

The primary aim of this chapter is to construct a framework for the optimization of ICI treatment schedules that considers both the positive and negative effects of ICIs. To this end, we develop in this section a simple model of tumour growth which displays the steady state behaviour we hypothesized in Section 5.1. In the absence of data, it is difficult to build a model which quantitatively predicts tumour growth in response to ICI treatment accurately. The simple tumour model used here serves as a place–holder within the theoretical framework and it can be easily exchanged for other, more biologically realistic models, as required.

5.2.1 Model equations

We denote tumour volume in cubic centimeters by T_u. Due to a lack of data from which biologically realistic tumour volumes can be estimated, it is currently not possible to assess what the size of the tumour should be at the saddle point and carrying capacity. We can, however, ensure that the model exhibits the qualitative steady

state behaviour and response to the presence of nivolumab (D_N) and ipilimumab (D_I) described in Section 5.1. We propose a simple model of tumour dynamics of the form

$$\frac{\mathrm{d}T_u}{\mathrm{d}t} = rT_u\left(\frac{T_u}{A_0(1+\gamma_5 D_N)(1+\gamma_6 D_I)} - 1\right)\left(1 - \frac{T_u}{K}\right), \quad T_u(0) = T_{u,0}, \tag{5.1}$$

where the parameters r, A_0, γ_5, γ_6, and K are positive constants, with r denoting the growth rate of the tumour, A_0 denoting the threshold volume for elimination of the tumour by the immune system in an ICI–free environment, K denoting the carrying capacity of the tumour and γ_5 and γ_6 capturing the pharmacodynamics of nivolumab and ipilimumab, respectively, where the presence of ICIs increases the tumour volume at the saddle point. Equation (5.1) admits three steady states, denoted $T_{u,*}$, which capture the behaviour described in Section 5.1: $T_{u,*} = 0$ cm^3, which is stable; $T_{u,*} = A_0(1+\gamma_5 D_N)(1+\gamma_6 D_I)$ cm^3, which is a saddle point; and $T_{u,*} = K$ cm^3, which is stable. Furthermore, as nivolumab and ipilimumab stimulate the immune response against the tumour, this should lead to larger tumours being eliminated by the immune system, or at least delayed in their growth, which is accomplished by the presence of these drugs affecting the position of the saddle point. This is illustrated in Fig. 5.1. In this work, we assume that at the start of the simulation the tumour has already escaped the immune system and thus $T_u(0) > A_0$. As ICIs are administered, two things can happen at a given time point: (1) the saddle point $A_0(1+\gamma_5 D_N)(1+\gamma_6 D_I)$ increases but does not become larger than $T_u(t)$ so $\mathrm{d}T_u/\mathrm{d}t$ remains positive and the tumour will keep growing, though at a slower rate; or (2) the saddle point $A_0(1+\gamma_5 D_N)(1+\gamma_6 D_I)$ grows larger than $T_u(t)$, so $\mathrm{d}T_u/\mathrm{d}t$ becomes negative and the tumour shrinks, at least temporarily. The ICI treatments that we consider in this work delay tumour growth but rarely if ever lead to tumour elimination [3, 47, 79]. For the remainder of this work, we will thus ensure that we are in a

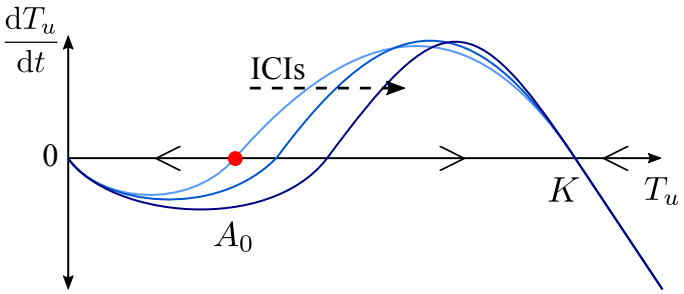


Figure 5.1: The RHS of Eq. (5.1) plotted as a function of the size of the tumour. For tumour volumes below the saddle point, A_0 (red dot), the RHS of Eq. (5.1) is negative and tumour volume will decrease. For tumour volumes above the saddle point, the RHS of Eq. (5.1) is positive and the tumour will grow to carrying capacity, K. Increasing the level of ICIs in the system (dashed arrow) increases the tumour volume at the saddle point.

regime where $T_u(0)$ is sufficiently large such that dT_u/dt remains positive.

5.2.2 Parameter values

We need to specify values for the parameters in the tumour growth model: r, A_0, K, γ_5, and γ_6. In the absence of data, we set the values for these parameters based on the required qualitative behaviour of the model. The parameters r, A_0, and K govern the growth of the tumour, and γ_5 and γ_6 translate the mass of nivolumab and ipilimumab, respectively, to their effect on anti–tumour immunity. In the absence of ICIs a tumour with a size greater than A_0 should grow to carrying capacity whereas the presence of nivolumab or ipilimumab should delay the growth of the tumour by a period on the order of months to fit with the average progression–free survival of patients treated with these ICIs [3, 47, 79]. For $r = 0.001$ cm^3 day^{-1}, $A_0 = 0.001$ cm^3, $K = 10$ cm^3, $\gamma_5 = 100$ grams^{-1} and $\gamma_6 = 25$ grams^{-1}, this behaviour is realized, as shown in Fig. 5.2 where tumour growth is simulated over six months from an initial tumour volume of 0.1 cm^3. We pick this initial tumour volume to use throughout this chapter as it is well above the ICI–free saddle point of 0.001 cm^3 and the tumour will thus grow initially even in scenarios where ICIs are applied. As the tumour volume at the saddle point increases as a consequence of ICIs being administered, tumour growth

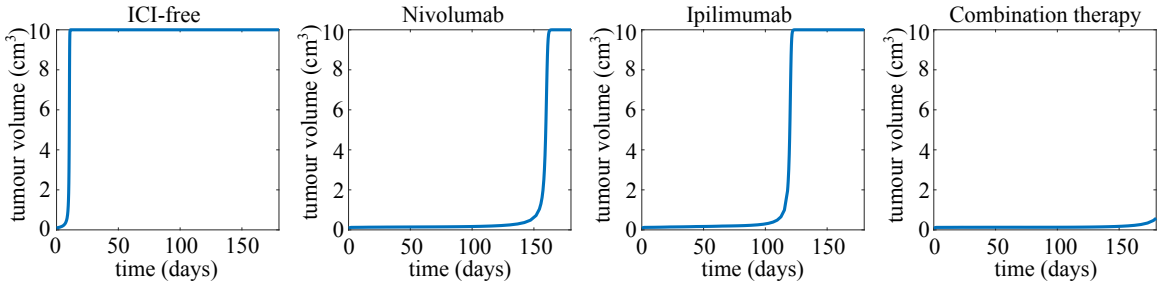


Figure 5.2: Time traces of tumour growth under ICI–free conditions, or with standard–of–care nivolumab, ipilimumab or combination therapy. The initial tumour volume is 0.1 cm^3, and initial ICI levels are zero. Standard–of–care nivolumab therapy is applied on days 2, 16, 30, 44, and 58 with $V_N = 0.25$ grams, as in Chapter 4. Standard–of–care ipilimumab therapy is applied on days 2, 23, 44, and 65 with $V_I = 1.0$ gram, as in Chapter 4. Standard–of–care combination therapy includes doses of nivolumab and ipilimumab on days 2, 23, 44, and 65 and a dose of nivolumab only on day 86 with $V_N = 0.25$ grams and $V_I = 0.1$ grams, as in Chapter 4. The parameter values listed in Table 5.1 are used unless otherwise noted.

will be delayed on the order of months, as shown in Fig. 5.2. This is consistent with observed patient responses to ICI therapy, so the simple tumour model as presented here shows the desired qualitative behaviour for this initial condition [3, 47, 79].

5.2.3 Optimal control of tumour growth

In this section, we use the optimal control algorithm from Castiglione and Piccoli [23] to find a schedule of ICI doses that delays tumour growth. We construct a payoff function that penalizes tumour growth at any time point in the simulation. The payoff function to be minimized is

$$J = \int_0^{t_f} T_u(t)\mathrm{d}t, \tag{5.2}$$

where t_f is the final time point of the simulation. Nivolumab and ipilimumab inhibit tumour growth, *i.e.* decrease $T_u(t)$, which is desirable and rewarded by the payoff function. As the payoff function does not include the levels of ICIs explicitly, this payoff function can be used to explore schedules for all three types of therapy.

In each case, we initialize the model with a growing tumour which is larger than the saddle point but small compared to the carrying capacity, *i.e.* $T_{u,0} = 0.1$ cm^3. To use the optimal control algorithm explained in Chapter 4 Section 4.3, we must set a convergence threshold, a step size for the gradient descent algorithm and the time point $t^-_{i_\alpha,\alpha}$ at which the variation equations are initialized. The convergence threshold is set to 0.1, and the gradient descent step size, h, is equal to 0.01. The time points $t_{i_\alpha,\alpha}$, where $\alpha \in \{N, I\}$, $i_N = [1, \ldots, n_N]$ and $i_I = [1, \ldots, n_I]$, are the time points at which the drug doses are administered. The time point at which the variation equations are initialized, $t^-_{i_\alpha,\alpha}$, is defined as $t^-_{i_\alpha,\alpha} = t_{i_\alpha,\alpha} - 10^{-10}$. All model parameters are set to their default values established in Section 5.2.2.

As in Chapter 4, we set the dose of nivolumab equal to $V_N = 0.25$ grams for monotherapy and combination therapy, and the dose of ipilimumab equal to $V_I = 1.0$ grams for monotherapy and $V_I = 0.1$ grams for combination therapy. The first dose of ICI therapy is administered on day 2 and is not part of the optimization. This is because autoimmune myocarditis as a side–effect of ICIs cannot develop prior to the first dose of ICIs being administered. This will be important in Section 5.3 when we also optimize for prevention of autoimmune myocarditis, so for consistency we also do not optimize the timing of the dose on day 2 here. As in Chapter 4, we let the treatment window start on day 10, so $t_{\min} = 10$. The upper bound on the treatment window, $t_{\max}$, depends on the type of ICI therapy being administered, and simulations are stopped 100 days after the end of the treatment window. Fig. 5.3 shows an overview of the timeline of the optimization problem.

We initialize the optimization of nivolumab monotherapy, ipilimumab monotherapy and combination therapy with their respective standard–of–care schedules. For all three types of therapy, we find that the optimized treatment schedule is nearly identical to the standard–of–care treatment schedule. For nivolumab monotherapy, the standard–of–care dosing schedule schedule is $S_0 = [16,\ 30,\ 44,\ 58]$ and the optimized

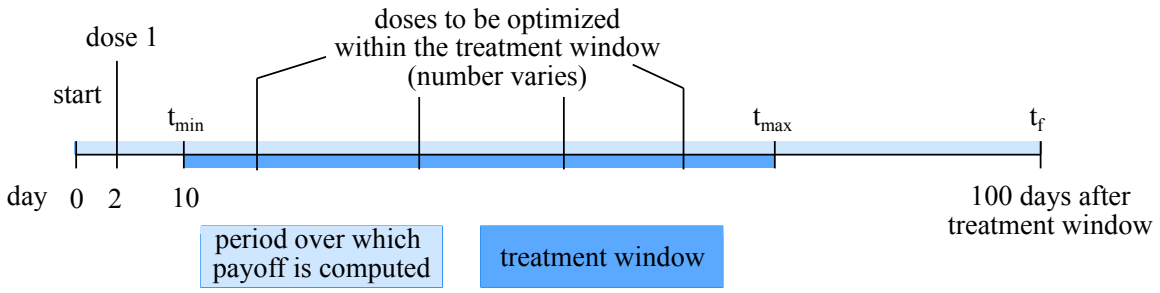


Figure 5.3: Timeline of the optimization problem for ICI treatment schedule design.

dosing schedule is $S = [16.0,\ 30.0,\ 44.0,\ 58.0]$. For ipilimumab monotherapy, the standard–of–care dosing schedule schedule is $S_0 = [23,\ 44,\ 65]$ and the optimized dosing schedule is $S = [23.0,\ 44.0,\ 65.0]$. For combination therapy, the standard–of–care dosing schedule schedule is $S_{N,0} = [23,\ 44,\ 65,\ 86]$, and $S_{I,0} = [23,\ 44,\ 65]$ and the optimized dosing schedule is $S_N = [23.0,\ 44.0,\ 65.1,\ 86.1]$, and $S_I = [23.0,\ 44.0,\ 65.0]$.

Based on delay of tumour growth alone, the standard–of–care treatment schedule for ICI therapy thus appears to be optimized. In Section 5.4, we explore whether these treatment schedules are sufficient to simultaneously prevent autoimmune myocarditis and delay tumour growth in a model where the tumour growth model is coupled to the model of autoimmune myocarditis.

5.3 Two–compartment heart–tumour model

To construct a model that captures the positive effects of ICIs on tumour growth and the negative effects of ICIs on the heart, we now couple a heart compartment, which consists of the model of autoimmune myocarditis developed in Chapter 4, to a tumour compartment, which consists of the model of tumour growth developed in Section 5.2.

5.3.1 Coupling the tumour and heart compartments

Anti–tumour T cells that cross–react with cardiac antigens are a known potential trigger for autoimmune myocarditis as a side–effect of ICIs, so we connect the heart compartment to the tumour compartment by adding a small, unidirectional flow of pathogenic T cells from the tumour to the heart [117]. We assume this flow to be unidirectional as the activation of pathogenic T cells initially happens as part of the anti–tumour immune response, and these activated pathogenic T cells can then circulate through the body and end up in the heart. The number of active T cells coming from the heart to the tumour microenvironment is assumed to be negligible compared to the number of active T cells already in the tumour microenvironment.

We recall from Chapter 4 that the model of autoimmune myocarditis which explicitly captures the effects of nivolumab and ipilimumab describes the evolution of six state variables: dead/damaged cardiomyocytes (C); innate immune cells (I); pathogenic CD4+ T cells (P); regulatory T cells (R); nivolumab (D_N); and ipilimumab (D_I). We couple the tumour growth model (Eq. (5.1)) to the model of autoimmune myocarditis by a unidirectional flow of pathogenic CD4+ T cells. The unidirectional flow of pathogenic CD4+ T cells is represented in the model by the addition of a term to Eq. (5.5) which is highlighted here in red:

$$\frac{\mathrm{d}C}{\mathrm{d}t} = a_1\frac{P^2}{R_1+R} + a_2\frac{I}{R_2+R} - d_1C, \tag{5.3}$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = a_4\frac{P}{R_1+R} + a_5C - d_3I, \tag{5.4}$$

$$\begin{aligned}\frac{\mathrm{d}P}{\mathrm{d}t} &= a_6\phi(1+\gamma_1D_N)(1+\gamma_2D_I)\frac{C}{R+R_3} - \frac{d_4}{1+\gamma_3D_N}P \\ &\quad +\delta\left(1+\gamma_7\frac{D_N}{D_N+1}\right)\left(1+\gamma_8\frac{D_I}{D_I+1}\right)\left(1+\frac{T_u}{T_u+1}\right),\end{aligned} \tag{5.5}$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = a_6(1+D_N)(1+D_I)(1-\phi)\frac{C}{R+R_3} - d_5(1+\gamma_4D_I)R, \tag{5.6}$$

$$\frac{\mathrm{d}D_N}{\mathrm{d}t} = -d_6D_N, \text{ with } D_N(t=t_{i_N,N}) = D_N(t=t^-_{i_N,N}) + V_N, \tag{5.7}$$

$$\frac{\mathrm{d}D_I}{\mathrm{d}t} = -d_7D_I, \text{ with } D_I(t=t_{i_I,I}) = D_I(t=t^-_{i_I,I}) + V_I, \tag{5.8}$$

$$\frac{\mathrm{d}T_u}{\mathrm{d}t} = rT_u\left(\frac{T_u}{A_0(1+\gamma_5D_N)(1+\gamma_6D_I)} - 1\right)\left(1-\frac{T_u}{K}\right), \tag{5.9}$$

where δ, γ_7 and γ_8 are positive, constant parameters. For a schedule consisting of $n = n_N + n_I$ total ICI doses, the time points $t_{i_\alpha,\alpha}$, where $\alpha \in \{N, I\}$, $i_N = [1, \ldots, n_N]$ and $i_I = [1, \ldots, n_I]$, are the time points at which a dose of ICIs is administered (nivolumab if $\alpha = N$ and ipilimumab if $\alpha = I$). The parameter δ represents the baseline flow of cardiac–antigen specific T cells that circulate through the body in a healthy state [52, 80]. The presence of a tumour is assumed to increase this flow as the tumour provides a source of activated cross–reacting T cells. The number of active T cells in the tumour compartment is assumed to correlate with tumour volume. Because nivolumab and ipilimumab both stimulate the anti–tumour immune response, the presence of these ICIs is assumed to further increase the level of T cells in the tumour compartment and thus in circulation throughout the body. The parameters γ_7 and γ_8 translate the mass of nivolumab and ipilimumab in the system to the effect on stimulating T cell activation in the tumour compartment. The system of ODEs is closed by defining the initial conditions $[C(0),\ I(0),\ P(0),\ R(0)\ D_N(0),\ D_I(0),\ T_u(0)] = [C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0},\ T_{u,0}]$.

5.3.2 Steady states

We derive a polynomial that describes the non–zero, ICI–free steady states of the model by setting the RHS of the Eq. (5.3)–(5.9) equal to zero and rearranging. From Eq. (5.7) and (5.8), we obtain that the steady state values for nivolumab, denoted $D_{N,*}$, and ipilimumab, denoted $D_{I,*}$, are $D_{N,*} = 0$ grams and $D_{I,*} = 0$ grams, so any steady state will be ICI–free. From the steady states of the tumour growth model noted in Section 5.2.1 we obtain that the ICI–free steady state values for tumour volume are $T_{u,*} = 0$ cm^3, $T_{u,*} = A_0$ cm^3, and $T_{u,*} = K$ cm^3. The following polynomial describes the non–zero, ICI–free steady state cell counts of pathogenic T cells, P_*, in terms of tumour volume, $T_{u,*}$:

$$\begin{aligned} 0 = {} & d_1 d_4 p_1 P_*^4 \\ & + (d_1 d_4 p_2 - K_F d_1 p_1 - a_1 a_6 d_1 d_3 \phi K_{R1})\, P_*^3 \\ & + (d_1 d_4 p_3 - K_F d_1 p_2 - a_1 a_6 d_1 d_3 \phi (R_2 - K_{R2}))\, P_*^2 \\ & + (d_1 d_4 p_4 - K_F d_1 p_3 - a_2 a_4 a_6 d_1 \phi)\, P_* \\ & - K_F d_1 p_4, \end{aligned}$$

where

$$\begin{aligned} p_1 &= d_1 d_3 K_{R1}^3, \\ p_2 &= K_{R2}^2 (d_1 d_3 R_2 - d_1 d_3 K_{R2} - a_2 a_5 + d_1 d_3 (R_3 - K_{R2}) + d_1 d_3 (R_1 - K_{R2})), \\ p_3 &= K_{R_2} ((R_1 + R_3 - 2K_{R2})(d_1 d_3 R_2 - d_1 d_3 K_{R2} - a_2 a_5) + d_1 d_3 (R_1 - K_{R2})(R_3 - K_{R2})), \\ p_4 &= (R_1 - K_{R2})(R_3 - K_{R2})(d_1 d_3 R_2 - d_1 d_3 K_{R2} - a_2 a_5), \end{aligned}$$

and

$$
\begin{aligned}
K_{R1} &= \frac{(1-\phi)d_4}{\phi d_5}, \\
K_{R2} &= \frac{1-\phi}{\phi d_5} K_F, \\
K_F &= \delta \left(1 + \frac{T_{u,*}}{1+T_{u,*}}\right).
\end{aligned}
$$

The ICI–free steady state values of the remaining variables are then given in terms of P_* by

$$
\begin{aligned}
R_* &= K_{R1} P_* - K_{R2}, \\
I_* &= \frac{a_4 d_1 P_* + a_1 a_5 P_*^2}{(R_* + R_1)(d_1 d_3 - a_2 a_5/(R_* + R_2))}, \\
C_* &= \frac{a_1}{d_1} \frac{P_*^2}{R_* + R_1} + \frac{a_2 I_*}{d_1 (R_* + R_2)}.
\end{aligned}
$$

As in Chapter 4, we say that autoimmune myocarditis has developed as a consequence of ICI therapy if, at the final time point of a simulation, the cell counts in the system are higher than those at the ICI–free saddle point, which can be computed from the equations above.

5.3.3 Parameter values

We need to specify parameter values for the parameters which govern the flow of pathogenic T cells: δ, γ_7 and γ_8. In the absence of data, we set the values for these parameters based on the required qualitative behaviour of the model. Here we must ensure that δ is sufficiently small that autoimmune myocarditis does not develop in an ICI–free environment, regardless of tumour volume, as there is no evidence that the presence of a tumour alone increases the risk of autoimmune myocarditis developing. Furthermore, we require the values of γ_7 and γ_8 to be such that autoimmune

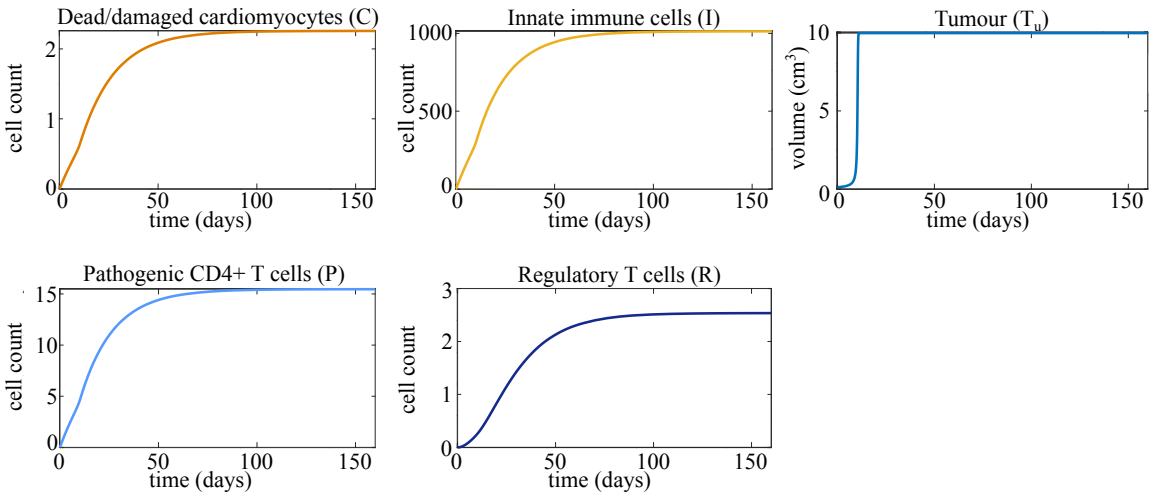


Figure 5.4: Time traces of state variables under ICI–free conditions. The initial tumour volume is 0.1 cm^3. The parameter values listed in Table 5.1 are used unless otherwise noted.

myocarditis develops for all three types of therapy, but at the earliest after the second dose. The reasoning for this is as follows. As in Chapter 4, we assume here that an increase in inflammatory cells due to ICI therapy can be detected a number of days after the first dose of ICIs has been administered and that adjustment to the dosing schedule can be made from that point forward. If the increase in inflammation after the first dose is so large that autoimmune myocarditis immediately develops, it is not possible to avoid this side–effect by optimizing subsequent doses because once autoimmune myocarditis has developed the system cannot return to the healthy steady state even in an ICI–free environment. To be in a position where we can use optimal control theory to adjust dosing schedules to prevent autoimmune myocarditis we must thus be in a scenario where myocarditis develops as a consequence of the second or later doses.

The simulations in Fig. 5.4 show that autoimmune myocarditis does not develop when the tumour grows in ICI–free conditions for $\delta = 0.5$ cells day^{-1}, $\gamma_7 = 5$ $grams^{-1}$ and $\gamma_8 = 5$ $grams^{-1}$. Fig. 5.5 shows that when standard–of–care ICI treatment is applied with the same values of V_N and V_I as in Chapter 4, tumour growth is

delayed, but autoimmune myocarditis develops later on in treatment for all three types of therapy. Thus, with the parameter values established here for δ, γ_7 and γ_8 and the values of V_N and V_I established in Chapter 4, we are in a parameter regime where we can use optimal control theory to investigate the problem of balancing tumour treatment with myocarditis prevention. A list of all parameter values is provided in Table 5.1. For these default parameter values, the saddle point in the heart compartment for a tumour of size $T_u = K$ cm^3 in an ICI–free environment is $[C,\ I,\ P,\ R] = [125,\ 1.90 \times 10^4,\ 317,\ 132]$ cells. Because the initial size of the tumour is well above the separatrix, the tumour will grow to carrying capacity even when ICI treatment is applied, as the doses of ICIs used here are too small to cause the saddle point to become bigger than the growing tumour. The cell counts at the ICI–free saddle point $[C,\ I,\ P,\ R] = [125,\ 1.90 \times 10^4,\ 317,\ 132]$ cells thus form a threshold above which the system is guaranteed to approach the diseased steady state. We conclude that autoimmune myocarditis will develop as a consequence of the applied dosing schedule if all cell counts in the system lie above the ICI–free saddle point at the end of a simulation. This is an overly strict criterion for development of autoimmune myocarditis as autoimmune myocarditis will develop as long as the cell counts in the system lie above the separatrix. However, for ease of implementation this definition is sufficient for our purposes.

We have defined a model consisting of our model of autoimmune myocarditis in a compartment representing the heart, connected to a second compartment representing a tumour microenvironment by a flow of pathogenic T cells. In the next section, we use this model to revisit the question of optimized ICI treatment design that we explored in Chapter 4.

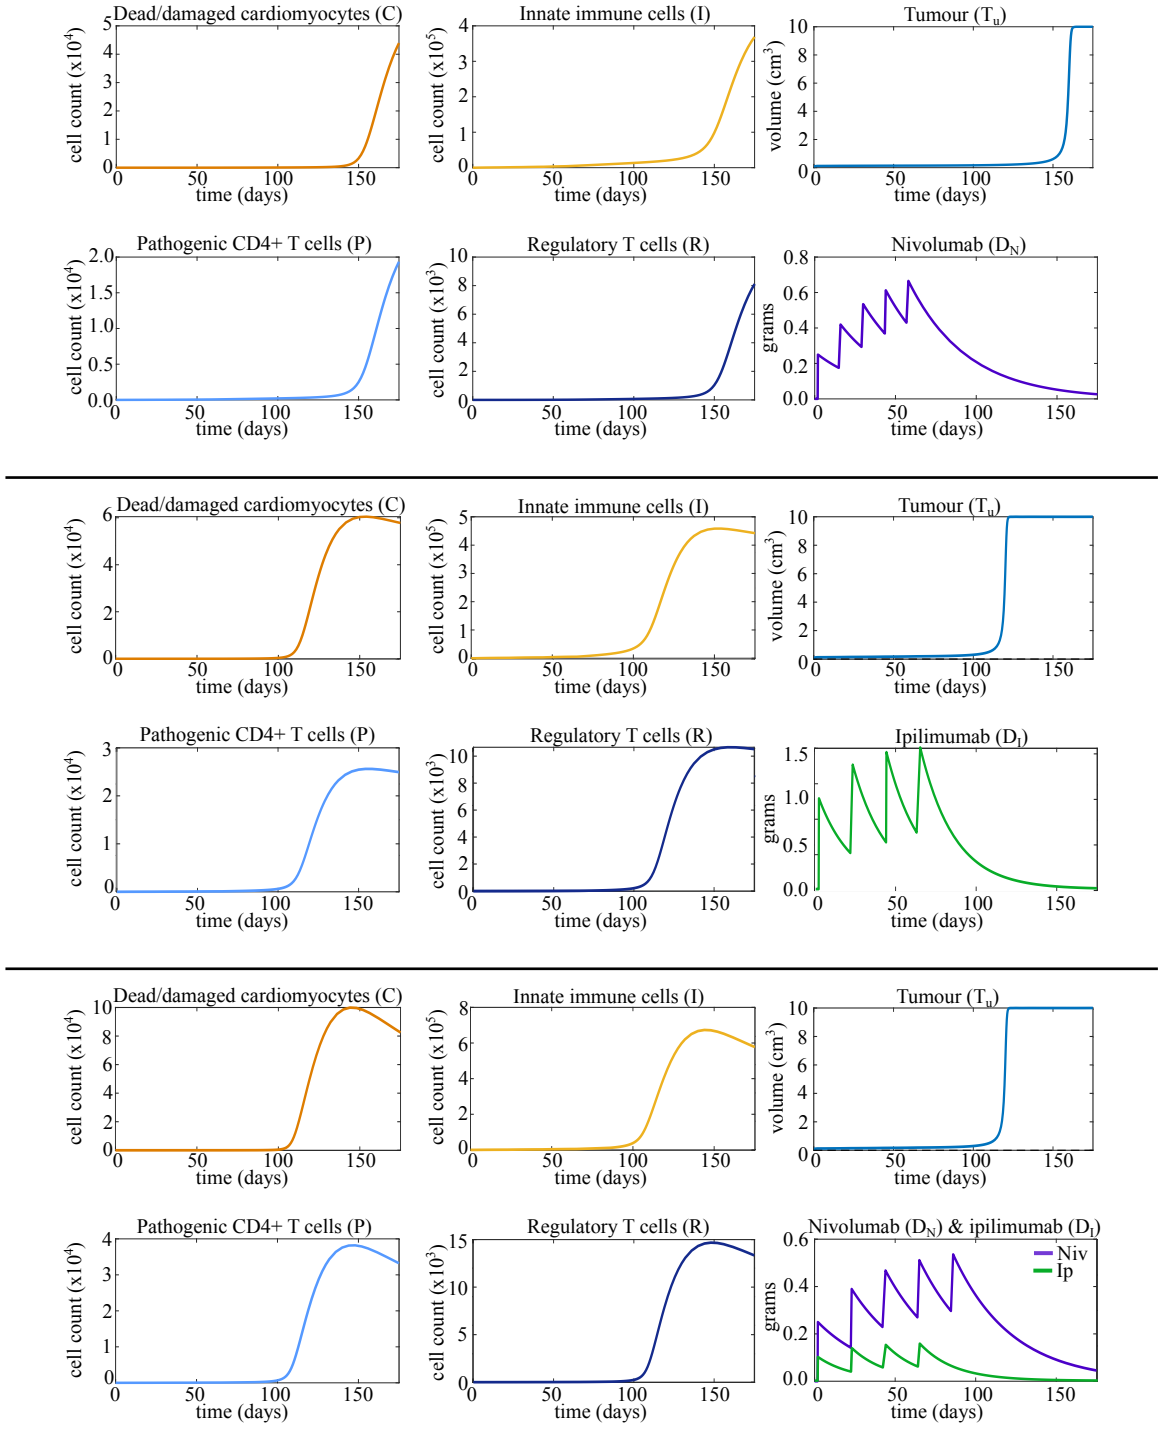


Figure 5.5: Time traces with standard–of–care nivolumab (top), ipilimumab (middle) or combination therapy (bottom). The initial tumour volume is 0.1 cm^3, and all other variables are initialized at zero. Standard–of–care nivolumab therapy is applied on days 2, 16, 30, 44, and 58 with $V_N = 0.25$ grams. Standard–of–care ipilimumab therapy is applied on days 2, 23, 44, and 65 with $V_I = 1.0$ gram. Standard–of–care combination therapy includes doses of nivolumab and ipilimumab on days 2, 23, 44, and 65 and a dose of nivolumab only on day 86 with $V_N = 0.25$ grams and $V_I = 0.1$ grams. The parameter values listed in Table 5.1 are used unless otherwise noted.

Param.	Definition	Value	Units
a_1	Rate at which pathogenic CD4+ T cells damage/kill cardiomyocytes	80	day^{-1}
a_2	Rate at which innate immune cells damage/kill cardiomyocytes	6000	cells day^{-1}
a_4	Rate at which pathogenic CD4+ T cells activate innate immune cells	500,000	cells day^{-1}
a_5	Rate at which damaged/dead cardiomyocytes activate innate immune cells	35	day^{-1}
a_6	Rate at which damaged/dead cardiomyocytes activate pathogenic CD4+ T cells	750	cells day^{-1}
d_1	Clearance rate of damaged/dead cardiomyocytes	75	day^{-1}
d_3	Innate immune cell death rate	7	day^{-1}
d_4	Pathogenic CD4+ T cell death rate	0.1	day^{-1}
d_5	Regulatory T cell death rate	0.1	day^{-1}
d_6	Degradation rate of nivolumab	$\log(2)/25$	day^{-1}
d_7	Degradation rate of ipilimumab	$\log(2)/14.7$	day^{-1}
R_1	Number of regulatory immune cells for 50% inhibition of pathogenic CD4+ T cell activity	1100	cells
R_2	Number of regulatory immune cells for 50% inhibition of innate immune cells	40000	cells
R_3	Number of regulatory immune cells for 50% inhibition of pathogenic CD4+ T cell activation	2000	cells
ϕ	Fraction activated CD4+ T cells that take on a pathogenic phenotype	0.7	unitless
δ	Rate of transport of T cells from tumour to heart	0.5	cells day^{-1}
r	Maximal growth rate of the tumour	0.001	cm^3 day^{-1}
A_0	Maximum tumour volume for immune system control	0.001	cm^3
K	Carrying capacity of tumour	10	cm^3
γ_1	Relative effect of nivolumab on T cell activation	1	grams^{-1}
γ_2	Relative effect of ipilimumab on T cell activation	1	grams^{-1}
γ_3	Relative effect of nivolumab on pathogenic T cell death	1	grams^{-1}
γ_4	Relative effect of ipilimumab on regulatory T cell death	1	grams^{-1}
γ_5	Relative effect of nivolumab on T cell activation and transport from tumour site	100	grams^{-1}
γ_6	Relative effect of ipilimumab on T cell activation and transport from tumour site	25	grams^{-1}
γ_7	Relative effect of nivolumab on T cell activation against the tumour	5	grams^{-1}
γ_8	Relative effect of ipilimumab on T cell activation against the tumour	5	grams^{-1}

Table 5.1: Definitions, values and units of model parameters.

5.4 Optimized dosing schedules for ICI therapy

In this section, we use the optimal control algorithm from Castiglione and Piccoli [23] to find a schedule of ICI doses that prevents autoimmune myocarditis from developing at the same time as delaying tumour growth. The payoff landscape explored by the gradient descent algorithm is large and high–dimensional. It is thus unlikely that initializing with the standard–of–care dosing schedule will lead to the algorithm converging to the global minimum which represents the optimal dosing schedule. However, our primary aim is to identify a schedule which prevents autoimmune myocarditis and delays tumour growth to some extent compared to no treatment, which may be achieved by a schedule represented by a local minimum. An optimized, though not necessarily globally optimal, schedule would thus suffice to achieve our objective and it is this we seek in this chapter. Future research can focus on selecting a dosing schedule that best inhibits tumour growth from among the optimized schedules that prevent autoimmune myocarditis obtained here.

The payoff function will need to account for both the positive and negative effects of ICI treatment. The presence of nivolumab and ipilimumab will delay tumour growth, which is a desired effect and should thus be rewarded by the payoff function. This is achieved by including the tumour volume, $T_u(t)$, integrated over time in the payoff function. Nivolumab and ipilimumab also increase the risk of autoimmune myocarditis developing, and this side–effect is penalized by including the number of dead/damaged cardiomyocytes, $C(t)$, integrated over time in the payoff function. The payoff function to be minimized is

$$J = \int_0^{t_f} \left[w_1 C(t) + w_2 T_u(t) \right] \mathrm{d}t. \tag{5.10}$$

This payoff function ensures that a penalty is incurred if autoimmune myocarditis develops or the tumour grows quickly. The presence of nivolumab and ipilimumab

has the opposite effect on the first term in the integral compared to the second term as the presence of ICIs delays tumour growth, *i.e.* decreases $T_u(t)$, and stimulates damage to the myocardium, *i.e.* increases $C(t)$. The first is desirable and rewarded by the payoff function whereas the second is not. An optimized dosing schedule will strike a careful balance between delaying tumour growth and preventing autoimmune myocarditis to minimize this payoff function. As the payoff function does not include the levels of ICIs explicitly, this payoff function applies to all three types of therapy.

In each case, we initialize the model with an ICI–free, inflammation–free heart compartment and a growing tumour which is small compared to the carrying capacity, *i.e.* $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0},\ T_{u,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0,\ 0.1]$. We set the weights $w_1 = 1$ cells^{-1} and $w_2 = 1$ cm^{-3} for simplicity. The convergence threshold is set to 0.1, and the gradient descent step size, h, is equal to 0.01. In Appendix B, we show that the optimization results for nivolumab and ipilimumab monotherapy are qualitatively robust to changes in w_2 and h, but that care should be taken in selecting h if a different value of w_2 is selected for combination therapy. The time point $t^-_{i_\alpha,\alpha}$ used in the initial conditions of the variation equations (see Chapter 4 Section 4.3) is $t^-_{i_\alpha,\alpha} = t_{i_\alpha,\alpha} - 10^{-10}$.

We set up the optimization problem as we did in Section 5.2.3, with the first dose of ICI therapy administered on day 2, which is not part of the optimization. We assume that some time is required after the first dose of ICIs is administered for the immune response in the heart to develop and be clinically detectable, after which action can be taken to adjust the subsequent doses to prevent autoimmune myocarditis from developing. To reflect this, we let the treatment window start on day 10, so $t_{\min} = 10$ days. The upper bound on the treatment window varies by type of ICI therapy, and simulations are stopped 100 days after the end of the treatment window.

Our aim in this chapter is to consider the effects of ICIs in the heart and tumour

simultaneously, and the way the optimization problem has been set–up above is one way to do this. By adding a tumour compartment to the model and redefining the payoff function to include tumour volume, the optimization problem has changed significantly from the set–up we had in Chapter 4 and a quantitative comparison between the optimized dosing schedules obtained in this chapter and the schedules found in Chapter 4 will thus not be possible. We can, however, explore whether including the tumour in the model and optimizing for both the positive and negative effects of ICIs simultaneously has a significant qualitative effect on what the optimized dosing schedules look like.

In Sections 5.4.1–5.4.3, we explore in turn the optimized treatment schedule for nivolumab monotherapy, ipilimumab monotherapy and combination therapy.

5.4.1 Nivolumab monotherapy

We initialize the dosing schedule at the standard–of–care schedule of nivolumab monotherapy, which is a dose every two weeks starting from day 2. Since we do not optimize the timing of the first dose, we have $S_0 = [16,\ 30,\ 44,\ 58]$. The final day of the treatment window is day 60 (*i.e.* $t_{\max} = 60$), and we keep $V_N = 0.25$ grams as in Chapter 4.

Fig. 5.6 shows the results of the optimization over the iterations of the gradient descent algorithm. It takes 410 iterations to find a dosing schedule that does not cause autoimmune myocarditis to develop, and a total 3114 iterations for the algorithm to converge with the required tolerance. The optimized dosing schedule is $S = [10.0,\ 22.7,\ 59.7,\ 60.0]$.

Compared to Fig. 4.6, the optimized treatment schedule is qualitatively different if the tumour dynamics are included in the model and the payoff function. Without a tumour compartment present, the only objective is to prevent myocarditis from developing, which is best achieved by delaying the administration of nivolumab as

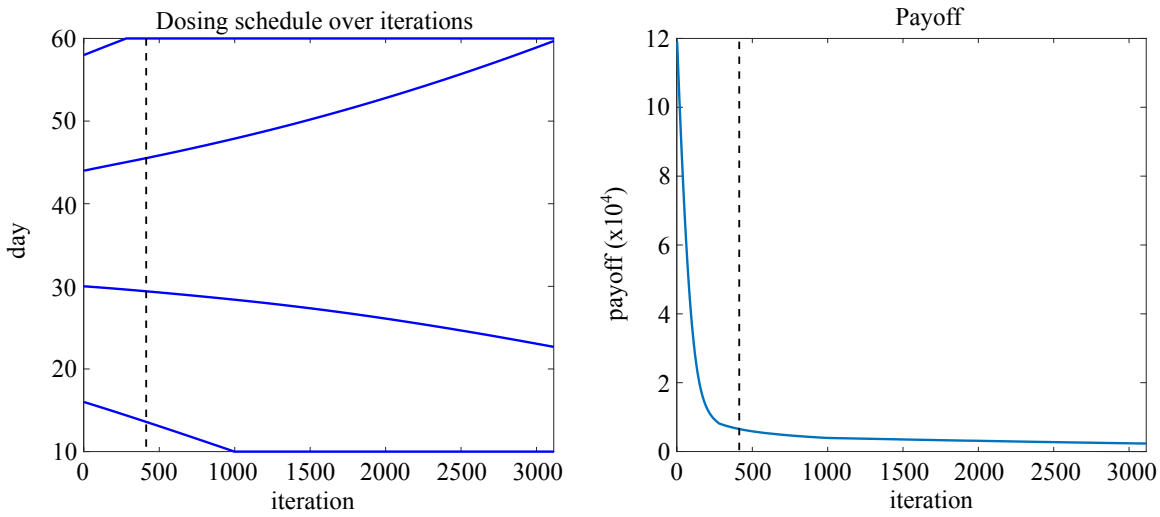


Figure 5.6: Left: the dosing schedule of nivolumab monotherapy evolving over the iterations of the optimization algorithm. Right: the payoff evolving over the iterations of the optimization algorithm. The dashed line indicates the first iteration at which autoimmune myocarditis does not develop.

long as possible, as shown in Fig. 4.6. However, such a large period of time without any nivolumab doses being administered will allow the tumour to grow uninhibited. The left–hand panel of Fig. 5.6 shows that to optimize for the prevention of autoimmune myocarditis as well as delay of tumour growth the dose timings are spread out to cover the whole treatment window. Regularly dosing keeps tumour growth under control, while the increased time between the second and third doses, not counting the dose on day 2, compared to the standard–of–care schedule allows immune cell counts to decrease, which prevents autoimmune myocarditis from developing.

Fig. 5.7 shows the time traces of the model variables for the optimized dosing schedule. When we compare tumour growth under this optimized dosing schedule to tumour growth under the standard–of–care schedule (see Fig. 5.2), there is no significant difference. However, Fig. 5.7 also shows that autoimmune myocarditis does not develop when the optimized treatment schedule is applied, as shown by the low numbers of inflammatory cells at the final time point. The optimized treatment schedule thus prevents autoimmune myocarditis from developing without diminishing the effect of treatment on the tumour.

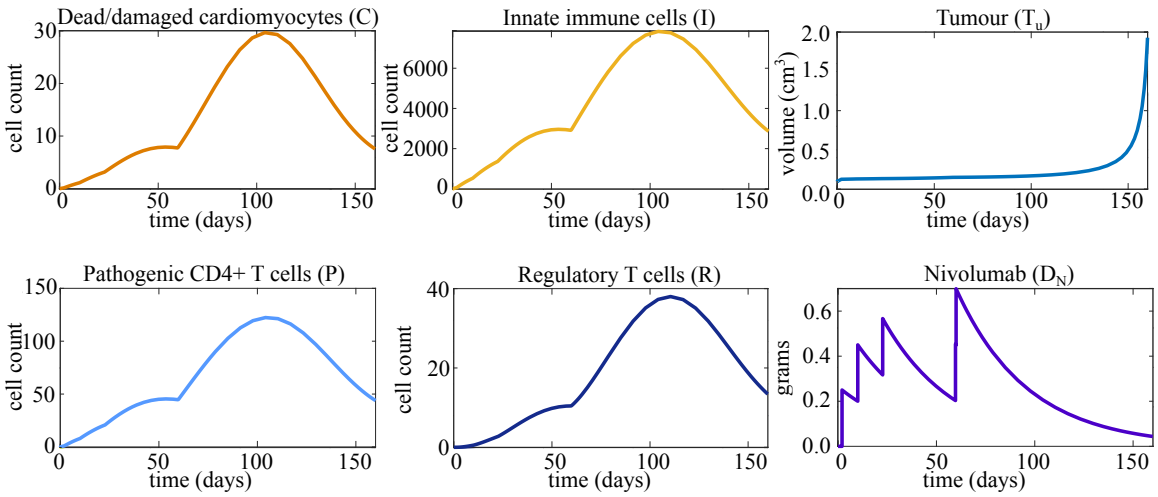


Figure 5.7: Time traces of model variables when the optimized dosing schedule for nivolumab monotherapy is applied. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0},\ T_{u,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0,\ 0.1]$, with the parameter values listed in Table 5.1 used.

In Appendix B.1, we explore how the results discussed here change as the dose size V_N, the length of the treatment window, the step parameter h and the relative values of the payoff function weights w_1 and w_2 are varied. We show that there is a very narrow range of values of V_N where autoimmune myocarditis does develop when a standard–of–care schedule is applied but can be prevented by an optimized dosing schedule. We further show that our results are qualitatively robust to changes in the length of the treatment window, the step parameter h and the payoff function weights.

5.4.2 Ipilimumab monotherapy

We set up the optimization problem for ipilimumab monotherapy similarly to that for nivolumab monotherapy. We initialize the schedule at the standard–of–care treatment schedule so $S_0 = [23,\ 44,\ 65]$ with a dose on day 2 which is not adjusted. The final day of the treatment window ($t_{\max}$) is day 70, and we keep $V_I = 1.0$ grams as in Chapter 4.

The algorithm converges after 2049 iterations and the optimized ipilimumab treat-

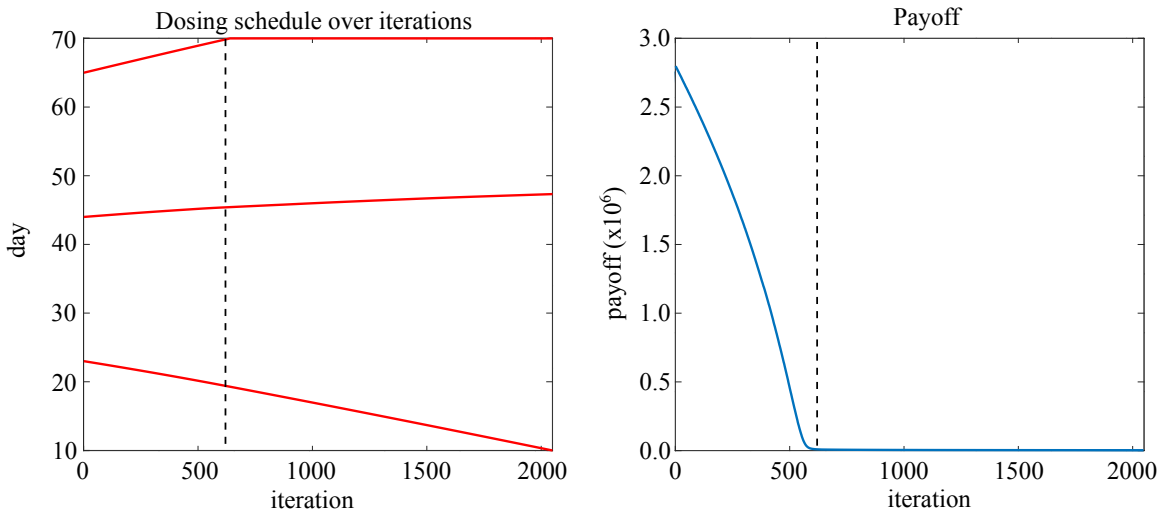


Figure 5.8: Left: the dosing schedule of ipilimumab monotherapy evolving over the iterations of the optimization algorithm. Right: the payoff evolving over the iterations of the optimization algorithm. The dashed line indicates the first iteration at which autoimmune myocarditis does not develop.

ment schedule is $S = [10.0,\ 47.4,\ 70.0]$. When we consider the results of the optimization in Fig. 5.8, we see a similar trend to what we saw for nivolumab monotherapy. For ipilimumab monotherapy in the heart–only model in Chapter 4 the main change in the dosing schedule was the delay of all doses compared to the standard–of–care treatment, as shown in Fig. 4.8. When we include the effect of ipilimumab on the tumour in the model, we see that the timings of the doses spread out to cover nearly the whole treatment window. This again allows for more time between doses during which the immune cell counts can decrease, while still keeping tumour growth under control (see Fig. 5.9). The variety of schedules that achieve the objective of both delaying tumour growth and reducing inflammation is smaller than when only myocarditis is considered, emphasizing again the tight constraints set by the need to balance the positive and negative effects of ICIs.

In Appendix B.2, we explore how the results discussed here change as the ipilimumab dose V_I, the length of the treatment window, the step parameter h and the relative values of the payoff function weights w_1 and w_2 are varied. We show that there exists a range of V_I for which autoimmune myocarditis develops under

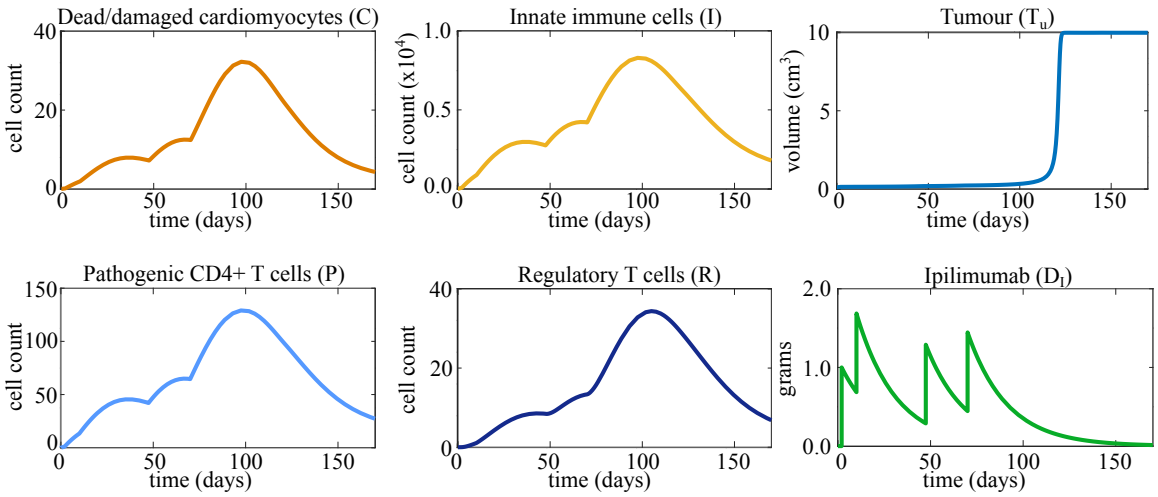


Figure 5.9: Time traces of model variables when the optimized dosing schedule for ipilimumab monotherapy is applied. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0},\ T_{u,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0,\ 0.1]$, with the parameter values listed in Table 5.1 used.

the standard–of–care dosing schedule, but can be prevented by an optimized dosing schedule. We further show that the results are robust to changes in length of the treatment window, the step size parameter h and the payoff function weights.

5.4.3 Combination therapy

Lastly, we turn to the optimization of treatment schedules for combination therapy. We use the standard–of–care schedule as a starting point for our optimization with a dose on day 2 that we do not change, so $S_{N,0} = [23,\ 44,\ 65,\ 86]$, and $S_{I,0} = [23,\ 44,\ 65]$. As in Chapter 4, we set $V_N = 0.25$ grams and $V_I = 0.1$ grams, and the end of the treatment window ($t_{\max}$) at day 90. The gradient descent algorithm takes a long time to converge for this problem. To alleviate this issue we initially set $h = 0.1$ until the payoff drops below 10^4 at which point we set $h = 0.01$ to be consistent with the other therapies.

It takes 6,132 iterations for the algorithm to converge to the optimized dosing schedule where $S_N = [10.0,\ 90.0,\ 90.0,\ 90.0]$ and $S_I = [10.0,\ 24.9,\ 63.7]$. Fig. 5.10 shows the dosing schedule and the payoff over the iterations. When we compare this

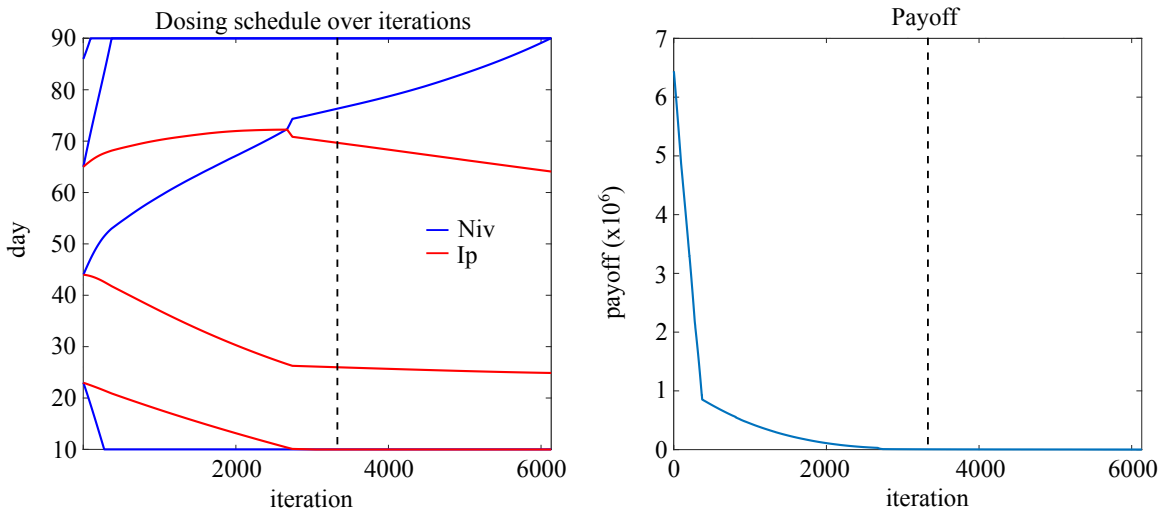


Figure 5.10: Left: the dosing schedule of combination therapy evolving over the iterations of the optimization algorithm. Right: the payoff evolving over the iterations of the optimization algorithm. The dashed line indicates the first iteration at which autoimmune myocarditis does not develop.

to the results for the heart–only model (see Fig. 4.10), we see that that the ICI doses are again more spread out over the treatment window, similarly to what we observe for nivolumab and ipilimumab monotherapy. As the three doses of ipilimumab were already regularly spaced out in the optimized dosing schedule for the heart–only model, the effect of adding the tumour compartment is less striking for this ICI than for nivolumab. Where the doses of nivolumab were all administered very late in the treatment window in Fig. 4.11, the optimized dosing schedule we obtained for the heart–tumour model administers one dose of nivolumab as early as possible on day 10, and three doses as late as possible, on day 90 (see Fig. 5.11). As we noted for the monotherapies, this spreading of ICI doses over the treatment window allows for a balance between the delay of tumour growth, which requires regular dosing, and preventing autoimmune myocarditis, which requires time spans without dosing so immune cell counts can decrease.

In Appendix B.3, we explore how the results discussed here change as the dose sizes V_N and V_I, the length of the treatment window, the step parameter h and the relative values of the payoff function weights w_1 and w_2 are varied. We show that there

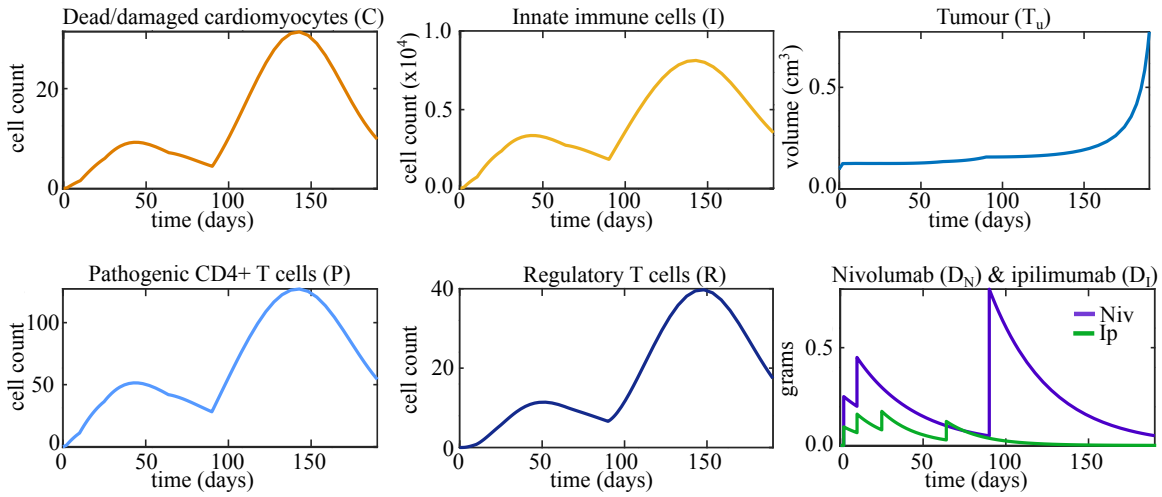


Figure 5.11: Time traces of model variables when the optimized dosing schedule for combination therapy is applied. The initial conditions are $[C_0,\ I_0,\ P_0,\ R_0,\ D_{N,0},\ D_{I,0},\ T_{u,0}] = [0,\ 0,\ 0,\ 0,\ 0,\ 0,\ 0.1]$, with the parameter values listed in Table 5.1 used.

exists a combined range of V_N and V_I for which autoimmune myocarditis develops under the standard–of–care dosing schedule, but can be prevented by an optimized schedule. We further show that the results are robust to changes in the length of the treatment window, the step size h, but that certain combinations of values for the weights w_1 and w_2 require careful calibration of the parameter h to ensure that the algorithm converges to a myocarditis–free solution when one exists.

In summary, in this section, we have shown that consideration of the effect of ICIs on tumour growth in addition to the risk of autoimmune myocarditis qualitatively changes what an optimized dosing schedule looks like. When only the prevention of autoimmune myocarditis is considered, ICI doses are delayed as long as possible because the primary objective of the payoff function is to decrease the levels of dead/damaged cardiomyocytes. When the delay of tumour growth is explicitly considered as well, ICI doses are spread out more evenly over the treatment window, striking a careful balance between regular dosing, so as to control tumour growth, and keeping ICI levels low enough such that the stimulation of the immune system does not cause the development of autoimmune myocarditis.

5.5 Discussion

In this chapter we have extended our model of autoimmune myocarditis by connecting it to a tumour growth model. This allows us to take a more holistic approach to optimizing cancer treatment where both the positive and negative effects of the ICIs are considered. We find that including the tumour compartment in our model has a significant impact on the treatment schedules compared to the optimized treatment schedules obtained in Chapter 4 where only prevention of autoimmune myocarditis was considered. When we solely consider the effect of ICIs on the development of autoimmune myocarditis, the optimized schedules mainly delay administering ICIs after an immune perturbation in order to allow immune cell counts to decrease. We showed in this chapter that optimizing for the delay of tumour growth alone does not result in dosing schedules that are significantly different from the standard–of–care. When we consider the effect of ICIs on the tumour and the heart simultaneously, however, we obtain schedules that spread out the doses as much as possible over the allowed treatment window, striking a balance between keeping tumour growth under control and preventing autoimmune myocarditis from developing.

More generally, this work shows that in designing treatment regimens it is essential to consider the effect of cancer treatment outside of, as well as within, the tumour microenvironment. If we only consider the effects of ICIs on the heart, we obtain dosing schedules that do not control tumour growth sufficiently. On the other hand, if we only consider the effects of ICIs on the tumour, the optimized dosing schedules cause autoimmune myocarditis to develop. We must thus consider both the positive and negative effects of ICI therapy explicitly to design treatment schedules that are both safe and effective.

In this chapter, our primary aim was to identify a schedule which prevents autoimmune myocarditis and delays tumour growth compared to the treatment–free case. We focused on achieving our objective by finding optimized, rather than opti-

mal, dosing schedules, and we did not seek to find the global minimum in the payoff landscape. In the future, the optimization framework we have presented here can be further developed to not only prevent myocarditis, but to potentially also maximize inhibition of tumour growth by finding the optimal dosing schedule, *i.e.* determining the optimal number of doses, dose size and timing of doses.

Due to a lack of data, it is currently not possible to further pin down biologically realistic values for the model parameters. As we have shown in Chapter 3, the steady state behaviour of the model of autoimmune myocarditis can be sensitive to changes in parameter values. Since the prevention of autoimmune myocarditis depends on where the separatrix is in state space, a change in parameter values that affects the steady states of the system would thus impact the results presented in this chapter. Furthermore, the model of tumour growth presented in this chapter shows the qualitative behaviour we would expect, but the quantitative predictions are unlikely to be biologically realistic. To facilitate parameter inference in the future, practical identifiability analysis could be applied to the model presented here to determine what type of data should be obtained to confidently estimate parameter values [140]. Estimating model parameter values directly from data will make the model output more biologically realistic and the predictions of optimized schedules more clinically relevant. The models for autoimmune myocarditis and tumour growth used in this work are simple and serve as a starting point from which more complex and biologically detailed models can be built. Including more variables, representing for example specific immune cell types or cytokines to more accurately represent the immunological mechanisms in the model of autoimmune myocarditis, and adding more detail to the tumour compartment, will allow for more realistic simulation of the system. In combination with sufficient patient data, the framework developed in this chapter for designing dosing schedules could then be used in the clinic to personalize cancer treatment.

Chapter 6

Conclusion

As cancer treatments improve and patient survival increases, the side–effects of these treatments, such as autoimmune myocarditis, become of increasing concern [66]. For autoimmune myocarditis in particular there is currently no pre–clinical screening assay to test potential new drug compounds for this side–effect, which, in combination with the difficulty in diagnosing autoimmune myocarditis, makes prevention of autoimmune myocarditis and its potentially fatal outcomes challenging [171]. Moreover, the immense complexity of the immune system makes it difficult to discern what the primary immunological mechanisms underlying the development and progression of autoimmune myocarditis are, and what effects ICIs have on these processes. Understanding this complexity, however, will be essential in designing an *in vitro* cardiotoxicity screening assay for new drug compounds, and in identifying safe and effective treatments regimens.

To improve our understanding of the immunological mechanisms underlying autoimmune myocarditis, and the effects of immune checkpoint inhibitors on its development and progression, we developed in this thesis the first mathematical model of this disease. By developing this model, we take a first step towards two goals:

1. determining which specific immune cell types and interactions form a minimal set that can accurately capture the pathology of autoimmune myocarditis in a mathematical model;

2. developing a pipeline to design safe and effective ICI treatment schedules.

From the first goal we can suggest a set of immune cell types which could be included in the initial design of an *in vitro* cardiotoxicity assay. The second goal is a step towards personalized medicine in the clinic where the risk of autoimmune myocarditis can be assessed on a patient–by–patient basis and treatment schedules can be adjusted accordingly to make ICI therapies safer.

In Chapter 2 we presented an overview of the immunological processes which experimental research has suggested are important in the development and progression of autoimmune myocarditis, as well as the mechanisms of action of two ICIs for which autoimmune myocarditis is a known side–effect: nivolumab and ipilimumab. The experimental literature suggests that autoimmune myocarditis is primarily CD4+ T cell–driven, so we focused our discussion on CD4+ T cells and their interactions with innate immune cells, such as macrophages, fibroblasts and polymorphonuclear leukocytes. We derived a network of interactions between T cells, innate immune cells, antigen–presenting cells and cardiomyocytes which summarizes the current understanding of the mechanisms underlying autoimmune myocarditis. This network forms the basis on which we build the first mathematical model of autoimmune myocarditis and, further, motivates possible extensions to this model in the future.

In Chapter 3, we developed the first mathematical model of autoimmune myocarditis based on the network of immunological interactions developed in Chapter 2. Our aim was to keep the model as simple as possible to allow the application of analytical approaches to gain insight into the model behaviour. To achieve this, we further simplified the network presented in Chapter 2, grouping together immune cell types that perform similar functions. Four groups of cells emerged: dead/damaged

cardiomyocytes that release the antigens driving the autoimmune reaction; innate immune cells that damage the myocardium; pathogenic CD4+ T cells that drive the autoimmune reactions; and regulatory T cells that inhibit the immune response. The small network describing the interactions between these four groups of cells was translated into a model consisting of a system of four time–dependent ODEs where the functional forms of the different terms in the equations were kept as simple as possible without sacrificing biological realism.

Clinical and experimental observations of autoimmune myocarditis motivated the hypothesis that any model of this disease must have three non–negative steady states: a stable "healthy" steady state with low levels of immune cells and damage to the myocardium; a stable "diseased" steady state with high levels of immune cells and damage to the myocardium; and a saddle point which lies on a separatrix which separates the basins of attraction of the healthy and diseased steady states and provides a threshold below which the immune response can be controlled and above which the immune response results in autoimmune myocarditis. We identified model parameters that guaranteed the existence of these three non–negative steady states in an ICI–free environment. We further showed that when ICIs are introduced in the system, the saddle point, together the separatrix, moves towards the "healthy" steady state, thus lowering the threshold for immune perturbations to result in the development of autoimmune myocarditis, which agrees with the experimental literature on the effects of ICIs.

In the absence of experimental data, the model was parameterized by performing a parameter sweep based on five literature–derived criteria and the steady state behaviour we hypothesized a model of autoimmune myocarditis should exhibit. To determine what type and quality of data should be collected to enable us to estimate parameter values in the future, we applied parameter identifiability analysis to the model. We found that all four state variables of the model must be observed for the

model parameters to be identifiable, but computational limits prevented a detailed exploration of the requirements on the data set to ensure practical identifiability of all model parameters.

In Chapter 4, we applied optimal control theory to our model of autoimmune myocarditis to answer the question: if a patient experiences a perturbation in the number of cardiac–specific pathogenic CD4+ T cells while on standard–of–care treatment with ICIs, can we adjust the scheduling of the subsequent doses of ICIs such that the patient does not develop autoimmune myocarditis? In our framework, we conclude that autoimmune myocarditis develops as a consequence of a given dosing schedule if, at the final time point, all cell counts in the system are above the cell counts at the ICI–free saddle point. To stimulate an immune response, we considered a perturbation to the system at the "healthy" steady state consisting of a finite number of pathogenic T cells. The payoff function that we minimize by gradient descent depends only on the number of dead/damaged cardiomyocytes in the system at the end of a simulation, thus favouring dosing schedules that do not cause autoimmune myocarditis to develop. Under the assumption that a given number of ICI doses are administered, we used optimal control approaches to optimize the timing of these doses in order to prevent the development of autoimmune myocarditis. This is a more biologically and clinically realistic optimal control problem compared to using optimal control approaches to directly control the concentration of ICIs in the patient.

For monotherapy with nivolumab or ipilimumab, as well as combination therapy, we showed that autoimmune myocarditis does not develop when ICIs are administered over a range of optimized dosing schedules whereby the patient receives the same number of ICI doses as the standard–of–care treatment regime. Compared to the standard–of–care dosing schedule, the common trend among all the myocarditis–free schedules is a delay in dosing with ICIs after the immune perturbation occurs to allow the number of immune cells in the heart to decrease and inflammation to be resolved.

This suggests that regular testing of patients for increased counts of active, cardiac antigen–specific immune cells will be an essential part of any strategy to prevent autoimmune myocarditis as a side–effect of ICIs using the framework presented in this chapter.

Lastly, in Chapter 5, we extended our model of autoimmune myocarditis by connecting it to a tumour growth model, enabling a more holistic approach to optimizing cancer treatment where both the positive and negative effects of the ICIs are considered. The payoff function to be minimized reflects this by counting the number of dead/damaged cardiomyocytes as well as tumour volume over time. Using the same optimal control approach as Chapter 4, we find that including the tumour compartment has a significant impact on the optimized treatment schedules. When we solely consider the effect of ICIs on the development of autoimmune myocarditis, the optimized schedules mainly delay administering ICIs after an immune perturbation to allow immune cell counts to decrease. Optimizing for inhibition of tumour growth alone on the other hand does not result in dosing schedules that are significantly different from the standard–of–care. However, when we consider the effect of ICIs on the tumour and the heart simultaneously, we obtain dosing schedules that spread out the doses over the treatment window, striking a balance between keeping tumour growth under control and preventing autoimmune myocarditis from developing.

6.1 Future work

In this thesis we have taken important first steps towards answering the two questions we posed. This work motivates many avenues of research that may be pursued in the future to further develop the mathematical model of autoimmune myocarditis and the framework for optimizing ICI dosing schedules.

The immunology presented in Chapter 2 is a necessary simplification of biological

reality as the complexity of the immunological mechanisms underlying autoimmune myocarditis is immense, and many of the underlying mechanisms are incompletely understood. Many details on the functions of immune cells, cell–cell interactions, and cytokines have necessarily been omitted to present a summary that is useful as a basis for the development of a simple first mathematical model. As the field of experimental myocarditis research develops, and our understanding of the immune system improves, the network of interactions presented in Chapter 2 will need to be refined. Although the overview provided here is sufficiently detailed to develop the first mathematical model of autoimmune myocarditis, future iterations of the model will need to include more detailed interactions between different specific cell types in order to suggest which specific immune cell types should be included in an *in vitro* cardiotoxicity assay for potential new drug compounds. Our mathematical model of autoimmune myocarditis captures the desired qualitative behaviour displayed by this disease but its state variables currently represent classes of immune cell types rather than individual immune cell types. To suggest specific immune cell types for the design of an *in vitro* assay our model can be expanded by including, for example, macrophages, fibroblasts and polymorphonuclear leukocytes separately instead of grouping them together as innate immune cells.

A model of autoimmune myocarditis that includes the dynamics of specific immune cell types and shows biologically realistic behaviour can be used to suggest an initial design for an *in vitro* cardiotoxicity assay. If a given set of immune cells can reproduce the characteristics of the disease we observe in patients *in silico*, this set of cells could also be sufficient to reproduce the risk of cardiotoxicity *in vitro* and can thus serve as a starting point for the design of the *in vitro* assay.

To be able to use our mathematical model of autoimmune myocarditis to answer the two questions posed above, we must first overcome the current inability to validate the model. There is, to our knowledge, no *in vitro* or *in vivo* data available, whether

from animal or human studies, with which the model can be validated. Resolving this problem will be challenging as autoimmune myocarditis is rare and difficult to diagnose, hence clinical data will be nearly impossible to obtain, and an *in vitro* assay to generate data with is not currently available. Validation with mouse or other animal data could potentially be an option, although one should then consider to what extent a model validated with such data is applicable to humans.

Another major challenge related to model validation is the difficulty in determining what type of data would be required to confidently estimate the parameter values in the model. Currently, performing practical identifiability analysis with the full model is computationally infeasible, as shown in Chapter 3, and we cannot determine the number of data points required, the time points at which they should be collected, or the level of noise that can be tolerated, to allow confident estimation of the parameter values.

However, ensuring that we have access to data with which the model parameter values can be confidently estimated is essential for validating the model, as well as developing it further in combination with the *in vitro* cardiotoxicity assay. Furthermore, it is important for moving towards use of the model in combination with the framework developed in Chapters 4 and 5 in personalized healthcare. This would require that we can determine parameter values based on patient data, so we must know the quality and quantity of data that needs to be collected in the clinic. Resolving the challenges around model validation and practical parameter identifiability analysis will require exploration of different options for data collection, as well as more advanced techniques for parameter identifiability analysis and improvement of the computational costs of existing algorithms.

Once we have overcome the issues surrounding validation and parameter identifiability, we can work towards a mathematical model of autoimmune myocarditis that can be used to inform an initial experimental design for an *in vitro* cardiotoxicity as-

say. We can then initiate a model–predict–refine cycle where the mathematical model and *in vitro* assay can be used to verify and improve each other. The mathematical model can be used to suggest an initial design for the cardiotoxicity assay, which can be performed *in vitro*. The data collected from these experiments can be compared with the model predictions, and based on this comparison the model can be refined. The refined model may suggest adjustments to the experimental set–up, which can be tested again, and so forth. This leads to a simultaneous improvement of both the *in vitro* assay and the mathematical model.

The work in Chapters 4 and 5 showed that it is essential to consider both the positive and negative effects of ICI treatment when designing safe and effective dosing schedules. When the negative effects of ICIs in the heart are considered explicitly in treatment schedule design, this alters the dosing schedules significantly compared to standard–of–care schedules. If we consider the positive and negative effects of ICIs simultaneously, we obtain dosing schedules that are different from the standard–of–care *and* different from when negative effects of ICIs are considered alone. These results suggest that we must take a more holistic approach to ICI treatment schedule design. We must consider the positive and negative effects of ICIs explicitly in order to obtain dosing schedules that are both safe and effective.

The heart is not the only organ or tissue affected by off–target effects of cancer therapy. The gastrointestinal tract, skin and liver are just a few of the other organs that are known to experience side–effects of ICI therapy [117]. To model the range of effects of ICI therapy on the whole body and take a holistic approach to treatment optimization, the heart–tumour model must be extended with compartments representing other parts of the body that are affected by off–target effects of ICI treatment. Such a modelling framework also allows for the inclusion of other types of therapy, such as radiation therapy, which are administered to patients alone or in combination with ICIs or other therapies. A model which includes all effects of cancer treatment

in combination with the optimal control framework we have presented in Chapters 4 and 5 will be an important step towards the design of personalized dosing schedules for cancer therapies in the clinic.

In this thesis we have developed the first mathematical model of autoimmune myocarditis, which, when combined with an optimal control framework, can be utilized to determine ICI treatment regimes that can control tumour growth while minimizing the risk of autoimmune myocarditis developing as a side–effect. Pursuing each of the avenues of research presented above will move us ever closer to safe and effective personalized cancer treatments for everyone.

Appendices

Appendix A

Exploration of optimal control parameters for autoimmune myocarditis model

In this appendix, we explore the effects of varying the values of the parameters used in the optimal control algorithm with which treatment is optimized in Chapter 4. We show that the optimization results are robust to changes in the gradient descent step size, h. We also explore how increasing the length of the treatment window affects the optimized dosing schedule identified by the algorithm and the effect of varying the ICI doses V_N and V_I and the parameters γ_3 and γ_4 which govern the pharmacodynamics of the ICIs in the model. Furthermore, we explore what happens when we change the functional form of the payoff function and minimize the integral of the dead/damaged cardiomyocytes over time instead of only considering the number of these cells at the final time point. This takes into account not only if the patient ultimately develops autoimmune myocarditis, but the total damage done to the cardiac tissue as a result of the immune perturbation and the application of ICIs. Lastly, we investigate the effect of the size of the immune perturbation and the timing of the immune perturbation

relative to the ICI doses.

Unless otherwise noted, we use the default parameter values used in Chapter 4. The default value of V_N is $V_N = 0.25$ grams, the default value of V_I is $V_I = 1.0$ grams for ipilimumab monotherapy and $V_I = 0.1$ grams for combination therapy. The final day of the treatment window, $t_{\max}$, is by default day 60 for nivolumab monotherapy, day 70 for ipilimumab monotherapy and day 90 for combination therapy. The default values of the other parameters are $h = 0.5$, $\gamma_3 = 1$ grams^{-1} and $\gamma_4 = 1$ grams^{-1}.

A.1 Nivolumab monotherapy

A.1.1 Time traces for varying dosing schedules

To illustrate the wide range of dosing schedules that do not lead to autoimmune myocarditis, we include here a set of time traces for simulations with dosing schedules which all prevent autoimmune myocarditis from developing, but differ significantly in timing of nivolumab administration (see Fig. A.1). These dosing schedules are achieved after 50, 100 and 150 iterations when the optimization is initialized with the standard–of–care treatment schedule.

A.1.2 The effect of changes to V_N

The value of $V_N = 0.25$ grams was chosen because it causes autoimmune myocarditis to develop when a standard–of–care nivolumab dosing schedule is applied (see Section 4.2.1). We explore the effect of the choice of V_N by computing the payoff of the optimized drug regime while varying V_N between 0.01 grams and 0.5 grams. We initialize the schedule at the standard–of–care when $V_N \leq 0.22$. For $V_N = m$, $m = 0.23, 0.24, \ldots, 0.36$ we initialize the schedule at the optimized schedule for $V_N = m - 0.01$ to speed up convergence. For $V_N \geq 0.37$ grams, we initialize at the schedule $S_0 = [10.0, 10.0, 10.0, 10.0]$ to ensure that the algorithm converges towards

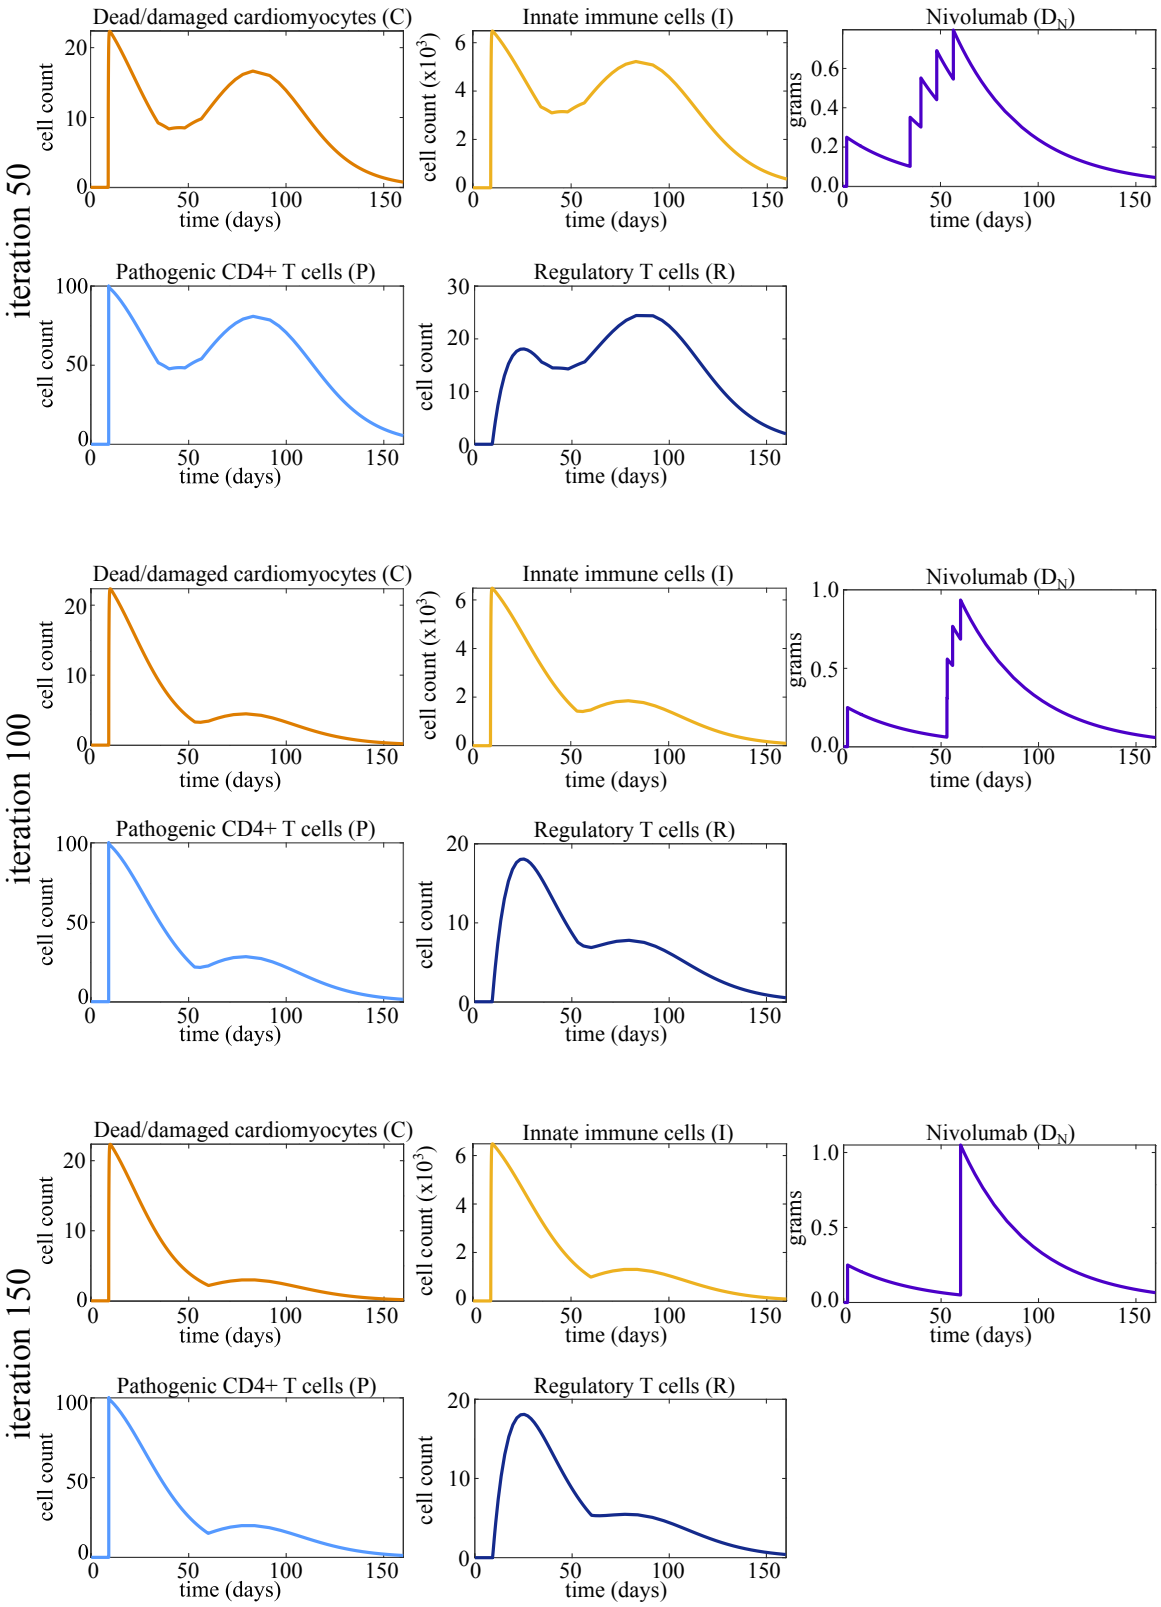


Figure A.1: Time traces of state variables when nivolumab dosing schedules are obtained after 50 (top), 100 (middle) or 150 (bottom) iterations. None of these dosing schedules lead to autoimmune myocarditis. Default parameter values are used.

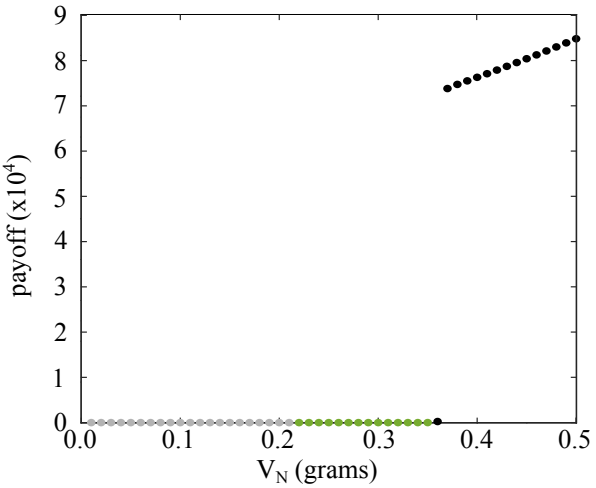


Figure A.2: The payoff of the optimized treatment schedule for nivolumab monotherapy over different values of V_N. The grey dots indicate values of V_N for which autoimmune myocarditis never develops. The green dots indicate values of V_N for which autoimmune myocarditis is side–effect of standard–of–care nivolumab therapy but it can be prevented by an optimized dosing schedule. The black dots indicate values of V_N for which autoimmune myocarditis cannot be prevented within the constraints set here.

the local minimum. All other parameters are set to their default values.

The payoff of the optimized dosing schedule is plotted as a function of V_N in Fig. A.2. We find that the value of V_N influences the outcome of ICI therapy significantly. For $V_N \leq 0.21$ grams, myocarditis does not develop when the standard–of–care schedule is applied. For values of V_N in the range 0.22–0.35 grams, autoimmune myocarditis does occur as a side–effect of standard–of–care treatment but it can be prevented by optimizing the dosing schedule. For $V_N \geq 0.36$ grams, autoimmune myocarditis occurs as a side–effect of the standard–of–care treatment schedule and cannot be prevented by optimizing the schedule within the constraints set here.

In Fig. A.3, we show the times traces when the optimized schedule for $V_N = 0.4$ grams is applied, with default parameter settings as described above and in Chapter 4. Here we see that, with the chosen parameter set and constraints, myocarditis cannot be avoided. Note that for values of $V_N \geq 0.37$ grams, the optimized dosing schedule is to apply all doses on the first day of the treatment window, day 10, as shown in

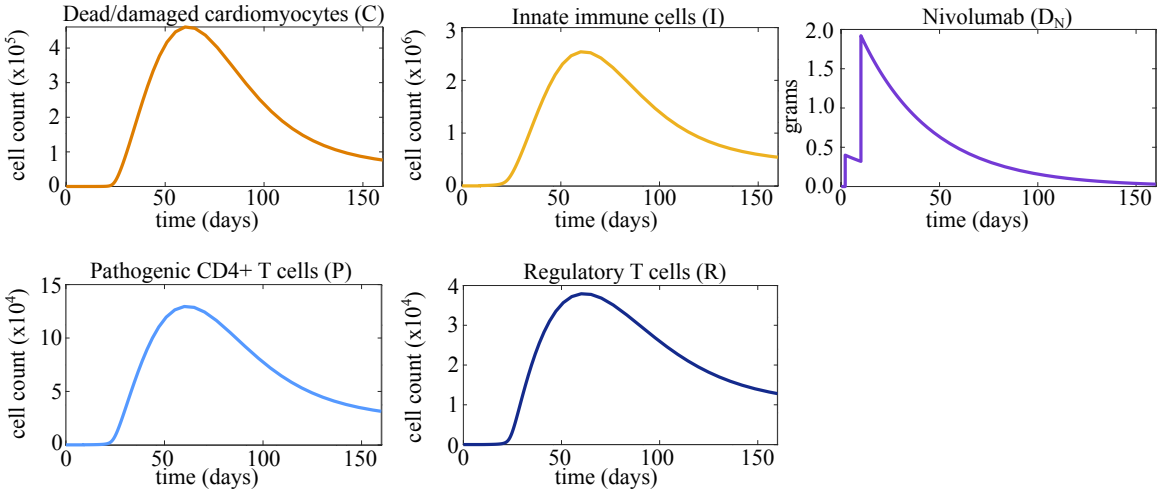


Figure A.3: Time traces for the optimized dosing schedule for $V_N = 0.4$ grams, with default settings for all other parameters.

Fig. A.3. This is because an early dose gives the system more time between the time point of dosing and the final time point of the simulation to approach the disease steady. As shown in Fig. A.3, the number of dead/damaged cardiomyocytes at the diseased steady state is less than the number of dead/damaged cardiomyocytes at the peak value after dosing. Since only the final count of dead/damaged cardiomyocytes is included in the payoff function, it is thus beneficial in this scenario where myocarditis is inevitable to dose sooner rather than later to give allow more time for cell counts to decrease and the system to approach the diseased steady state.

A.1.3 The effect of extending the treatment window

Next, we explore the effect of varying the final time point of the treatment window, $t_{\max}$. The default setting for this parameter for nivolumab monotherapy is day $t_{\max} = 60$. We explore what happens to the optimal dosing schedule when the treatment window is extended, up to and including a treatment window which ends on day $t_{\max} = 110$. We use the standard–of–care schedule as the initial schedule for $t_{\max} = 60$, and thereafter set the initial schedule for $t_{\max} = m$, $m = 61, 62, \ldots, 110$, to be the optimized schedule for $t_{\max} = m - 1$. All other model and algorithm parameters are

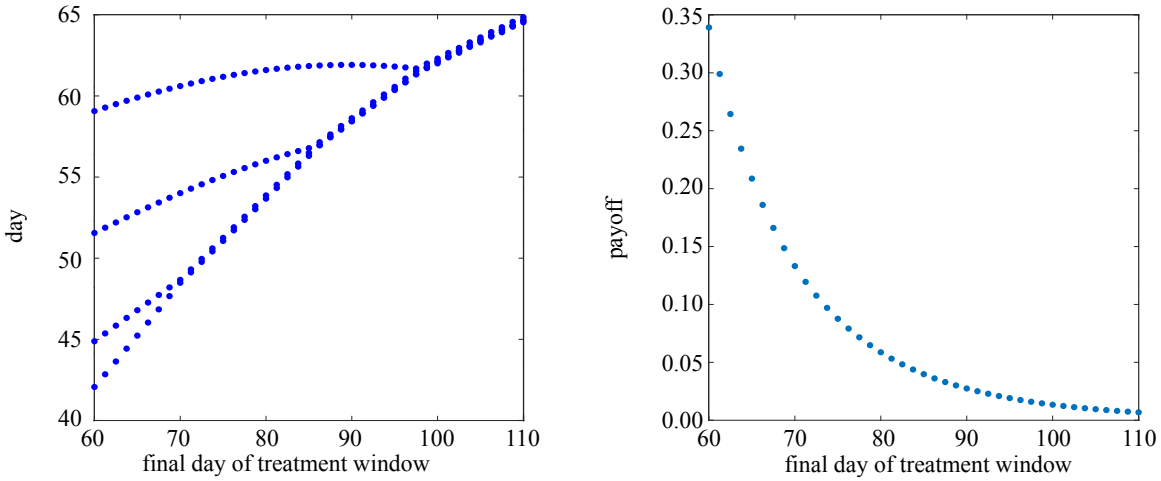


Figure A.4: Optimized nivolumab dosing schedule (left) and payoff (right) for different lengths of the treatment window. Default parameter values are used unless otherwise noted.

kept at their default values. The optimized solutions that prevent disease for shorter treatment windows will also not lead to myocarditis when the treatment window is extended. However, extending the treatment window might lead the algorithm to converge to even better solutions as the increased length of time over which ICI doses can be administered can allow solutions which decrease the payoff even further.

The results of this analysis are shown in Fig. A.4. As the treatment window is extended, optimal solutions of increasingly lower payoff are found. As the left–hand panel of Fig. A.4 shows, the doses of nivolumab are administered increasingly later and closer together as the treatment window is extended, though doses are never administered after day 65. When $t_{\max} = 60$, the four doses are administered separately, whereas for $t_{\max} \geq 99$ all four doses are administered almost simultaneously. The length of the treatment window thus qualitatively affects the optimized dosing schedule, but, as the right–hand panel of Fig. A.4 shows, the decrease in the payoff as the treatment window is extended is tiny compared to the payoff difference between the standard–of–care schedule and any of these optimized schedules. In terms of achieving our objective of preventing autoimmune myocarditis, the length of the treatment window is thus not important. If other factors, such as the effect of ICIs on

h	Schedule nivolumab	Payoff
0.1	34.5, 40.1, 48.3, 57.0	0.7298
1	48.0, 48.3, 53.7, 60.0	0.2519
10	60.0, 60.0, 60.0, 60.0	0.1561

Table A.1: The optimized dosing schedule of nivolumab monotherapy and the payoff vary as the size of the step size parameter, h, is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

the tumour or other side–effects, are considered, the length of the treatment window might be relevant to consider.

A.1.4 The effect of changes to h

We need to ensure that our results are robust when the gradient descent step parameter, h, is varied. The default value for this parameter is $h = 0.5$. We now vary the value of h while keeping all other parameters at their default values for nivolumab monotherapy and initializing all optimizations with the standard–of–care schedule. The results in Table A.1 shows that optimization of the dosing schedule for nivolumab monotherapy is robust to changes in h in terms of achieving our objective of identifying a dosing schedule that does not cause autoimmune myocarditis to develop.

A.1.5 The effect of varying γ_3

We now explore the effect on the model simulations of nivolumab monotherapy when the value of γ_3 is varied. This parameter translates the mass of nivolumab in the system to the effect on the deactivation rate of pathogenic CD4+ T cells, and captures two effects. The first is the direct pharmacodynamic effect of nivolumab on the deactivation rate of pathogenic CD4+ T cells, and the second is how a dose of nivolumab affects this deactivation rate relative to the effect of that same dose on the activation of T cells. The second effect is governed by the ratio γ_1/γ_3. The case

where γ_1 and γ_3 are varied simultaneously is covered by the analysis in Section A.1.2 where the value of V_N is varied. We explore the consequences of varying the relative effect of a dose nivolumab on these two immunological processes here by varying the value of γ_3 while keeping γ_1 constant at its default value of one grams^{-1}.

We first consider the effect that varying γ_3 has on the outcome of standard–of–care treatment. We keep $V_N = 0.25$ grams. When $\gamma_3 = 0.1$ grams^{-1}, i.e. when the effect of nivolumab on T cell activation is ten times stronger than its effect on pathogenic CD4+ T cell deactivation, standard–of–care treatment does not cause autoimmune myocarditis to develop. When $\gamma_3 = 10$ grams^{-1}, i.e. when the effect of nivolumab on pathogenic CD4+ T cell deactivation is ten times stronger than its effect on T cell activation, standard–of–care treatment does cause autoimmune myocarditis to develop.

Since optimization of the treatment schedule is no longer necessary when $\gamma_3 = 0.1$ grams^{-1}, we now only explore for $\gamma_3 = 10$ grams^{-1} whether an optimized schedule can be designed which prevents autoimmune myocarditis from developing. We initialize the dosing schedule at the standard–of–care schedule and keep all other parameters at their default values. The optimized dosing schedule administers all doses of nivolumab on day $t_{\min} = 10$, the first day of the treatment window, and this schedule still causes autoimmune myocarditis. We thus have three different outcomes for three different values of γ_3. This suggests that it will be important to assess what the relative effects of nivolumab on the different immunological processes is, as this has significant consequences for whether or not a myocarditis–free dosing schedule can be identified within the constraints used here.

A.1.6 The effect of changes to the payoff function

In the default payoff function, only the number of dead/damaged cardiomyocytes at the final time point of the simulation is taken into consideration:

$$J = C(t_f). \tag{A.1}$$

This could allow for solutions where a large peak of dead/damaged cardiomyocytes occurs at an earlier time point, signifying significant damage to cardiac tissue. Although it is unlikely for such a peak to occur without the patient developing autoimmune myocarditis and the number of dead/damaged cardiomyocytes at the final time point also being high, we explore the effect on the results when the number of dead/damaged cardiomyocytes is integrated over time in the payoff function:

$$J = \int_0^{t_f} C(t)\mathrm{d}t. \tag{A.2}$$

All parameter values are kept at their default values, and the initial dosing schedule is the standard–of–care schedule. Fig. A.5 shows the result of the optimization. The optimized dosing schedule applies all doses of nivolumab on the last day of the treatment window. Although this was not the optimized schedule with the default payoff function, the evolution of the dosing schedule and the payoff over the iterations of the algorithm follow a nearly identical trend where all ICI doses are delayed, and many different dosing schedules prevent autoimmune myocarditis, as shown by the large flat region when the payoff is plotted over the iterations. Whether the count of dead/damaged cardiomyocytes is integrated over time or only assessed at the final time point thus does not appear to have a significant effect on the results of the optimization of the dosing schedules for nivolumab.

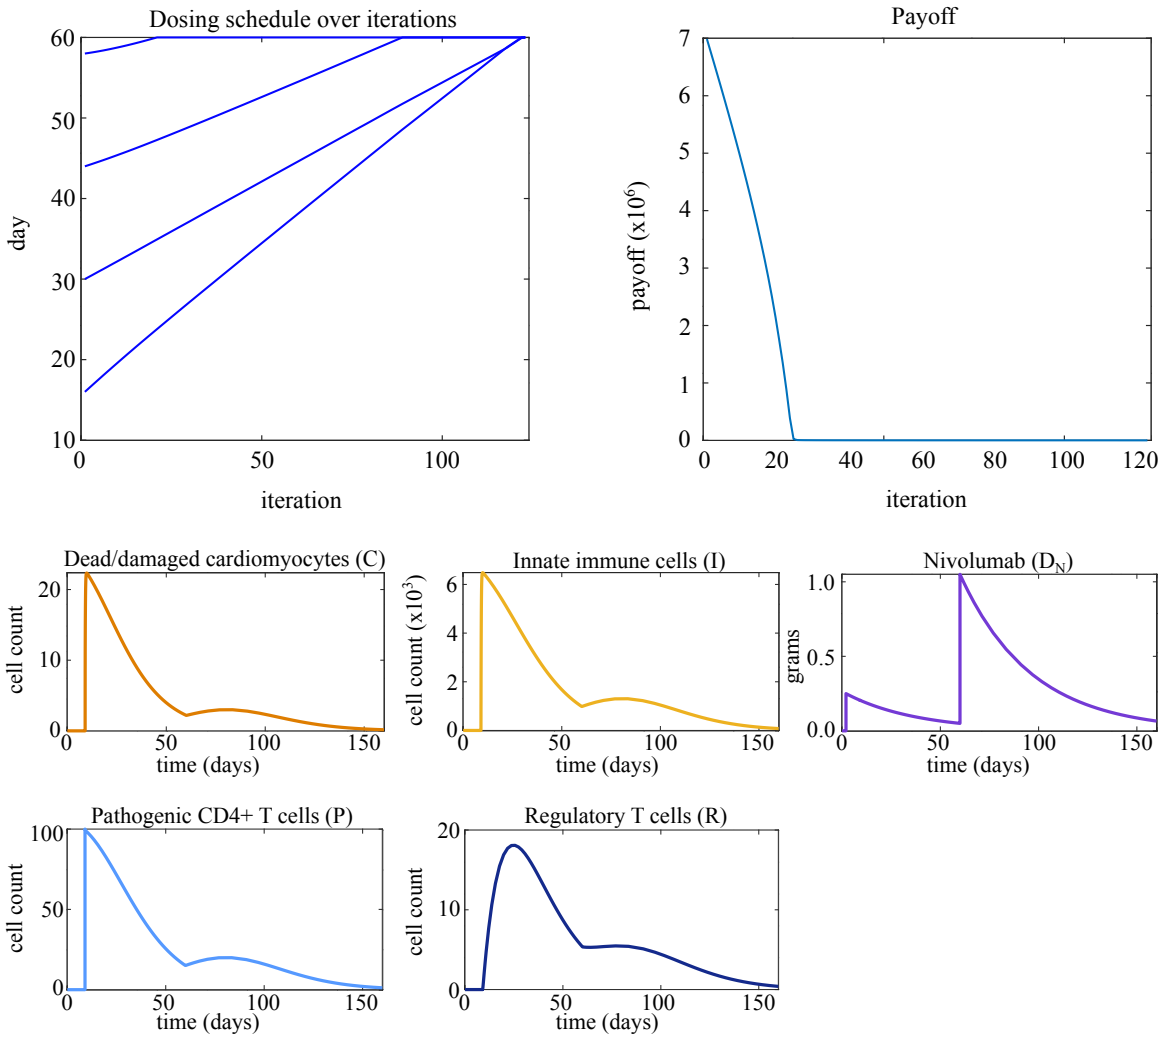


Figure A.5: Top: evolution of the nivolumab dosing schedule and payoff over iterations when the number of dead/damaged cardiomyocytes is integrated over time in the payoff function (see Eq. (A.2)). Bottom: time traces of state variables when the optimized nivolumab dosing schedule is applied in the case where the number of dead/damaged cardiomyocytes is integrated over time in the payoff function (see Eq. (A.2)). All default parameter values are used.

A.1.7 The effect of the timing of the immune perturbation and detection

In the scenario considered in Chapter 4, an immune perturbation consisting of 100 pathogenic CD4+ T cells occurs on day 9 to simulate the immune perturbations that a patient being treated with ICIs might experience. We now consider the effects of the timing of the immune perturbation on the optimization of the dosing schedule.

In standard–of–care treatment with nivolumab, a dose is applied every two weeks. In between doses, the drug is slowly cleared from the body, but not all of it is cleared within two weeks as the half–life of nivolumab in the body is 25 days [27]. Thus, as the second dose is applied, there is a small amount of nivolumab still present and the level of nivolumab in the body after the second dose is thus higher than after the first dose. This remains true for approximately the first six doses of nivolumab, after which the level of nivolumab is sufficiently high that the clearance rate balances the amount that is administered in a single dose. As the level of nivolumab influences the position of the separatrix and thus the threshold for an immune response to lead to autoimmune myocarditis, when the immune perturbation occurs across the first six doses of nivolumab could potentially have an impact on whether or not autoimmune myocarditis develops.

To explore this we let the immune perturbation occur on day 37, thus applying three doses of nivolumab prior to the immune perturbation, with the third dose being administered on day 30. We find that when the immune perturbation occurs on day 37 autoimmune myocarditis can still be prevented by an optimized schedule despite the separatrix being lower at the time when the immune system is perturbed compared to when the immune perturbation occurs on day 9. However, a longer treatment window is required. When the immune perturbation happens on day 37, the subsequent four doses that are optimized are applied on days 44, 58, 72 and 86 in the standard–of–care schedule. We find that the treatment window needs to be extended to at least

day 96 in order for an optimized dosing schedule which does not lead to autoimmune myocarditis to be possible. The optimized schedule is then to apply all four doses on the last day of the treatment window.

It is unlikely that an immune perturbation is detected immediately. Therefore, we investigate the effects on the optimization of the dosing schedule when the immune perturbation occurs on day 9 but is not detected until after the third dose has been applied on day 30. In this set–up we optimize the subsequent four doses which are applied on days 44, 58, 72 and 86 in standard–of–care treatment. We find that autoimmune myocarditis can be prevented in this scenario if all doses of nivolumab are applied on day 97 or later. Thus, similar to delaying the immune perturbation itself, delayed detection of the immune perturbation requires a longer treatment window in order to find a dosing schedule for nivolumab which does not lead to autoimmune myocarditis.

A.2 Ipilimumab monotherapy

A.2.1 The effect of changes to V_I

We explore the effect of varying V_I on the optimization of the dosing schedule for ipilimumab monotherapy. The default value of V_I is 1.0 grams for ipilimumab monotherapy. To make the optimization more efficient and ensure that the gradient descent algorithm converges towards a myocarditis–free local minimum, when one exists, we initialize the dosing schedule with the standard–of–care schedule $S_0 = [23,\ 44,\ 65]$ when $V_I < 0.95$ grams and with $S_0 = [10,\ 10,\ 10]$ if $V_I \geq 1.52$ grams. For $V_I = m$, $m = 0.98, 1.01, \ldots, 1.49$, we initialize the schedule at the optimized schedule for $V_I = m - 0.03$. The payoff as a function of V_I is shown in Fig. A.6. For $V_I < 0.95$ grams, the standard–of–care dosing schedule does not lead to autoimmune myocarditis. For $0.95 \leq V_I \leq 1.49$ grams, the standard–of–care dosing schedule does lead

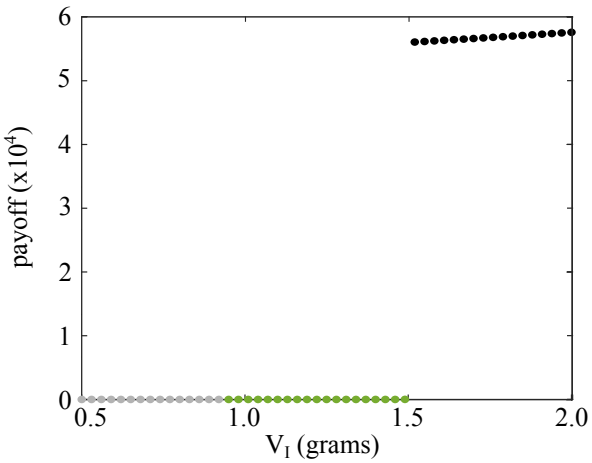


Figure A.6: Payoff of the optimized dosing schedule as a function of V_I. All parameters are set to default values unless otherwise noted. The grey dots indicate values of V_I for which autoimmune myocarditis never develops. The green dots indicate values of V_I for which autoimmune myocarditis is side–effect of standard–of–care ipilimumab therapy but it can be prevented by an optimized dosing schedule. The black dots indicate values of V_I for which autoimmune myocarditis cannot be prevented within the constraints set here.

to autoimmune myocarditis but an optimized dosing schedule which does not can be identified. For $V_I \geq 1.52$ grams, no dosing schedule can be found within the constraints of the default set–up and parameter values which does not lead to autoimmune myocarditis.

A.2.2 The effect of extending the treatment window

We now explore the effect on the results of the optimization when the length of the treatment window is increased. The last day of the treatment window, $t_{\max}$, is day 70 by default. We explore what happens to the optimal dosing schedule when the treatment window is extended, up to and including a treatment window which ends on day $t_{\max} = 120$. We use the standard–of–care schedule as the initial schedule for $t_{\max} = 70$, and thereafter set the initial schedule for $t_{\max} = m$, $m = 71, 72, \ldots, 120$, to be the optimized schedule for $t_{\max} = m-1$. Default parameter settings are used unless otherwise noted. Fig. A.7 shows the results as the treatment window is extended to

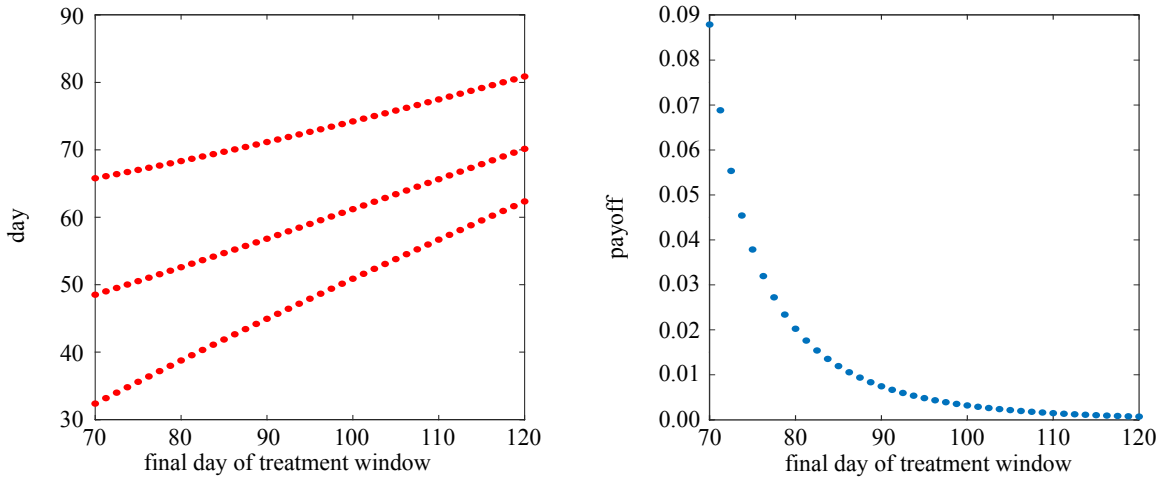


Figure A.7: The optimized dosing schedule (left) and payoff (right) as a function of the length of the treatment window for ipilimumab monotherapy. All default parameter settings are used, unless otherwise noted.

up to day $t_{\max} = 120$. As the length of the treatment window increases, the doses in the optimized dosing schedules are administered later and closer together. Similar to what we saw for nivolumab monotherapy, the length of the treatment window thus qualitatively affects the optimized dosing schedule, but the decrease in the payoff as the treatment window is extended is tiny compared to the payoff difference between the standard–of–care schedule and any of these optimized schedules. In terms of achieving our objective of preventing autoimmune myocarditis, the length of the treatment window is thus not important. If other factors, such as the effect of ICIs on the tumour or other side–effects, are considered, the length of the treatment window might be relevant to consider.

A.2.3 The effect of h

We need to ensure that our results are robust when the gradient descent step parameter, h, is varied. The default value for this parameter is $h = 0.5$. We now vary the value of h while keeping all other parameters at their default values for ipilimumab monotherapy and using the standard–of–care dosing schedule as the initial schedule

h	Schedule ipilimumab	Payoff
0.1	29.5, 46.9, 65.0	0.2001
1	34.8, 49.9, 66.4	0.0624
10	60.1, 67.1, 70.0	0.0221

Table A.2: The optimized dosing schedule of ipilimumab monotherapy and the payoff vary as the size of the step size parameter, h, is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

for all three optimizations. The results in Table A.2 show that optimization of the dosing schedule for ipilimumab monotherapy is robust to changes in h in terms of achieving our objective of preventing autoimmune myocarditis from developing.

A.2.4 The effect of varying γ_4

We explore the effect on the model simulations of ipilimumab monotherapy when the value of γ_4 is varied. This parameter translates the mass of ipilimumab in the system to the effect on the deactivation rate of regulatory T cells, and captures two effects. The first is the direct pharmacodynamic effect of ipilimumab on the deactivation rate of regulatory T cells, and the second is how a dose of ipilimumab affects this deactivation rate relative to the effect of that same dose on the activation of T cells. The second effect is governed by the ratio γ_2/γ_4. The case where γ_2 and γ_4 are varied simultaneously is covered by the analysis in Section A.2.1 where the value of V_I is varied. Here, we explore the consequences of varying the relative effect of a dose of ipilimumab on these two immunological processes by varying the value of γ_4 while keeping γ_2 constant at its default value of one grams^{-1}.

We first consider the effect that varying γ_4 has on the outcome of standard–of–care treatment. We keep $V_I = 1.0$ grams. In the case where $\gamma_4 = 0.1$ grams^{-1} as well as when $\gamma_4 = 10$ grams^{-1}, the standard–of–care treatment schedule for ipilimumab monotherapy causes autoimmune myocarditis to develop.

We optimize the ipilimumab treatment schedule for $\gamma_4 = 0.1$ grams^{-1} and $\gamma_4 = 10$

grams^{-1}, initializing the dosing schedule at the standard–of–care schedule and keeping all other parameters at their default values. For $\gamma_4 = 0.1$ grams^{-1} as well as $\gamma_4 = 10$ grams^{-1}, we see that the algorithm identifies a myocarditis–free dosing schedule. For $\gamma_4 = 0.1$ grams^{-1}, the optimized dosing schedule is $S = [29.9,\ 47.6,\ 66.7]$. For $\gamma_4 = 10$ grams^{-1}, the optimized dosing schedule is $S = [36.5,\ 49.9,\ 63.8]$. These schedules are qualitatively similar to the optimized dosing schedule for $\gamma_4 = 1$ grams^{-1}. The optimized ipilimumab monotherapy schedule is thus not sensitive to the relative strength of the effect of ipilimumab on T cell activation compared to regulatory T cell deactivation.

A.3 Combination therapy

A.3.1 The effect of changes to V_N and V_I

We explore the effect of varying V_N and V_I simultaneously on the optimization of the dosing schedule for combination therapy. As opposed to the monotherapies explored above, we now consider a two–dimensional parameter space where we plot outcome over the values of V_N and V_I. We vary V_N across the range 0.1–0.5 grams and V_I across the range 0–0.5 grams. For each combination of values of V_N and V_I we set all other parameters to their default values and use the standard–of–care dosing schedule as the initial schedule for the optimization.

Fig. A.8 shows treatment outcome as a function of the values of V_N and V_I. We distinguish three regions in this plot space: (1) combinations of V_N and V_I for which autoimmune myocarditis does not develop when the standard–of–care schedule for combination therapy is applied (white in Fig. A.8); (2) combinations of V_N and V_I for which autoimmune myocarditis does develop when the standard–of–care schedule for combination therapy is applied, but an optimized dosing schedule for which this side–effect is prevented can be identified (green in Fig. A.8); and (3) combinations of

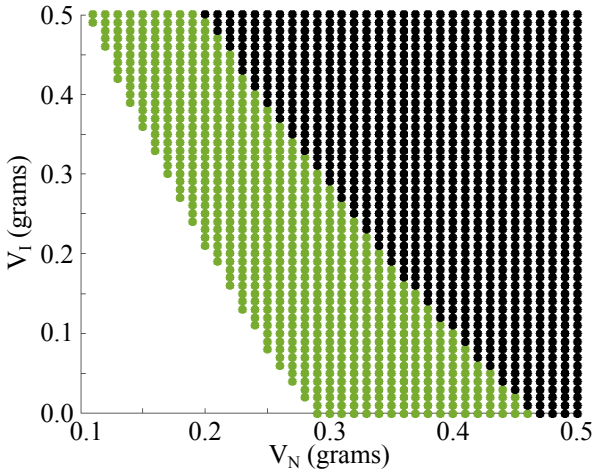


Figure A.8: Treatment outcome as a function of the values of V_N and V_I in combination therapy. For all other parameters the default values are used. In the white region of values of V_N and V_I autoimmune myocarditis does not develop as a side–effect of standard–of–care treatment. In the green region of values of V_N and V_I autoimmune myocarditis does develop as a side–effect of standard–of–care treatment but an optimized schedule which prevents this side–effect can be identified. In the black region of values of V_N and V_I autoimmune myocarditis cannot be prevented.

V_N and V_I for which autoimmune myocarditis does develop when the standard–of–care schedule for combination therapy is applied, and no alternative dosing schedules for which this side–effect is prevented can be identified (black in Fig. A.8).

A.3.2 The effect of extending the treatment window

We now explore how extending the treatment window affects the optimized dosing schedule. The default value of $t_{\max}$ for combination therapy is day 90. We now increase $t_{\max}$ from day 90 to day 125 in steps of one day. All optimizations are initialized with the standard–of–care schedules $S_{0,N} = [23,\ 44,\ 65,\ 86]$ for nivolumab, and at $S_{0,I} = [23,\ 44,\ 65]$ for ipilimumab. Fig. A.9 shows the results. Extending the dosing window leads to optimized dosing schedules in which ICI doses are increasingly later and doses of the same drug are administered closer together. This effect is more pronounced for nivolumab than for ipilimumab. The dose size of ipilimumab is relatively small for this type of treatment, so the timing of ipilimumab doses may not

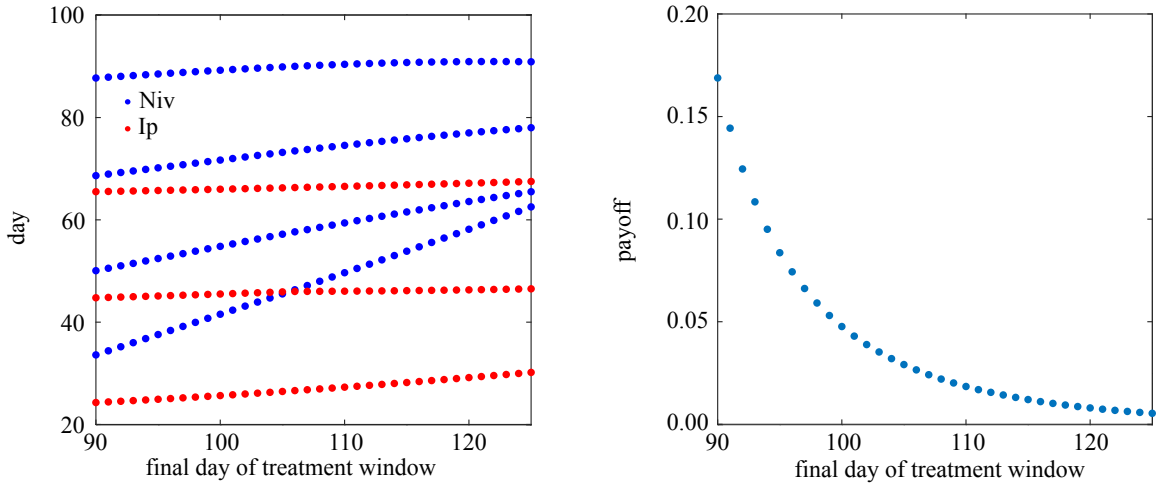


Figure A.9: The optimized dosing schedule (left) and payoff (right) as a function of the length of the treatment window of combination therapy. All optimizations initialized at $S_{0,N} = [23,\ 44,\ 65,\ 86]$ for nivolumab, and at $S_{0,I} = [23,\ 44,\ 65]$ for ipilimumab. Default settings are used for all other parameters.

have a large impact on the outcome of the immune response compared to the impact of nivolumab. The results in Fig. A.9 are similar to what we observed for nivolumab and ipilimumab monotherapy, where increasing the length of treatment window qualitatively changes the optimized dosing schedule but only affects a relatively small change in the associated payoff.

A.3.3 The effect of h

We need to ensure that our results are robust when the gradient descent step parameter, h, is varied. The default value for this parameter is $h = 0.5$. We now vary the value of h while keeping all other parameters their default values for combination therapy, and using the standard–of–care dosing schedule for combination therapy as the initial schedule for all three optimizations. The results in Table A.3 shows that optimization of the dosing schedule for combination therapy is robust to changes in h in terms of achieving our objective of identifying a dosing schedule which does not cause autoimmune myocarditis.

h	Schedule nivolumab	Schedule ipilimumab	Payoff
0.1	29.1, 47.4, 67.0, 86.9	23.7, 44.4, 65.3	0.3789
1	36.8, 52.0, 69.8, 88.3	24.8, 45.0, 65.7	0.1258
10	70.5, 72.7, 80.6, 90.0	32.0, 46.9, 68.0	0.0447

Table A.3: The optimized dosing schedule for combination therapy and the payoff vary as the size of the step parameter, h, is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

A.3.4 The effect of varying γ_3 and γ_4

We explore the effect on the model simulations of combination therapy when the values of γ_3 and γ_4 are varied. The parameter γ_3 translates the mass of nivolumab in the system to the effect on the deactivation rate of pathogenic CD4+ T cells. The parameter γ_4 translates the mass of ipilimumab in the system to the effect on the deactivation rate of regulatory cells. The case where all γ_i, $i = 1, \ldots, 4$, are varied simultaneously is covered by the analysis in Section A.3.1 where the values of V_N and V_I are varied. We explore the effect of varying the relative strength with which a dose of ICIs affects different immunological processes by varying γ_3 and γ_4 and keeping γ_1 and γ_2 constant at one grams^{-1}.

We first consider the effect that varying γ_3 and γ_4 has on the outcome of standard–of–care treatment. When $\gamma_3 = 0.1$ grams^{-1}, the standard–of–care treatment does not cause autoimmune myocarditis to develop, independent of whether γ_4 is equal to 0.1, 1, or 10 grams^{-1}. When γ_3 is equal to one grams^{-1} or 10 grams^{-1}, standard–of–care treatment causes autoimmune myocarditis for γ_4 equal to 0.1, 1 or 10 grams^{-1}. The outcome thus appears primarily sensitive to the value of γ_3.

We optimize combination therapy treatment for $\gamma_3 = 1$ grams^{-1} and $\gamma_3 = 10$ grams^{-1} while varying γ_4 between 0.1, 1 and 10 grams^{-1}. We initialize the dosing schedule at the standard–of–care schedule and keep all other parameters at their default values. When $\gamma_3 = 1$ grams^{-1}, we can identify a myocarditis–free dosing

schedule for all values of γ_4. For $\gamma_4 = 0.1$ grams^{-1} and $\gamma_4 = 10$ grams^{-1}, the evolution of the dosing schedules over the iterations, as well as the optimized dosing schedule, are qualitatively similar to the those found for the default value of $\gamma_4 = 1$ grams^{-1}. When $\gamma_3 = 10$ grams^{-1}, no myocarditis–free dosing schedule can be identified for any values of γ_4.

We thus find that the optimization and outcomes of combination therapy are sensitive to varying the value of γ_3 but not to variations in the value of γ_4. This is consistent with what we found for nivolumab and ipilimumab monotherapy. Thus, it is particularly important that future work includes an exploration of the relative strengths of the effect of nivolumab on different immunological processes.

Appendix B

Exploration of optimal control parameters for heart–tumour model

In Chapter 5 we designed optimized dosing schedules for nivolumab and ipilimumab monotherapy as well as combination therapy using our heart–tumour model. Here we explore the effect of changing the values of the doses V_N and V_I, the length of the treatment window, the step parameter h, and the weights w_1 and w_2 in the payoff function. The default value of V_N is $V_N = 0.25$ grams, the default value of V_I is $V_I = 1.0$ grams for ipilimumab monotherapy and $V_I = 0.1$ grams for combination therapy. The final day of the treatment window, $t_{\max}$, is by default day 60 for nivolumab monotherapy, day 70 for ipilimumab monotherapy and day 90 for combination therapy. The default values of the other parameters are $h = 0.01$, and the weights are $w_1 = 1$ cells^{-1}, $w_2 = 1$ cm^{-3}.

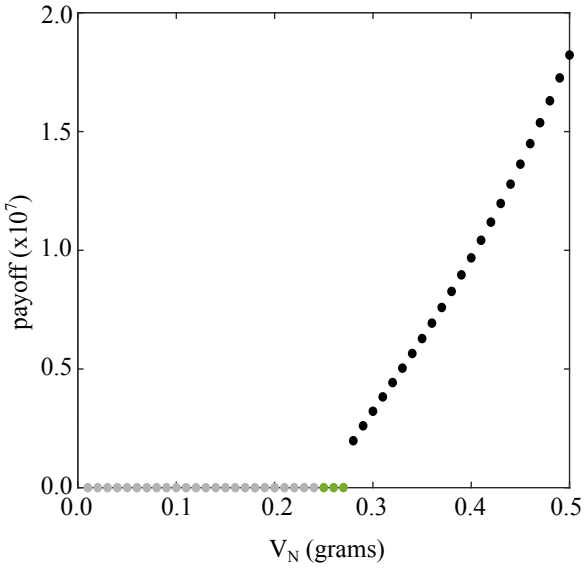


Figure B.1: The payoff of the optimized nivolumab dosing schedule as V_N is varied. A grey dot indicates a value of V_N for which autoimmune myocarditis does not develop when a standard–of–care treatment schedule is applied. A green dot indicates a value of V_N for which autoimmune myocarditis does develop when a standard–of–care treatment schedule is applied, but we can find an optimized treatment schedule for which autoimmune myocarditis is prevented. A black dot indicates a value of V_N for which autoimmune myocarditis does develop when a standard–of–care treatment schedule is applied, and this cannot be prevented by optimizing the dosing schedule.

B.1 Nivolumab monotherapy

B.1.1 The effect of changes to V_N

We explore the effects of changing the value of V_N on the development of autoimmune myocarditis. We vary the value of V_N and keep all other parameter values at their default values listed in Table 5.1. We then use the set–up and algorithm as explained in Chapter 5 to optimize the dosing schedule. We use the standard–of–care schedule as the initial schedule for $V_N = 0.01$ grams, and set the initial schedule for $V_N = m$ grams, $m = 0.02, 0.03, \ldots, 0.27$, to be the optimized schedule for $V_N = m - 0.01$ grams. For $V_N > 0.27$, we set the initial schedule equal $S_0 = [60, 60, 60, 60]$ to speed up convergence.

As shown in Fig. B.1, there is a very small range of values of V_N for which au-

toimmune myocarditis develops when the standard–of–care schedule is applied, but does not develop for an optimized dosing schedule identified by the optimal control algorithm (green). For $V_N \leq 0.24$ grams, the standard–of–care schedule does not cause autoimmune myocarditis. For $V_N \geq 0.28$ grams, the standard–of–care schedule causes autoimmune myocarditis to develop and a myocarditis–free solution cannot be identified. That leaves the range $0.25 \leq V_N \leq 0.27$ grams for which we can optimize the dosing schedule such that it does not cause autoimmune myocarditis where the standard–of–care dosing schedule does, within the given constraints.

The range of V_N for which autoimmune myocarditis occurs as a side–effect of standard–of–care treatment but can be prevented by an optimized schedule is noticeably smaller for the heart–tumour model compared to the heart–only model where autoimmune myocarditis was caused by the standard–of–care schedule for $V_N \geq 0.22$ grams and could be prevented for $V_N \leq 0.355$ grams. The smaller range of V_N for which we can prevent autoimmune myocarditis in the heart–tumour model through optimization of the dosing schedule is likely due to the extra constraints placed on the dosing schedule by the explicit consideration of tumour growth.

B.1.2 The effect of changes to the final dosing day, $t_{\max}$

In this section, we explore the effects of increasing the length of the treatment window on the dosing schedule. We set all other parameters to the values in Table 5.1 and let the final day of the dosing window, $t_{\max}$, vary from day 60 to day 110. We use the standard–of–care schedule as the initial schedule for $t_{\max} = 60$, and thereafter set the initial schedule for $t_{\max} = m$, $m = 61, 62, \ldots, 110$, to be the optimized schedule for $t_{\max} = m - 1$.

Fig. B.2 shows the results. Note that for every value of $t_{\max}$ the last dose in the schedule consists of two doses administered simultaneously. As the length of the treatment window is extended to $t_{\max} = 82$, the last doses of the schedule are delayed

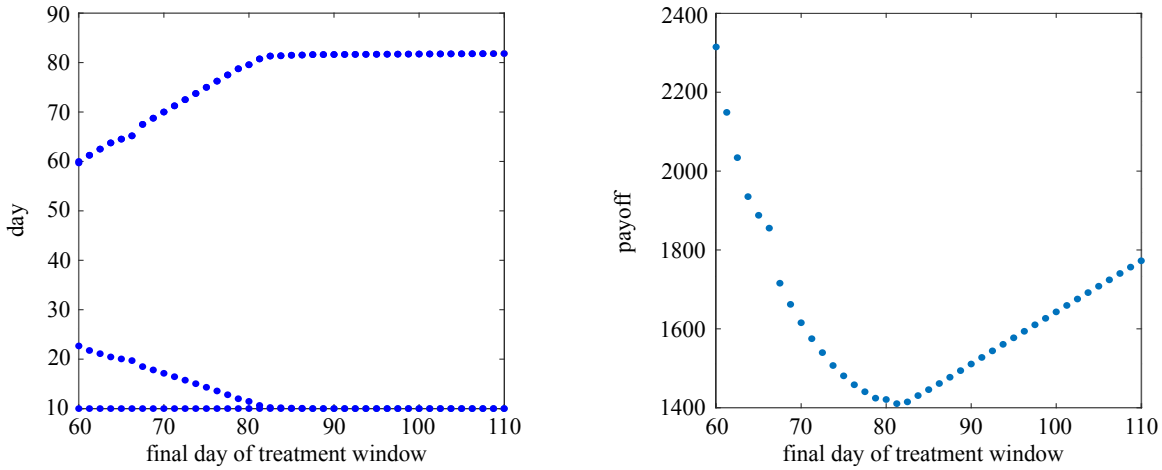


Figure B.2: Optimized nivolumab dosing schedule (left) and payoff (right) for different lengths of the treatment window. Parameter values in Table 5.1 are used unless otherwise noted.

to the last day of the dosing window. For treatment windows where $t_{\max} > 82$, the last two doses are consistently administered on day 82. This can be explained by the fact that dosing later than day 82 would leave too much time between doses during which the tumour can grow uninhibited. The first two doses in the optimized part of the dosing schedule are administered separately early in the treatment window for shorter treatment windows, and simultaneously on day 10 from $t_{\max} = 82$. Overall, we see that the trend of dosing early as well as late in the treatment window such that tumour growth is inhibited by regular dosing and sufficient time is left between doses to prevent autoimmune myocarditis holds as the treatment window is extended.

B.1.3 The effect of changes to h

We now vary the the step size parameter h to ensure robustness of our results with respect to this algorithm parameter. We keep all other parameters at the default values listed in Table 5.1 and initialize with the standard–of–care dosing schedule.

The results in Table B.1 show that the value of h does not qualitatively affect the results of the optimization. For all values of h used here, we see that the optimized dosing schedule for nivolumab monotherapy consists of two early doses administered

h	Schedule nivolumab	Payoff
0.001	10.0, 28.4, 47.9, 60.0	3.9415×10^3
0.005	10.0, 22.5, 60.0, 60.0	2.2870×10^3
0.01	10.0, 22.7, 59.7, 60.0	2.3148×10^3
0.05	10.0, 10.0, 60.0, 60.0	1.8514×10^3
0.1	10.0, 10.0, 60.0, 60.0	1.8514×10^3

Table B.1: The optimized schedule of nivolumab doses and the payoff vary as the size of the step size parameter, h, is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

before day 30 and two late doses administered on day 47 or thereafter. None of the dosing schedules in Table B.1 cause autoimmune myocarditis to develop and all inhibit tumour growth, so although the exact timing of doses varies somewhat, the main objectives of inhibiting tumour growth and preventing autoimmune myocarditis are achieved for all values of h.

B.1.4 The effect of payoff function weights

Lastly, we explore the effects of varying the weights in the payoff function. The payoff function used for all types of ICI therapy in Chapter 5 is

$$J = \int_0^{t_f} \left[w_1 C(t) + w_2 T_u(t) \right] \mathrm{d}t. \tag{B.1}$$

Since only the relative values of the weights matters, it suffices to vary only w_2 and keep $w_1 = 1$ cells^{-1}. We vary w_2 across the range [0.01, 0.1, 1, 10, 100] cm^{-3} and optimize the treatment schedule for each value in this range. We keep all other parameters at the values listed in Table 5.1 and initialize the dosing schedule as the standard–of–care schedule.

Table B.2 shows that for each value of w_2 the optimized dosing schedule is qualitatively the same. The payoff also remains on the same order of magnitude. Because the weights appear in the payoff function, we would expect the payoff to vary some-

w_2	Schedule nivolumab	Payoff
0.001	10.0, 22.6, 59.7, 60.0	2.2770×10^3
0.1	10.0, 22.6, 59.7, 60.0	2.2808×10^3
1	10.0, 22.7, 59.7, 60.0	2.3148×10^3
10	10.0, 23.1, 60.0, 60.0	2.6256×10^3
100	10.0, 26.6, 60.0, 60.0	5.8737×10^3

Table B.2: The optimized schedule of nivolumab doses and the payoff vary as the weight w_2 is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

what as different parts of the payoff function are prioritized and weighted differently. For all values of w_2, the algorithm identifies a dosing schedule which does not cause autoimmune myocarditis to develop and inhibits tumour growth. We can thus conclude that the optimization of nivolumab dosing schedules is robust to changes in the weights in the payoff function.

B.2 Ipilimumab monotherapy

B.2.1 The effect of changes to V_I

We vary the value of V_I to establish the range of values for which autoimmune myocarditis develops as a side–effect of the standard–of–care schedule, but can be prevented by an optimized dosing schedule. We set all other parameters to the values listed in Table 5.1. We use the standard–of–care schedule as the initial schedule for $V_I = 0.75$ grams, and thereafter set the initial schedule for $V_I = m$ grams, $m = 0.76, 0.77, \ldots, 1.25$, to be the optimized schedule for $V_I = m - 0.01$ grams.

Fig. B.3 shows that we again obtain three regions with different outcomes. For $V_I \leq 0.95$ grams, autoimmune myocarditis does not develop when the standard–of–care schedule is applied. For $V_I \geq 1.09$ grams, autoimmune myocarditis is a side–effect when the standard–of–care schedule is applied, and it cannot be prevented

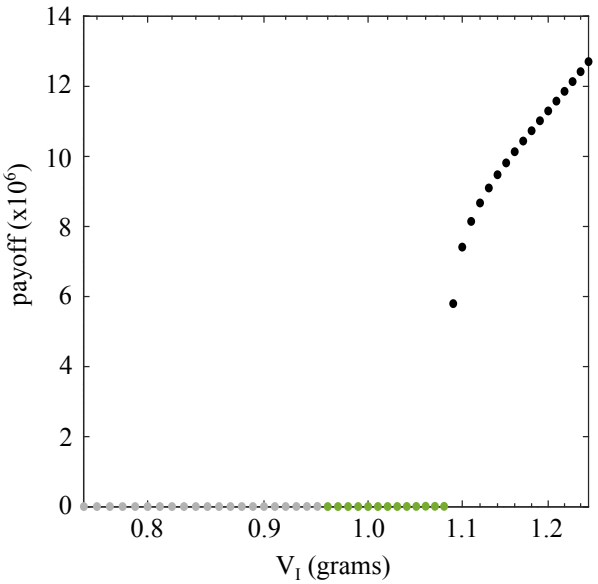


Figure B.3: The payoff of the optimized ipilimumab dosing schedule as V_I is varied. A grey dot indicates a value of V_I for which autoimmune myocarditis does not develop when a standard–of–care treatment schedule is applied. A green dot indicates a value of V_I for which autoimmune myocarditis does develop when a standard–of–care treatment schedule is applied, but we can find an optimized treatment schedule for which autoimmune myocarditis is prevented. A black dot indicates a value of V_I for which autoimmune myocarditis does develop when a standard–of–care treatment schedule is applied, and this cannot be prevented by optimizing the dosing schedule.

by an optimized schedule. This leaves the range $0.96 \leq V_I \leq 1.08$ grams for which autoimmune myocarditis is a side–effect of the standard–of–care treatment schedule, but it can be prevented by an optimized dosing schedule.

The range of values of V_I for which we are able to prevent myocarditis from developing as a side–effect when it did occur with the standard–of–care schedule in the heart–only model is $0.95 < V_I < 1.5$ grams, which is a significantly larger range of values than $0.96 \leq V_I \leq 1.08$ grams which we obtained for the heart–tumour model. Similar to what we saw for nivolumab monotherapy, it is thus clear that the explicit consideration of the effects of ipilimumab on the tumour constrains the range of values of V_I for which we can prevent autoimmune myocarditis from developing.

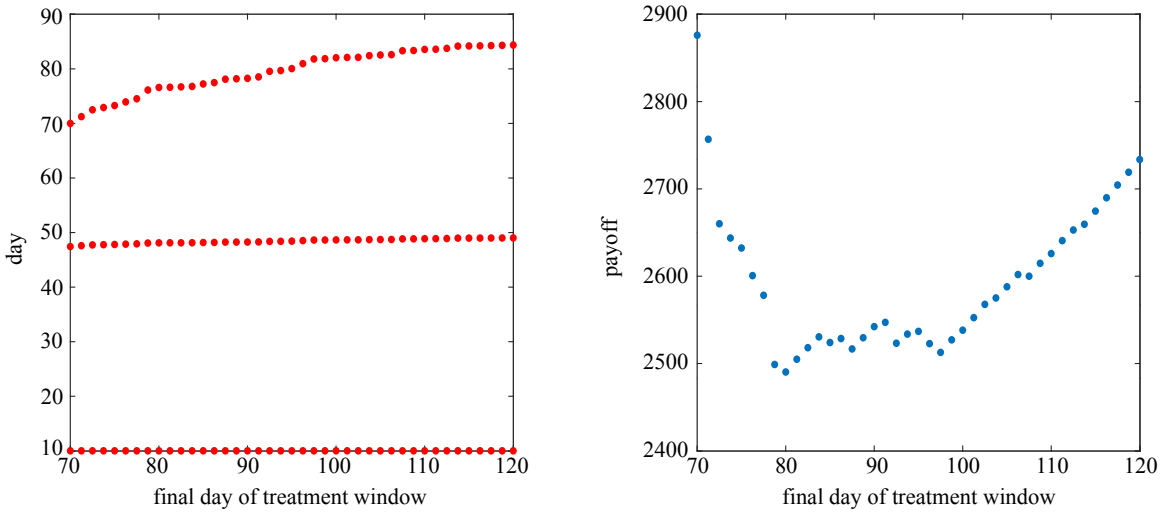


Figure B.4: Optimized ipilimumab dosing schedule (left) and payoff (right) for different lengths of the treatment window. Parameter values listed in Table 5.1 are used unless otherwise noted.

B.2.2 The effect of changes to the final dosing day

We now turn to the effect of the length of the treatment window on the optimization of treatment schedules for ipilimumab monotherapy. The last day of the treatment window, $t_{\max}$, is day 70 by default. We now let $t_{\max}$ increase from 70 to 120, while keeping all parameter values as listed in Table 5.1. We use the standard–of–care schedule as the initial schedule for $t_{\max} = 70$, and thereafter set the initial schedule for $t_{\max} = m$, $m = 71, 72, \ldots, 120$, to be the optimized schedule for $t_{\max} = m - 1$.

Fig. B.4 shows that the results of the optimization are qualitatively robust to changes in the length of the dosing window. The second and thirds dose, counting the dose administered on day 2 as the first dose, are administered approximately on the same day, regardless of the length of the treatment window. The fourth dose ipilimumab is administered later as the treatment window increases. This even spread of the second, third and fourth dose over the treatment window controls tumour growth while leaving enough time between doses for immune cell counts to decrease, thus preventing autoimmune myocarditis.

h	Schedule ipilimumab	Payoff
0.001	15.1, 46.6, 70.0	3.8528×10^3
0.005	13.7, 46.7, 70.0	3.4658×10^3
0.01	10.0, 47.4, 70.0	2.8761×10^3
0.05	10.0, 47.4, 70.0	2.8759×10^3
0.1	10.0, 47.5, 70.0	2.8749×10^3

Table B.3: The schedule of ipilimumab doses and the payoff vary as the size of the step size parameter, h, is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

B.2.3 The effect of h

We explore the robustness of the optimization results for ipilimumab monotherapy with respect to changes in the step parameter h. We keep all other parameters at the values listed in Table 5.1 while varying h from 0.001 to 0.1. We initialize the dosing schedule at the standard–of–care schedule.

The results in Table B.3 show that the value of h does not qualitatively affect the results of the optimization. For all values of h used here, we see that the optimized dosing schedule for ipilimumab monotherapy consists of one early dose, one dose around day 47 and one late dose on day 70, thus covering the whole treatment window. None of the dosing schedules in Table B.3 cause autoimmune myocarditis to develop and all inhibit tumour growth.

B.2.4 The effect of payoff function weights

Lastly, we explore the effects of varying the weights in the payoff function on the optimization of ipilimumab monotherapy. The payoff function used for all types of therapy in Chapter 5 is

$$J = \int_0^{t_f} \left[w_1 C(t) + w_2 T_u(t) \right] \mathrm{d}t. \tag{B.2}$$

w_2	Schedule ipilimumab	Payoff
0.01	10.0, 47.4, 70.0	0.2353×10^4
0.1	10.0, 47.4, 70.0	0.2400×10^4
1	10.0, 47.4, 70.0	0.2876×10^4
10	10.0, 47.5, 70.0	0.7590×10^4
100	16.2, 46.6, 70.0	5.5136×10^4

Table B.4: The schedule of ipilimumab doses and the payoff vary as the weight w_2 is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

Since only the relative values of the weights matters, it suffices to vary only w_2, while keeping $w_1 = 1$ cells^{-1}. We vary w_2 across the range [0.01, 0.1, 1, 10, 100] cm^{-3} and optimize the treatment schedule for each value in this range. We keep all other parameters at the values listed in Table 5.1 and initialize the dosing schedule as the standard–of–care schedule.

Table B.4 shows that for each value of w_2, the optimized dosing schedules remain qualitatively similar. Because the weights appear in the payoff function, we would expect the payoff to vary somewhat as different parts of the payoff function are prioritized and weighted differently, and this reflect in Table B.4. For all values of w_2, the algorithm identifies a dosing schedule which does not cause autoimmune myocarditis to develop and inhibits tumour growth. We can thus conclude that the optimization of ipilimumab dosing schedules is robust to changes in the weights in the payoff function.

B.3 Combination therapy

B.3.1 The effect of changes to V_N and V_I

We explore the effects of varying the values of V_N and V_I simultaneously over the ranges 0.1–0.3 grams and 0.0–0.3 grams, respectively. We keep all other parameters

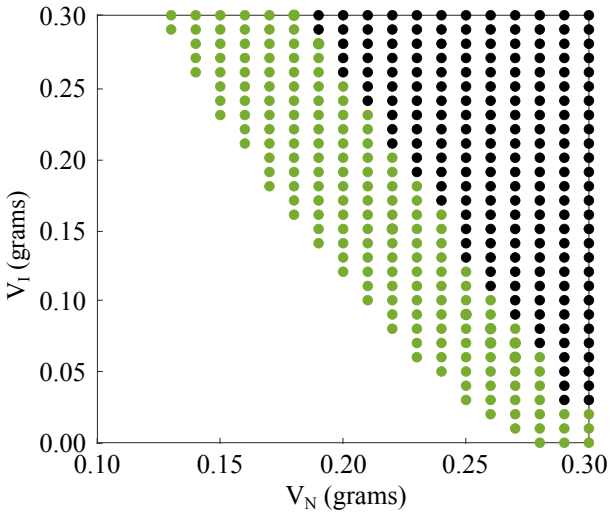


Figure B.5: Treatment outcome of the optimized combination therapy dosing schedule as V_N and V_I are varied simultaneously. In the white region of values of V_N and V_I autoimmune myocarditis never develops. In the green region of values of V_N and V_I autoimmune myocarditis is a possible side–effect of combination therapy but it can be prevented by an optimized dosing schedule. In the black region of values of V_N and V_I autoimmune myocarditis cannot be prevented as a side–effect.

at the values in Table 5.1. To speed up the optimization, for each value of V_N we initialize the dosing schedule at the standard–of–care schedule if $V_I = 0$ grams, and thereafter set the initial schedule for $V_I = m$ grams, $m = 0.01, 0.02, \ldots, 0.30$, to be the optimized schedule for $V_I = m - 0.01$ grams.

We find again that there are three types of outcome as we vary V_N and V_I. The white area in Fig. B.5 indicates the values for which autoimmune myocarditis does not develop as a side–effect of the standard–of–care treatment schedule. The green area indicates the values for which myocarditis does develop as side–effect of the standard–of–care treatment schedule, but we can optimize the treatment schedule such that we prevent this. The black area indicates those values for which autoimmune myocarditis cannot be prevented within the constraints set here. Again, we observe that the ranges of V_N and V_I for which we can prevent autoimmune myocarditis from developing are much smaller for the heart–tumour model compared to the heart–only model (see Fig. A.8), emphasizing the stricter constraints placed on the optimization by the consideration of the effects of ICIs on tumour growth.

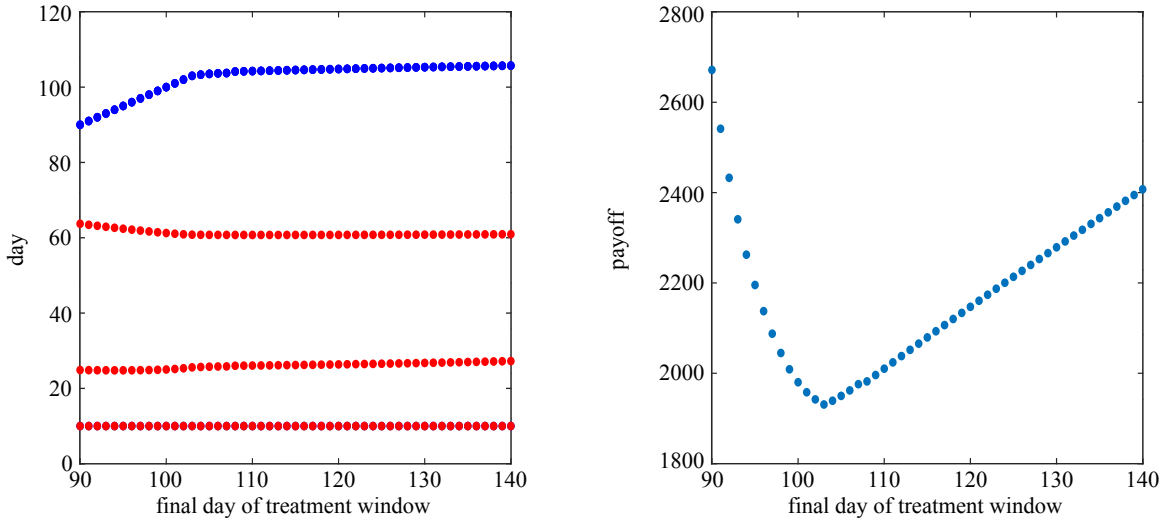


Figure B.6: Optimized combination therapy dosing schedule (left) and payoff (right) for different lengths of the treatment window. Parameter values listed in Table 5.1 are used unless otherwise noted. For the optimization of the dosing schedule for each length of treatment window, the initial dosing schedule, S_0, is set to be the standard–of–care dosing schedule.

B.3.2 The effect of extending the treatment window

We now consider the effect that extending the treatment window has on the optimization of combination therapy treatment schedules. All parameters are set to the values listed in Table 5.1. The final day of the treatment window, $t_{\max}$, is varied from 90 to 140. To speed up the optimization, we use the standard–of–care schedule as the initial schedule for $t_{\max} = 90$, and thereafter set the initial schedule for $t_{\max} = m$, $m = 91, 92, \ldots, 140$, to be the optimized schedule for $t_{\max} = m - 1$.

Fig. B.6 shows the results. Note that the last dose of nivolumab (blue) consists of three doses administered simultaneously, and the second dose, administered on day 10 combines one dose of nivolumab and one dose of ipilimumab. As Fig. B.6 shows, the optimization of the treatment schedule for combination therapy is robust to changes in the length of the treatment window. As the treatment window is extended, the optimized schedule does not change qualitatively. The final three doses of nivolumab, which are administered simultaneously, are delayed slightly but the

h	Schedule nivolumab	Schedule ipilimumab	Payoff
0.001	10.0, 77.2, 90.0, 90.0	10.0, 25.9, 69.1	5647.9
0.005	10.0, 89.8, 90.0, 90.0	10.0, 24.9, 63.8	2695.7
0.01	10.0, 90.0, 90.0, 90.0	10.0, 24.9, 63.7	2671.5
0.05	10.0, 90.0, 90.0, 90.0	10.0, 24.2, 58.3	2558.6
0.1	10.0, 90.0, 90.0, 90.0	10.0, 23.4, 52.9	2480.4

Table B.5: The schedule of combination therapy and the payoff vary as the size of the step size parameter, h, is varied. None of the dosing schedules in this table cause autoimmune myocarditis to develop.

timing of the other doses remains mostly constant as $t_{\max}$ is varied. The payoff of each of these schedules also does not change significantly taking into account that a longer simulation is expected to lead to a higher payoff as the integral is computed over a longer time span.

B.3.3 The effect of h

Next we investigate robustness of the optimization results to varying the step size parameter h. To speed up the optimization of combination therapy dosing schedules, we set $h = 0.1$ initially and then change to $h = 0.01$ once the payoff is less than 10^4. Here we vary the value of h once the payoff is less than 10^4 and leave $h = 0.1$ above this threshold. All other parameters are set to the values listed in Table 5.1 and we initialize the dosing schedule as the standard–of–care schedule for all h.

None of the schedules in Table B.5 cause autoimmune myocarditis to develop. Furthermore, the schedules do not change qualitatively as the value of h is varied. The results of the optimization of combination therapy are thus robust to our choice of h.

B.3.4 The effect of payoff function weights

Lastly, we explore the effect of changing the weights in the payoff function on the optimization results. The payoff function used for all types of therapy in Chapter 5 is

$$J = \int_0^{t_f} [w_1 C(t) + w_2 T_u(t)] \, \mathrm{d}t. \tag{B.3}$$

Since only the relative values of the weights matter, it suffices to vary only w_2 and keep $w_1 = 1$ cells^{-1}. We vary each w_2 across the range [0.01, 0.1, 1, 10, 100] cm^{-3} and optimize the treatment schedule for each value of w_2. We keep all other parameters at the values listed in Table 5.1 and initialize the dosing schedule as the standard–of–care schedule.

Results in Table B.6 show that we need to take some care when picking the weights in the payoff function for combination therapy. For $w_2 = 0.01$ the optimized schedule still causes autoimmune myocarditis to develop. The evolution of the dosing schedule and the payoff over the iterations of the gradient descent algorithm with this value of w_2 show that, up to convergence, the evolution of the dosing schedule is qualitatively the same as for those where the algorithm does find a myocarditis–free solution. The problem thus appears to be that convergence happens too far away from a local minimum in the payoff landscape, likely due to the flatness of the payoff landscape in a large neighborhood around the minimum. This indicates that careful calibration of the step parameter h, the payoff functions weights and the convergence threshold is necessary to ensure convergence to a myocarditis–free solution.

In this appendix we have shown that the ranges of V_N and V_I for which we can prevent autoimmune myocarditis when it occurs as a side–effect of the standard–of–care schedule are more narrow when the effects of ICIs on tumour growth are considered in addition to the prevention of autoimmune myocarditis. We have also

w_2	Schedule nivolumab	Schedule ipilimumab	Payoff
0.01	10.0, 68.8, 90.0, 90.0	14.5, 32.2, 71.6	1.8297×10^5
0.1	10.0, 90.0, 90.0, 90.0	10.0, 24.9, 63.7	2.6407×10^3
1	10.0, 90.0, 90.0, 90.0	10.0, 24.9, 63.7	2.6716×10^3
10	10.0, 89.8, 90.0, 90.0	10.0, 25.0, 63.7	3.0039×10^3
100	10.0, 80.4, 90.0, 90.0	10.0, 25.7, 67.2	7.4791×10^3

Table B.6: The schedule of combination therapy and the payoff vary as the weight w_2 is varied. The optimized schedule for $w_2 = 0.01$ is the only schedule which causes autoimmune myocarditis to develop.

shown that for nivolumab and ipilimumab monotherapy, the results of optimization of the dosing schedule are robust to changes in the length of the treatment window, the step parameter h, and the weights in the payoff function. For combination therapy, the optimized schedule is not significantly affected by changes to length of the treatment window or the step parameter h if all other parameters are set to their default values, but care should be taken when setting the values of h in combination with certain payoff function weights. The weights of the payoff function are ideally set by the biological context of the optimization problem and we have shown that for certain values of the weights the algorithm converges to within the tolerance threshold before a myocarditis–free solution is found, in a scenario where one does exist. This suggests that the default value of h that we have used in Chapter 5 and this appendix is not appropriate for all values of the weights and should thus be carefully considered.

Bibliography

[1] A. Abbas, J. Lohr, et al. T cell tolerance and autoimmunity. *Autoimmunity Reviews*, 3(7):471–475, 2004.

[2] A. Abdel-Gadir, A. Massoud, et al. Antigen–specific Treg cells in immunological tolerance: implications for allergic diseases. *F1000Research*, 7(38), 2018.

[3] Y. Adachi, A. Tamiya, et al. Predictive factors for progression–free survival in non–small cell lung cancer patients receiving nivolumab based on performance status. *Cancer Medicine*, 9(4):1383–1391, 2020.

[4] M. Afanasyeva, D. Georgakopoulos, et al. Autoimmune myocarditis: Cellular mediators of cardiac dysfunction. *Autoimmunity Reviews*, 3:476–486, 2004.

[5] N. Agrawal, A. Khunger, et al. Cardiac toxicity associated with immune checkpoint inhibitors: Case series and review of the literature. *Case Reports in Oncology*, 12(1):260–276, 2019.

[6] B. Alberts, A. Johnson, et al. *Molecular Biology of the Cell*. Garland Science, New York, NY, USA, 4 edition, 2002.

[7] B. P. Amoah, H. Yang, et al. Immunopathogenesis of myocarditis: The interplay between cardiac fibroblast cells, dendritic cells, macrophages and CD4+T cells. *Scandinavian Journal of Immunology*, 82(1):1–9, 2015.

[8] P. Anagnostopoulou, D. Hasler, et al. A mathematical model to understand the airway clearance in cystic fibrosis. *European Respiratory Journal*, 54(suppl 63), 2019.

[9] C. Atri, F. Guerfali, et al. Role of human macrophage polarization in inflammation during infectious disease. *International Journal of Molecular Sciences*, 19(6), 2018.

[10] J. Barin and D. Cihakova. Control of inflammatory disease by CD4+ T cells. *Annals of the New York Academy of Science*, 1285(1):80–96, 2013.

[11] J. Barin, N. Rose, et al. Macrophage diversity in cardiac inflammation: A review. *Immunobiology*, 217(5):468–475, 2012.

[12] J. Barin, M. Talor, et al. Mechanisms of IFNγ regulation of autoimmune myocarditis. *Experimental and Molecular Pathology*, 89(2):83–91, 2010.

[13] G. Bellu, M. P. Saccomani, et al. DAISY: a new software tool to test global identifiability of biological and physiological systems. *Computer methods and programs in biomedicine*, 88(1):52–61, 2007.

[14] M. Bertsch, B. Franchi, et al. Alzheimer's disease: a mathematical model for onset and progression. *Mathematical Medicine and Biology*, 34(2):193–214, 2017.

[15] W. Bracamonte-Baran and D. Cihakova. *Advances in Experimental Medicine and Biology*, volume 1003, chapter Cardiac autoimmunity: Myocarditis, pages 187–221. Springer New York LLC, 2017.

[16] R. Brady and H. Enderling. Mathematical models of cancer: When to predict novel therapies, and when not to. *Bulletin of Mathematical Biology*, 81:3722–3731, 2019.

[17] E. Buchbinder and A. Desai. CTLA–4 and PD–1 pathways. *American Journal of Clinical Oncology*, 39(1):98–106, 2016.

[18] M. Burnet. Cancer–a biological approach. *British Medical Journal*, 1(5023):841–847, 1957.

[19] M. Burnet. Immunological surveillance in neoplasia. *Transplantation reviews*, 7:3–25, 1971.

[20] P. Błyszczuk. Myocarditis in humans and in experimental animal models. *Frontiers in Cardiovascular Medicine*, 6:64, 2019.

[21] F. Carbonell, Y. Iturria-Medina, et al. Mathematical modeling of protein misfolding mechanisms in neurological diseases: A historical overview. *Frontiers in Neurology*, 9:37, 2018.

[22] J. Carneiro, T. Paixão, et al. Immunological self–tolerance: Lessons from mathematical modeling. *Journal of Computational and Applied Mathematics*, 184(1):77–100, 2005.

[23] F. Castiglione and B. Piccoli. Optimal control in a model of dendritic cell transfection cancer immunotherapy. *Bulletin of Mathematical Biology*, 68(2):255–274, 2006.

[24] F. Castiglione and B. Piccoli. Cancer immunotherapy, mathematical modeling and optimal control. *Journal of Theoretical Biology*, 247(4):723–732, 2007.

[25] M. Castro and R. J. de Boer. Testing structural identifiability by a simple scaling method. *PLoS Computational Biology*, 16(11), 2020.

[26] S. Celli, M. Day, et al. How many dendritic cells are required to initiate a T–cell response? *Blood*, 120(19):3945–3948, 2012.

[27] M. Centanni, D. J. A. R. Moes, et al. Clinical pharmacokinetics and pharmacodynamics of immune checkpoint inhibitors. *Clinical Pharmacokinetics*, 58(7):835–857, 2019.

[28] X. Chen, T. Hickling, et al. A mechanistic, multiscale mathematical model of immunogenicity for therapeutic proteins: Part 1 – theoretical model. *CPT: Pharmacometrics and Systems Pharmacology*, 3(e133), 2014.

[29] F. Cheng and J. Loscalzo. Autoimmune cardiotoxicity of cancer immunotherapy. *Trends in Immunology*, 38(2):77–78, 2017.

[30] J. C. Chimal-Eguia, J. C. Rangel-Reyes, et al. Improving convergence in therapy scheduling optimization: A simulation study. *Mathematics*, 8(12):1–17, 2020.

[31] L. Chow, Y. Ye, et al. Phenotypic analysis of infiltrating cells in human myocarditis: An immunohistochemical study in paraffin–embedded tissue. *Archives of pathology & laboratory medicine*, 113(12):1357–1362, 1989.

[32] A. Corthay. Does the immune system naturally protect against cancer? *Frontiers in Immunology*, 5, 2014.

[33] K. Couper, D. Blount, et al. IL–10: The master regulator of immunity to infection. *The Journal of Immunology*, 180(9):5771–5777, 2008.

[34] E. M. DeFilippis, S. Narain, et al. Rapidly progressive cardiac failure due to giant cell myocarditis. *HSS Journal*, 11(2):182–186, 2015.

[35] M. Delitala, U. Dianzani, et al. A mathematical model for immune and autoimmune response mediated by T–cells. *Computers & Mathematics with Applications*, 66(6):1010–1023, 2013.

[36] S. Dick, J. Macklin, et al. Self–renewing resident cardiac macrophages limit

adverse remodeling following myocardial infarction. *Nature Immunology*, 20:29–39, 2019.

[37] N. Diny, N. Rose, et al. Eosinophils in autoimmune diseases. *Frontiers in Immunology*, 8:484, 2017.

[38] H. Dritschel, S. Waters, et al. A mathematical model of cytotoxic and helper T cell interactions in a tumour microenvironment. *Letters in Biomathematics*, 5(sup1):S34–S68, 2018.

[39] M. Duchosal. B–cell development and differentiation. *Seminars in Hematology*, 34(1 Suppl 1):2–12, 1997.

[40] G. P. Dunn, A. T. Bruce, et al. Cancer immunoediting: from immuno–surveillance to tumor escape. *Nature Immunology*, 3(11):991–998, 2002.

[41] G. P. Dunn, L. J. Old, et al. The three Es of cancer immunoediting. *Annual Review of Immunology*, 22:329–360, 2004.

[42] S. Epelman, K. Lavine, et al. Embryonic and adult–derived resident cardiac macrophages are maintained through distinct mechanisms at steady state and during inflammation. *Immunity*, 40(1):91–104, 2014.

[43] European Medicines Agency. OPDIVO Annex I – Summary of product characteristics.

[44] European Medicines Agency. YERVOY Annex I – Summary of product characteristics.

[45] N. Fachada, V. Lopes, et al. Agent based modelling and simulation of the immune system: a review. In *13º Encontro Português de Inteligência Artificial (EPIA 2007)*, pages 300–315, 2007.

[46] F. Fatehi, S. N. Kyrychko, et al. Stochastic effects in autoimmune dynamics. *Frontiers in Physiology*, 9:45, 2018.

[47] D. A. Fennell, S. Ewings, et al. Nivolumab versus placebo in patients with relapsed malignant mesothelioma (confirm): a multicentre, double–blind, randomised, phase 3 trial. *The Lancet Oncology*, 22(11):1530–1540, 2021.

[48] D. F. Fiorentino, A. Zlotnik, et al. IL–10 acts on the antigen–presenting cell to inhibit cytokine production by Th1 cells. *The Journal of Immunology*, 146(10):3444–3451, 1991.

[49] S. Fulda, A. M. Gorman, et al. Cellular stress responses: Cell survival and cell death. *International Journal of Cell Biology*, 2010, 2010.

[50] G. Fung, H. Luo, et al. Myocarditis. *Circulation Research*, 118(3):496–514, 2016.

[51] S. J. Gaudino and P. Kumar. Cross–talk between antigen presenting cells and T cells impacts intestinal homeostasis, bacterial infections, and tumorigenesis. *Frontiers in Immunology*, 10:360, 2019.

[52] C. Gil-Cruz, C. Perez-Shibayama, et al. Microbiota–derived peptide mimics drive lethal inflammatory cardiomyopathy. *Science*, 366(6467):881–886, 2019.

[53] M. Gonzalez-Meija and A. Doseff. Regulation of monocytes and macrophages cell fate. *Frontiers in Bioscience*, 14:2413–2431, 2009.

[54] C. Granier, E. D. Guillebon, et al. Mechanisms of action and rationale for the use of checkpoint inhibitors in cancer. *ESMO Open*, 2(2), 2017.

[55] J. L. Guerriero. Chapter three – Macrophages: Their untold story in T cell activation and function. In L. Galluzzi and N.-P. Rudqvist, editors, *Biology of*

T Cells – Part B, volume 342 of *International Review of Cell and Molecular Biology*, pages 73–93. Academic Press, 2019.

[56] H. R. Hampton and T. Chtanova. Lymphatic migration of immune cells. *Frontiers in Immunology*, 10:1168, 2019.

[57] W. Hao and A. Friedman. Mathematical model on Alzheimer's disease. *BMC Systems Biology*, 10(108), 2016.

[58] A. Haslam and V. Prasad. Estimation of the percentage of US patients with cancer who are eligible for and respond to checkpoint inhibitor immunotherapy drugs. *JAMA Network Open*, 2(5):e192535–e192535, 2019.

[59] L. Heinzerling, E. de Toni, et al. Checkpoint Inhibitors. *Deutsch Arzteblatt International*, 116(8):119–126, 2019.

[60] L. Heinzerling, P. Ott, et al. Cardiotoxicity associated with CTLA4 and PD1 blocking immunotherapy. *Journal for ImmunoTherapy of Cancer*, 4(50), 2016.

[61] S. Heymans, U. Eriksson, et al. The quest for new approaches in myocarditis and inflammatory cardiomyopathy. *Journal of the American College of Cardiology*, 68:2348–2364, 2016.

[62] X. Hou, G. Chen, et al. The cardiac microenvironment instructs divergent monocyte fates and functions in myocarditis. *Cell Reports*, 28(1):172–189.e7, 2019.

[63] W.-Z. Huang, W.-H. Hu, et al. A mathematical modelling of initiation of dendritic cells–induced T cell immune response. *International Journal of Biological Sciences*, 15:1396–1403, 2019.

[64] S. Hugues. Dynamics of dendritic cell–T cell interactions: a role in T cell outcome. *Seminars in Immunopathology*, 32(3):227–238, 2010.

[65] P. Italiani and D. Boraschi. From monocytes to M1/M2 macrophages: Phenotypical vs. functional differentiation. *Frontiers in Immunology*, 5:514, 2014.

[66] D. Jain, R. Russell, et al. Cardiac complications of cancer therapy: Pathophysiology, identification, prevention, treatment, and future directions. *Current Cardiology Report*, 19(36), 2017.

[67] C. Janeway, P. Travers, et al. *Immunobiology: The Immune System in Health and Disease*. Garland Science, New York, NY, USA, 5th edition, 2001.

[68] D. Johnson, J. Balko, et al. Fulminant myocarditis with combination immune checkpoint blockade. *New England Journal of Medicine*, 375(18):1749–1755, 2016.

[69] S. Josefowicz, L.-F. Lu, et al. Regulatory T cells: Mechanisms of differentiation and function. *Annual Review of Immunology*, 30(1):531–564, 2012.

[70] R. Kandolin, J. Lehtonen, et al. Diagnosis, treatment, and outcome of giant–cell myocarditis in the era of combined immunosuppression. *Circulation: Heart Failure*, 6(1):15–22, 2013.

[71] M. Kaplan. Role of neutrophils in systemic autoimmune diseases. *Arthritis Research and Therapy*, 15(219), 2013.

[72] M. K. Kim and J. Kim. Properties of immature and mature dendritic cells: phenotype, morphology, phagocytosis, and migration. *RSC Advances*, 9(20):11230–11238, 2019.

[73] D. E. Kirk. *Optimal control theory: An introduction*. Prentice–Hall, 1970.

[74] S. Klein, J. Ziello, et al. PD–1 blockade activates conventional CD4 T cells and the innate immune response during glioblastoma eradication. *The Journal of Immunology*, 200(1 Supplement), 2018.

[75] E. Kolaczkowska and P. Kubes. Neutrophil recruitment and function in health and inflammation. *Nature Reviews Immunology*, 13(3):159–175, 2013.

[76] B. Kumar, T. Connors, et al. Human T cell development, localization, and function throughout life. *Immunity*, 48(2):202–213, 2018.

[77] W. Kwok, V. Tan, et al. Frequency of epitope–specific naive CD4+ T cells correlates with immunodominance in the human memory repertoire. *The Journal of Immunology*, 188(6):2537–2544, 2012.

[78] C. Lappas, J. Rieger, et al. A2A adenosine receptor induction inhibits IFN–γ production in murine CD4+ T cells. *The Journal of Immunology*, 174(2):1073–1080, 2005.

[79] J. Larkin, V. Chiarion-Sileni, et al. Five–year survival with combined nivolumab and ipilimumab in advanced melanoma. *New England Journal of Medicine*, 381(16):1535–1546, 2019.

[80] A. Lichtman. The heart of the matter: Protection of the myocardium from T cells. *Journal of Autoimmunity*, 45:90–96, 2013.

[81] Z. Liu and C. Yang. A mathematical model of cancer treatment by radiotherapy. *Computational and Mathematical Methods in Medicine*, 2014, 2014.

[82] A. G. López, J. Sabuco, et al. Avoiding healthy cells extinction in a cancer model. *Journal of Theoretical Biology*, 349:74–81, 2014.

[83] D. Lu, Z. Ni, et al. Beyond T cells: Understanding the role of PD–1/PD–L1 in tumor–associated macrophages. *Journal of Immunology Research*, 2019, 2019.

[84] R. Luckheeram, R. Zhou, et al. CD4+ T cells: differentiation and functions. *Clinical and Developmental Immunology*, 2012(925135), 2012.

[85] H. Lv, E. Havari, et al. Impaired thymic tolerance to α–myosin directs autoimmunity to the heart in mice and humans. *Journal of Clinical Investigation*, 121(4):1561–1573, 2011.

[86] G. E. Mahlbacher, K. C. Reihmer, et al. Mathematical modeling of tumor–immune cell interactions. *Journal of Theoretical Biology*, 469:47–60, 2019.

[87] S. Marino, I. Hogue, et al. A methodology for performing global uncertainty and sensitivity analysis in systems biology. *Journal of Theoretical Biology*, 254:178–196, 2008.

[88] R. B. Martin. Optimal control drug scheduling of cancer chemotherapy. *Automatica*, 28(6):1113–1123, 1992.

[89] I. V. Martínez, M. E. Hernando, et al. Definition of an agent–based model of the autoimmune response in Type 1 diabetes. In *7th Iberian Conference on Information Systems and Technologies (CISTI 2012)*, pages 1–4, 2012.

[90] J. Mason. Myocarditis and dilated cardiomyopathy: An inflammatory link. *Cardiovascular Research*, 60(1):5–10, 2003.

[91] T. Mayadas, X. Cullere, et al. The multifaceted functions of neutrophils. *Annual Review of Pathology: Mechanisms of Disease*, 9(1):181–218, 2014.

[92] J. Mbongue, H. Nieves, et al. The role of dendritic cell maturation in the induction of insulin–dependent diabetes mellitus. *Frontiers in Immunology*, 8:327, 2017.

[93] X. Meng, J. Yang, et al. Regulatory T cells in cardiovascular diseases. *Nature Reviews Cardiology*, 13(3):167–179, 2016.

[94] J. Mestas and C. Hughes. Of mice and not men: Differences between mouse and human immunology. *The Journal of Immunology*, 172(5):2731–2738, 2004.

[95] L. Michel, T. Rassaf, et al. Cardiotoxicity from immune checkpoint inhibitors. *IJC Heart & Vasculature*, 25:100420, 2019.

[96] P. Momtaz, V. Park, et al. Safety of infusing ipilimumab over 30 minutes. *Journal of Clinical Oncology*, 33(30):3454–3458, 2015.

[97] M. Moro-García, J. Mayo, et al. Influence of inflammation in the process of T lymphocyte differentiation: Proliferative, metabolic, and oxidative changes. *Frontiers in Immunology*, 9:339, 2018.

[98] H. Murphy, H. Jaafari, et al. Differences in predictions of ODE models of tumor growth: A cautionary example. *BMC Cancer*, 16(1), 2016.

[99] B. Nelson. IL–2, regulatory T cells, and tolerance. *The Journal of Immunology*, 172(7):3983–3988, 2004.

[100] R. Nurieva, X. Liu, et al. Molecular mechanisms of T-cell tolerance. *Immunological Reviews*, 241(1):133–144, 2011.

[101] T. Németh and A. Mócsai. The role of neutrophils in autoimmune diseases. *Immunology Letters*, 143(1):9–19, 2012.

[102] R. Obst. The timing of T cell priming and cycling. *Frontiers in Immunology*, 6:563, 2015.

[103] T. Okazaki, Y. Tanaka, et al. Autoantibodies against cardiac troponin I are responsible for dilated cardiomyopathy in PD–1–deficient mice. *Nature Medicine*, 9:1477–1483, 2003.

[104] E. B. Okeke and J. E. Uzonna. The pivotal role of regulatory T cells in the regulation of innate immune cells. *Frontiers in Immunology*, 10:680, 2019.

[105] S. Ong, D. Ligons, et al. Natural killer cells limit cardiac inflammation and fibrosis by halting eosinophil infiltration. *American Journal of Pathology*, 185(3):847–861, 2015.

[106] S. Ong, N. Rose, et al. Natural killer cells in inflammatory heart disease. *Clinical Immunology*, 175:26–33, 2017.

[107] V. Ovcinnikovs, E. M. Ross, et al. CTLA–4–mediated transendocytosis of costimulatory molecules primarily targets migratory dendritic cells. *Science Immunology*, 4(35), 2019.

[108] N. Palaskas, J. Mattei, et al. Immune checkpoint inhibitor myocarditis: Pathophysiological characteristics, diagnosis, and treatment. *Journal of the American Heart Association*, 9(2):e013757, 2020.

[109] J. Pan, M. Zhang, et al. Interferon–γ is an autocrine mediator for dendritic cell maturation. *Immunology Letters*, 94(1):141–151, 2004.

[110] V. Papayannopoulos. Neutrophil extracellular traps in immunity and disease. *Nature Reviews Immunology*, 108:134–147, 2018.

[111] Y. Park and B. Bochner. Eosinophil survival and apoptosis in health and disease. *Allergy, Asthma and Immunology Research*, 2(2):87–101, 2010.

[112] Q. Peng, X. Qiu, et al. PD–L1 on dendritic cells attenuates T cell activation and regulates response to immune checkpoint blockade. *Nature Communications*, 11(1), 2020.

[113] N. Pennock, J. White, et al. T cell responses: naïve to memory and everything in between. *Advances in Physiology Education*, 37(4):273–283, 2013.

[114] G. Peretto, S. Sala, et al. Ventricular arrhythmias in myocarditis. *Journal of the American College of Cardiology*, 75(9):1046–1057, 2020.

[115] R. Peterson. Regulatory T–cells: Diverse phenotypes integral to immune homeostasis and suppression. *Toxicologic Pathology*, 40(2):186–204, 2012.

[116] B. Piccoli and F. Castiglione. Optimal vaccine scheduling in cancer immunotherapy. *Physica A: Statistical Mechanics and its Applications*, 370(2):672–680, 2006.

[117] M. A. Postow, R. Sidlow, et al. Immune–related adverse events associated with immune checkpoint blockade. *New England Journal of Medicine*, 378(2):158–168, 2018.

[118] S. Psarras, D. Beis, et al. Three in a box: Understanding cardiomyocyte, fibroblast, and innate immune cell interactions to orchestrate cardiac repair processes. *Frontiers in Cardiovascular Medicine*, 6:32, 2019.

[119] A. A. de los Reyes V, E. Jung, et al. Optimal control strategies of eradicating invisible glioblastoma cells after conventional surgery. *Journal of the Royal Society Interface*, 12(106), 2015.

[120] A. A. de los Reyes V and Y. Kim. Optimal regulation of tumour–associated neutrophils in cancer progression. *Royal Society Open Science*, 9(2), 2022.

[121] M. Robertson-Tessi, A. El-Kareh, et al. A mathematical model of tumor–immune interactions. *Journal of Theoretical Biology*, 294:56–73, 2012.

[122] K. Rock and H. Kono. The inflammatory response to cell death. *Annual Review of Pathology: Mechanisms of Disease*, 3(1):99–126, 2008.

[123] Y. Rubtsov and A. Rudensky. TGFβ signalling in control of T–cell–mediated self–reactivity. *Nature Reviews Immunology*, 7(6):443–453, 2007.

[124] C. Rudd, A. Taylor, et al. CD28 and CTLA–4 coreceptor expression and signal transduction. *Immunological Reviews*, 229(1):12–26, 2009.

[125] S. Sakaguchi, K. Wing, et al. Regulatory T cells: How do they suppress immune responses? *International Immunology*, 21(10):1105–1111, 2009.

[126] A. Salmaninejad, S. Valilou, et al. PD–1/PD–L1 pathway: Basic biology and role in cancer immunotherapy. *Journal of Cellular Physiology*, 234(10):16824–16837, 2019.

[127] C. I. Sandefur, R. C. Boucher, et al. Mathematical model reveals role of nucleotide signaling in airway surface liquid homeostasis and its disregulation in cystic fibrosis. *Proceedings of the National Academy of Sciences*, 114(35):E7272–E7281, 2017.

[128] S. Sanjabi, L. Zenewicz, et al. Anti–inflammatory and pro–inflammatory roles of TGF–β, IL–10, and IL–22 immunity and autoimmunity. *Current Opinion in Pharmacology*, 9(4):447–453, 2009.

[129] S. Sarv Ahrabi and A. Momenzadeh. Metaheuristics and pontryagin's minimum principle for optimal therapeutic protocols in cancer immunotherapy: a case study and methods comparison. *Journal of Mathematical Biology*, 81(2):691–723, 2020.

[130] G. Schett and M. Neurath. Resolution of chronic inflammatory disease: universal and tissue–specific concepts. *Nature Communications*, 9(3261), 2018.

[131] A. Schmidt, N. Oberle, et al. Molecular mechanisms of Treg–mediated T cell suppression. *Frontiers in Immunology*, 3:51, 2012.

[132] A. F. J. Schmidt, O. R. Rodrigues, et al. Mediastinal lymph node distribution, size and number: definitions based on an anatomical study. *Jornal Brasileiro de Pneumologia*, 33(2):134–140, 2007.

[133] C. Serhan, P. Ward, et al., editors. *Fundamentals of Inflammation*. Cambridge University Press, Camrbidge, UK, 2010.

[134] R. Serre, S. Benzekry, et al. Mathematical modeling of cancer immunotherapy and its synergy with radiotherapy. *Cancer Research*, 76(17):4931–4940, 2016.

[135] A. Shapouri-Moghaddam, S. Mohammadian, et al. Macrophage plasticity, polarization, and function in health and disease. *Journal of Cellular Physiology*, 233(9):6425–6440, 2018.

[136] J. A. Sharp, A. P. Browning, et al. Optimal control of acute myeloid leukaemia. *Journal of Theoretical Biology*, 2019.

[137] A. Shinde and N. Frangogiannis. Mechanisms of fibroblast activation in the remodeling myocardium. *Current Pathobiology Reports*, 5(2):145–152, 2017.

[138] A. Shirin, I. S. Klickstein, et al. Prediction of optimal drug schedules for controlling autophagy. *Scientific Reports*, 9(1), 2019.

[139] E. Sikora. Activation–induced and damage–induced cell death in aging human T cells. *Mechanisms of Ageing and Development*, 151:85–92, 2015.

[140] M. J. Simpson, R. E. Baker, et al. Practical parameter identifiability for spatio–temporal models of cell invasion. *Journal of the Royal Society Interface*, 17(164), 2020.

[141] D. Sojka, Y.-H. Huang, et al. Mechanisms of regulatory T–cell suppression – a diverse arsenal for a moving target. *Immunology*, 124(1):13–22, 2008.

[142] E. Sontag. A dynamic model of immune responses to antigen presentation predicts different regions of tumor or pathogen elimination. *Cell Systems*, 4(2):231–241.e11, 2017.

[143] C. Souders, S. Bowers, et al. Cardiac fibroblast. *Circulation Research*, 105(12):1164–1176, 2009.

[144] B. Su, W. Zhou, et al. Mathematical modelling of immune response in tissues. *Computational and Mathematical Methods in Medicine*, 10(1):9–38, 2009.

[145] F. Tang, X. Du, et al. Anti–CTLA–4 antibodies in cancer immunotherapy: selective depletion of intratumoral regulatory T cells or checkpoint blockade? *Cell & Bioscience*, 8(30), 2018.

[146] L. Thomas. On immunosurveillance in human cancer. *The Yale Journal of Biology and Medicine*, 55:329–333, 1982.

[147] C. Thäle and A. Kiderlen. Sources of interferon–gamma (IFN–γ) in early immune response to Listeria monocytogenes. *Immunobiology*, 210(9):673–683, 2005.

[148] D. Tirziu, F. J. Giordano, et al. Cell communications in the heart. *Circulation*, 122(9):928–937, 2010.

[149] M. Tocut, R. Brenner, et al. Autoimmune phenomena and disease in cancer patients treated with immune checkpoint inhibitors. *Autoimmunity Reviews*, 17(6):610–616, 2018.

[150] E. Trombetta, M. Ebersold, et al. Activation of lysosomal function during dendritic cell maturation. *Science*, 299(5611):1400–1403, 2003.

[151] R. Vagnozzi, A. Johansen, et al. CARdiac immunotherapy: T cells engineered to treat the fibrotic heart. *Molecular Therapy*, 27(11):1869–1871, 2019.

[152] S. A. van der Vegt, L. Polonchuk, et al. Mathematical modelling of autoimmune myocarditis and the effects of immune checkpoint inhibitors. *Journal of Theoretical Biology*, 537:111002, 2022.

[153] S. A. van der Vegt, Y.-J. Wang, et al. A model-informed approach to assess the risk of immune checkpoint inhibitor-induced autoimmune myocarditis. *Frontiers in Pharmacology*, 13, 2022.

[154] D. Vdovenko and U. Eriksson. Regulatory role of CD4+ T cells in myocarditis. *Journal of Immunological Research*, 2018:1–11, 2018.

[155] N. Vrisekoop, I. den Braber, et al. Sparse production but preferential incorporation of recently produced naïve T cells in the human peripheral pool. *Proceedings of the National Academy of Sciences*, 105(16):6115–6120, 2008.

[156] L. T. Weckbach, U. Grabmaier, et al. Midkine drives cardiac inflammation by promoting neutrophil trafficking and NETosis in myocarditis. *Journal of Experimental Medicine*, 216(2):350–368, 2019.

[157] S. Wei, C. Duffy, et al. Fundamental mechanisms of immune checkpoint blockade therapy. *Cancer Discovery*, 8(9):1069–1086, 2018.

[158] T. Wen and M. Rothenberg. The regulatory function of eosinophils. *Microbiology Spectrum*, 4(5), 2016.

[159] J. Westman, S. Grinstein, et al. Phagocytosis of necrotic debris at sites of injury and inflammation. *Frontiers in Immunology*, 10, 2020.

[160] E. Wherry. T cell exhaustion. *Nature Immunology*, 12:492–499, 2011.

[161] F.-G. Wieland, A. Hauber, et al. On structural and practical identifiability. *Current opinion in systems biology*, 25:60–69, 2021.

[162] J. D. Wolchok, B. Neyns, et al. Ipilimumab monotherapy in patients with pretreated advanced melanoma: a randomised, double–blind, multicentre, phase 2, dose–ranging study. *The Lancet Oncology*, 11(2):155–164, 2010.

[163] L. Wu, S. Ong, et al. Cardiac fibroblasts mediate IL–17A–driven inflammatory dilated cardiomyopathy. *Journal of Experimental Medicine*, 211(7):1449–1464, 2014.

[164] H. Yang, Y. Chen, et al. Interleukin–13 reduces cardiac injury and prevents heart dysfunction in viral myocarditis via enhanced M2 macrophage polarization. *Oncotarget*, 8(59):99495–99503, 2017.

[165] Y. Yang, G. Jiang, et al. Programmed cell death and its role in inflammation. *Military Medical Research*, 2(12), 2015.

[166] A. Yin, D. J. A. Moes, et al. A review of mathematical models for tumor dynamics and treatment resistance evolution of solid tumors. *CPT: Pharmacometrics and Systems Pharmacology*, 8(10):720–737, 2019.

[167] K. W. Yoon. Dead cell phagocytosis and innate immune checkpoint. *BMB Reports*, 50(10):496–503, 2017.

[168] J.-L. Yu and S.-J. Jang. A mathematical model of tumor–immune interactions with an immune checkpoint inhibitor. *Applied Mathematics and Computation*, 362(124523), 2019.

[169] M. Zarak-Crnkovic, G. Kania, et al. Heart non–specific effector CD4+ T cells protect from post–inflammatory fibrosis and cardiac dysfunction in experimental autoimmune myocarditis. *Basic Research in Cardiology*, 115(6), 2020.

[170] L. Zhang, M. Jones-O'Connor, et al. Cardiotoxicity of immune checkpoint inhibitors. *Current Treatment Options in Cardiovascular Medicine*, 21(7):32, 2019.

[171] Y. Zhang, M. Zhang, et al. Silencing MicroRNA–155 attenuates cardiac injury

and dysfunction in viral myocarditis via promotion of M2 phenotype polarization of macrophages. *Scientific Reports*, 6(22613), 2016.

[172] J. Zhu, H. Yamane, et al. Differentiation of effector CD4 T cell populations. *Annual Review of Immunology*, 28(1):445–489, 2010.