UNIVERSITY OF CALIFORNIA,
IRVINE

# MATHEMATICAL MODELING OF GROWTH, EVOLUTION, AND CONTROL IN BIOLOGICAL POPULATIONS

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Mathematics

by

Anthony Zamora

Dissertation Committee:
Professor Natalia Komarova, Chair
Professor Knut Solna
Professor German Enciso
Associate Professor Anna Ma

2026

# DEDICATION

To Natalie, my colleagues, and my friends and family.

# TABLE OF CONTENTS

# LIST OF FIGURES

# LIST OF TABLES

# ACKNOWLEDGMENTS

I would like to sincerely thank my advisor, Natalia Komarova, for her support and dedication to my academic journey during my graduate studies. I am truly appreciative of her mentorship and the many ways that she has contributed to my academic and professional growth.

I am indebted to my friends and peers from my graduate program, including, in no particular order, Arghya Raskshit, Alex Luna, Thanh Tran, Yibiao Ruan, Vignesh Iyer, and Michael Womack.

Lastly, I am grateful for my family for all the unconditional support that they have given me over the years. My achievements would not have been possible without their help and encouragement.

# VITA

## Anthony Zamora

### EDUCATION

**Doctor of Philosophy in Mathematics** **2026**
University of California, Irvine *Irvine, CA*

**Master of Science in Mathematics** **2019**
California State University, Los Angeles *Los Angeles, CA*

**Bachelor of Science in Mathematics** **2014**
University of California, Los Angeles *Los Angeles, CA*

**Mathematics**

### PUBLICATIONS

Zamora, A., Komarova, N. L., and Roy, S. (2026). Optimal melanoma treatment protocols: A bilinear control model. Mathematical Biosciences, 400, 109764.

Moreau, A., Nguyen, D.T., Hinbest, A.J. et al. Surface remodeling and inversion of cell-matrix interactions underlie community recognition and dispersal in Vibrio cholerae biofilms. Nat Commun 16, 327 (2025).

# ABSTRACT OF THE DISSERTATION

Mathematical Modeling of Growth, Evolution, and Control in Biological Populations

By

Anthony Zamora

Doctor of Philosophy in Mathematics

University of California, Irvine, 2026

Professor Natalia Komarova, Chair

Biological populations are heterogeneous, spatially structured, and responsive to changing environments. These features are especially important in tumors, where treatment alters the relative fitness of susceptible and resistant cells, and in bacterial biofilms, where environmental flow affects cells differently according to their phenotype and local surroundings. This dissertation develops deterministic and spatial stochastic models to investigate how growth, evolution, and environmental change shape population dynamics and control.

The first part of the dissertation considers treatment scheduling for melanoma containing treatment-susceptible and treatment-addicted resistant cells. A bilinear ordinary differential equation model is formulated in which treatment suppresses susceptible cells but promotes resistant cells, while treatment withdrawal reverses these effects. An optimal control problem balances cumulative and terminal tumor burden against treatment cost. Pontryagin's maximum principle is used to derive the switching function and characterize the possible structure of optimal protocols. For biologically relevant initial conditions with a small resistant population, the resulting schedules have a monotone structure: full treatment may be followed by an intermediate singular regimen and ultimately by permanent treatment cessation. Repeated transitions between treatment and no treatment do not arise within this

model. A sequential quadratic Hamiltonian method with finite-step convergence is then used to compute the controls numerically. Parameter estimation from published melanoma data provides a biologically informed application of the analytical and computational framework.

The dissertation next examines how spatial competition modifies melanoma dynamics. A two-dimensional stochastic agent-based model represents susceptible and resistant cells undergoing division and death during alternating treatment and no-treatment phases. Because division requires neighboring space, suppression of susceptible cells can release spatial constraints and facilitate resistant-cell expansion. Time-series data generated by the model are analyzed using sparse identification of nonlinear dynamics. Effective logistic equations capture important features of aggregate cell growth, providing a connection between individual-based spatial simulations and lower-dimensional models that may be suitable for future control formulations.

Finally, the same event-driven spatial stochastic framework is adapted to model producer and non-producer phenotypes in Vibrio cholerae biofilms. Matrix production imposes a reproductive cost but protects producer cells from flow-induced removal, particularly within producer-rich neighborhoods. Density-dependent phenotypic switching couples changes in population size to changes in composition. Simulations show that flow preferentially removes non-producers, reduces population density, and shifts phenotypic conversion toward the adhesive producer state, thereby maintaining producer-rich clusters. An extension to producer-cheater competition demonstrates that non-producing cells persist more effectively when they can benefit from matrix generated by neighboring producers.

These studies show how environmental conditions can reverse relative fitness, how spatial interactions generate population-level dynamics not captured by well-mixed models, and how mathematical analysis and computation can connect biological mechanisms to questions of population control.

# Chapter 0

# Introduction

Biological populations are rarely homogeneous. Tumors may contain cells with different levels of sensitivity to treatment, while microbial communities may contain phenotypes that differ in their production of shared resources. These differences affect not only the growth of individual cells but also the behavior of the population as a whole. Moreover, the relative fitness of each type is often dependent on its environment. A cancer cell that is resistant to a drug may grow more successfully during treatment but incur a fitness cost when treatment is withdrawn. A bacterium that produces an adhesive extracellular matrix may reproduce more slowly than a non-producer, yet remain attached more effectively when exposed to flow. Thus, growth, competition, and selection cannot always be understood as fixed properties of individual types; they emerge from interactions among population composition, environmental conditions, and spatial organization.

Mathematical models provide a framework for identifying the mechanisms that govern these systems and for determining which biological details are necessary to explain their observed behavior. The appropriate level of description depends on the question being asked. Ordinary differential equation models describe changes in population size through aggregate

variables and are particularly useful for analysis, parameter estimation, and optimization. Spatial stochastic and agent-based models instead represent individual cells and their local interactions, allowing crowding, neighborhood structure, and random events to influence population-level outcomes. These approaches are complementary. A low-dimensional model can expose general dynamical principles and make an optimal control problem analytically tractable, while a spatial model can test the consequences of assumptions that are hidden by a well-mixed approximation.

This dissertation develops mathematical models of heterogeneous biological populations in two applications: cancer (specifically, melanoma) and bacterial biofilms (specifically, V. Cholerae). Although these systems differ substantially in their biological setting, they share a common mathematical structure. Each contains multiple cell types whose growth rates depend on the environment. Each involves a tradeoff: drug resistance can protect a melanoma cell during treatment while reducing its fitness after drug withdrawal, and extracellular-matrix production can improve bacterial adhesion while imposing a reproductive cost. In each case, changes in the environment alter the balance between competing types. Treatment changes the relative fitness of susceptible and resistant tumor cells, whereas fluid flow and nutrient availability change the relative advantage of producer and non-producer bacteria. These common features make both systems natural settings to study how environmental variation, phenotypic heterogeneity, and spatial structure interact.

Three related themes run throughout the dissertation: growth, evolution, and control. Here, evolution is understood broadly as a change in the composition of a population through growth, death, removal, or mutations/phenotypic switching. In all the systems considered here, an external condition modifies the demographic rates of the population and thereby changes its composition and long-term behavior.

## 0.1 Growth and competition in heterogeneous populations

A central feature in modeling heterogeneous populations is that the type with the greatest short-term growth advantage need not dominate under all environmental conditions. Fitness is context dependent. In the melanoma models considered here, treatment-sensitive cells grow in the absence of treatment but decline when treatment is applied. Resistant cells exhibit the opposite response: they have a relatively higher growth rate during treatment but grow poorly, or decay, after treatment is removed. This phenomenon is often described as “addiction to treatment". Treatment therefore plays two roles. It suppresses the susceptible population, but it also creates an environment favorable to resistant cells. Removing treatment reverses these selective pressures.

A related tradeoff occurs in biofilms formed by Vibrio cholerae. Producer cells secrete extracellular material that promotes adhesion to a surface and to neighboring cells. Production is costly, however, and reduces their reproductive rate relative to non-producing cells. In the absence of strong removal, non-producers may benefit from their higher rate of division. Under flow, producers receive an opposing advantage because their adhesive properties reduce the probability of being washed away. The outcome of competition is therefore determined not by reproduction alone but by the combined effects of growth, removal, local population structure, and phenotype switching.

These examples illustrate a general principle: an environmental intervention acts as a selective force. Drug administration and external flow are very different biologically, but mathematically both modify type-dependent rates. The resulting dynamics depend on how long the system remains in a given environment, how rapidly the population responds, and whether interactions are global or local. This perspective provides the conceptual connection between the cancer and biofilm applications developed in this dissertation.

## 0.2 Population-level and individual-based descriptions

The first melanoma model is formulated as a system of ordinary differential equations. The state variables represent the total populations of treatment-susceptible and treatment-addicted resistant cells. Treatment enters as a time-dependent control that modifies the net growth rates of both populations. This formulation assumes that cells of the same type are interchangeable and that spatial position does not directly affect their growth. Its simplicity makes it possible to characterize treatment schedules using optimal control theory and to determine how the balance between tumor burden and treatment cost shapes the optimal protocol.

A well-mixed model cannot, however, represent competition for physical space. In a growing tumor, the death of susceptible cells may create space into which resistant cells can divide. Consequently, stronger suppression of one population can indirectly accelerate the expansion of another, even when the intrinsic growth parameters of the latter remain unchanged. Such effects depend on the spatial arrangement of cells and cannot be recovered from independent exponential growth equations without additional effective terms.

Agent-based models address this limitation by representing cells as individuals occupying sites on a lattice. Division depends on the availability of neighboring space, and all the processes such as divisions, deaths, removal, mutations and/or phenotype conversion occur through stochastic events. Population-level behavior then emerges from many local interactions. Spatial models can produce clustering, boundary-driven growth, and competition for empty sites, all of which may alter conclusions obtained from a well-mixed approximation.

At the same time, detailed spatial simulations can be difficult to analyze or optimize directly. This motivates a connection between the two modeling scales. Data generated by an agent-based model can be used to identify effective population-level dynamics, including nonlinear growth laws that account indirectly for spatial limitation. In this dissertation,

sparse identification methods are used to investigate whether the aggregate dynamics of the spatial melanoma model can be represented by a low-dimensional system of differential equations. This creates a pathway from spatial stochastic simulations back to models suitable for optimal control.

## 0.3 Environmental change, selection, and control

In cancer treatment, the external environment can be manipulated by the drugs. A treatment schedule determines when selective pressure is applied and how strongly it acts. An intuitive strategy might alternate between treatment and rest: treatment reduces susceptible cells, while rest periods exploit the reduced relative fitness of resistant cells. Such reasoning motivates intermittent therapy, but it does not determine the duration or timing of treatment intervals, nor does it account systematically for treatment toxicity and tumor burden.

Optimal control theory makes these tradeoffs explicit. The treatment schedule is treated as a function to be chosen, subject to the population dynamics and bounds on drug intensity. An objective functional assigns costs to tumor burden, terminal tumor size, and treatment exposure. The resulting problem asks not whether treatment should be continuous or intermittent in a predetermined pattern, but which admissible schedule best balances the competing objectives. The analysis in this dissertation shows that, for the bilinear model considered, optimal schedules have a constrained monotone structure: full treatment may be followed by an intermediate singular phase and then by permanent treatment cessation. Repeated cycles of treatment and withdrawal are not optimal within this modeling framework.

The biofilm model does not formulate an optimal control problem, but it addresses the same relationship between environment and composition. Flow changes the removal rates of producer and non-producer cells, while density-dependent switching changes the supply of each

phenotype. The model therefore asks how an externally imposed condition and an internally regulated transition mechanism jointly determine population persistence, spatial patterning, and coexistence. In this sense, the biofilm application complements the melanoma control problem: one studies how an intervention should be chosen, while the other examines how environmental pressure reorganizes a population through ecological and phenotypic processes.

## 0.4 Objectives and organization of the dissertation

This dissertation uses mathematical models to investigate how biological heterogeneity, environmental conditions, and spatial interactions shape the dynamics and control of evolving populations. The specific questions are:

- What qualitative forms can an optimal treatment protocol take when susceptible and resistant melanoma cells respond oppositely to treatment?
- How can these protocols be computed numerically, and how do they depend on model parameters and experimentally estimated growth rates?
- How does spatial competition alter the aggregate growth of susceptible and resistant melanoma cells?
- Can effective differential equations be identified from spatially generated tumor data and ultimately used in a control problem?
- How do reproductive cost, adhesion, flow, and phenotype switching determine the composition and spatial organization of a bacterial biofilm?

Chapter 1 develops the analytical theory for a bilinear optimal control model of melanoma containing treatment-susceptible and treatment-addicted resistant cells. Treatment is rep-

resented by a bounded, time-dependent control, and the objective functional accounts for cumulative tumor burden, terminal tumor size, and treatment cost. Pontryagin's maximum principle is used to derive the adjoint system and switching function. The geometry of the resulting trajectories is then analyzed to characterize the possible structure of optimal treatment protocols.

Chapter 2 develops the numerical and data-informed aspects of the melanoma control problem introduced in Chapter 1. A sequential quadratic Hamiltonian method is used to compute treatment protocols, and its convergence properties are examined. The numerical solutions are compared with the structures predicted by the analytical theory and are used to investigate the effects of treatment cost, initial tumor composition, and mutation. Published experimental data are then used to estimate the model's growth parameters and study the protocols obtained in a biologically informed parameter regime.

Chapter 3 introduces a spatial agent-based model of melanoma. Susceptible and resistant cells divide and die on a two-dimensional lattice under alternating treatment and no-treatment conditions. The spatial formulation makes it possible to investigate how competition for empty space couples the dynamics of the two populations. Time-series data generated by the simulations are subsequently analyzed using sparse identification methods to obtain effective ordinary differential equation models. This provides an initial connection between mechanistic spatial simulations and lower-dimensional models that may eventually be incorporated into an optimal control framework.

Chapter 4 continues with spatial stochastic models using the system of a Vibrio cholerae biofilm. Producer and non-producer cells occupy a two-dimensional lattice and undergo division, death, flow-induced removal, and density-dependent phenotypic conversion. The model represents both the reproductive cost of VPS production and the adhesive protection it provides under flow. Simulations are used to investigate how the onset of flow changes population composition, spatial organization, and the balance of conversion between the two

phenotypes. An extension distinguishes non-producers from cheaters that do not produce VPS but can benefit from matrix generated by neighboring producer cells.

These studies demonstrate how mathematical modeling can aid in understanding complex dynamics of biological systems undergoing growth, evolution, and environmental variation. Growth changes the abundance of competing types; environmental conditions determine which types are favored; spatial structure regulates access to resources and empty space; and intervention changes both the size and composition of the population. By moving between deterministic and stochastic descriptions, and between population-level and individual-based models, this research seeks to identify both general principles and application-specific mechanisms.

# Chapter 1

# Analytical Characterization of Optimal Melanoma Treatment Protocols

## 1.1 Biological motivation

Melanoma is a widespread cancer, comprising 1.7% of all new cancer diagnoses and 70% of all skin cancer-related deaths as of 2018 [29]. Once melanoma metastasizes, prognosis is poor; in 2012 the median survival time after a diagnosis of metastatic melanoma was only 6 to 10 months [32]. The advent of targeted therapies has dramatically increased this, however, to a median of 22 to 25 months for patients treated with a BRAFi and a mitogen-activated protein kinase-kinase (MEK) inhibitor [29]. Unfortunately, a complete response is seen in only 16% of the patients - the remainder experience tumor rebound with a drug-resistant phenotype.

The modern trajectory of treatment-strategy development in BRAF-mutant melanoma began with the introduction of selective BRAF inhibitors such as vemurafenib and dabrafenib in

2010–2011, when early trials demonstrated rapid and dramatic tumor regressions in patients with BRAF V600E melanoma [6]. Yet these responses were short-lived, and early mechanistic work quickly revealed multiple routes to MAPK pathway reactivation, including NRAS mutations, BRAF amplification, and alternative splicing [22, 30]. While the mechanisms by which resistance to the BRAF inhibitors occurs vary, approximately 70% of the resistance mechanisms have been reported to act by up-regulating the natural mitogen-activated protein kinase (MAPK) signaling pathway, either by up-regulating BRAF expression or removing downstream dependence on BRAF for kinase activation [15]. The realization that continuous targeted therapy almost inevitably led to resistance spurred efforts to find strategies capable of prolonging response, including the first explorations of drug “holidays" or non-standard dosing schedules.

The field took a major turn when it was reported that vemurafenib-resistant melanoma clones in mice displayed a striking form of drug addiction: resistant tumors grew poorly or regressed when the drug was removed because of toxic hyperactivation of ERK signaling [34]. In their xenograft models, intermittent dosing (four weeks on, two weeks off) significantly extended tumor control, suggesting that cycling therapy could exploit fitness costs of resistance. Similar findings emerged with dual BRAF/MEK inhibition, where resistant lines again showed reduced fitness in the absence of drug [21]. This set the stage for a pivotal conceptual shift: that the behavior of resistant clones might be exploited through dosing strategies rather than simply overcome by additional inhibitors.

In parallel, mathematical and theoretical studies, notably by A.R.A. Anderson and colleagues, provided an ecological foundation for treatment breaks by showing how tumor populations composed of sensitive and resistant cells could be manipulated through competition dynamics. Paper [18] established a general framework for virtual clinical trials by building mechanistic ODE models of phenotypic resistance, calibrating them with in vitro data, and generating heterogeneous virtual patient cohorts to identify more effective schedules

for combination therapy. In Ref [31], they advanced this approach by integrating single-cell transcriptional profiling, flow cytometry, in vitro assays, and xenograft experiments to reveal how transcriptional-state switching drives non-genetic resistance to BRAF inhibitors and to demonstrate, in vivo, that personalized adaptive therapy schedules guided by real-time tumor dynamics outperform continuous or fixed intermittent dosing. In [17], they extended their modeling strategy directly into the clinical domain, fitting competition and switching models to longitudinal measurements from metastatic melanoma patients treated with BRAF/MEK inhibitors and showing through simulation that adaptive therapy (see also [36, 14, 35]) could substantially prolong tumor control while reducing drug exposure.

As this conceptual foundation matured, more sophisticated experimental systems emerged, including integrating single-cell transcriptional measurements with xenograft studies [9]. Yet when fixed intermittent therapy reached clinical testing, the results diverged sharply from the preclinical promise. A randomized phase II trial comparing continuous dabrafenib–trametinib therapy with a 7-days-on, 7-days-off schedule found that continuous dosing achieved significantly longer progression-free survival, contradicting the predictions from mouse studies and raising questions about the clinical relevance of drug-addicted resistant clones [1].

In this chapter we suggest that the observed results of clinical trials could actually be consistent with the results of mathematical modeling in the context of "drug-addicted" resistant clones. To do this, we build on the work in [16], which develops and analyzes a bilinear control model that reflects the dynamics of susceptible and resisted ("addicted") cells in melanoma treatment, and includes treatment is either "on" or "off." The authors pose an optimal control problem that minimizes total tumor burden over a fixed treatment horizon and rigorously analyze it using the Pontryagin Maximum Principle, identifying both bang–bang controls (pure on/off phases) and singular controls (intermediate effective dosing that emerges mathematically) as potential optimal regimens.

Here, we extend the framework of [16] by using optimal control theory to uncover a treatment

regimen that minimizes overall side effects of the drug in addition to tumor burden on the patient. The relaxed optimal control problem is stated and solved in this chapter. Further, for the numerical implementation of the control problems, we use a PMP-based sequential quadratic Hamiltonian (SQH) method [2] and investigate its convergence properties. Initially inspired by needle variations in the proof of the PMP [24], this method has been used in context of optimal control problems arising in tomography (e.g., see [3, 7, 26, 27]) but has not been explored in the field of oncology. The implementation of the SQH method does not need differentiability of the governing model nor of the objective functional with respect to the optimal control variable, which makes it quite robust compared to traditional gradient-based optimal control methods. Furthermore, we show that the SQH algorithm converges in finitely many steps, which makes it computationally very efficient for handling high-dimensional control problems that usually arise in oncology.

The rest of this chapter is organized as follows. Section 1.2 introduces the ODE model that describes the system of susceptible and resistant (addicted) cancer cells under a treatment regimen consisting of alternating phases of treatment and no treatment. The initial statement of the optimal control problem does not have a solution. As in [16], we will consider the relaxed optimal control problem for which we will solve analytically. Section 1.3 presents the application of the Pontryagin maximum principle followed by the analytic solution to the optimal control problem. In Section 2.2, we present the SQH method used to solve for the optimal control and corroborate the analytic solution. In Section 2.3, we present the results to the optimal control problem with other types of resistant mutants, followed by concluding discussion and remarks in the final section.

## 1.2 Preliminaries

### 1.2.1 Model set-up

Consider a setting where cells are assumed to mix well and dynamics are deterministic, which can be described by ordinary differential equations. The preliminary modeling uses ODEs to describe the evolutionary co-dynamics of susceptible ($X$) and resistant ($Y$) melanoma cells:

$$\text{No treatment:} \quad \begin{cases} \frac{dX}{d\tau} = & L_x^{\text{off}}(1-\mu)x - D_x^{\text{off}}x, \\ \frac{dY}{d\tau} = & L_x^{\text{off}}\mu x + (D_y^{\text{off}} - D_y^{\text{off}})y, \end{cases} \tag{1.1}$$

$$\text{In treatment:} \quad \begin{cases} \frac{dX}{d\tau} = & L_x^{\text{on}}(1-\mu)x - D_x^{\text{on}}x, \\ \frac{dY}{d\tau} = & L_x^{\text{on}}\mu x + (L_y^{\text{on}} - D_y^{\text{on}})y. \end{cases} \tag{1.2}$$

Here, $\tau$ is the dimensional time, $X(\tau)$ and $Y(\tau)$ are numbers of susceptible and mutant cells, the symbols $L$ and $D$ stand for division and death rates of cells, the subscripts indicate which population the rate characterizes, and the subscripts refer to the "off" and "on" treatment regimens. The symbol $\mu$ refers to the mutation rate, where resistant cells are generated from the susceptible cells through *de novo* mutations.

Let us define the net growth rates as follows:

$$\Gamma_x^{\text{off}} = L_x^{\text{off}} - D_x^{\text{off}}, \quad \Gamma_y^{\text{off}} = L_y^{\text{off}} - D_y^{\text{off}}, \quad \Gamma_x^{\text{on}} = L_x^{\text{on}} - D_x^{\text{on}}, \quad \Gamma_y^{\text{on}} = L_y^{\text{on}} - D_y^{\text{on}}. \tag{1.3}$$

To non-dimensionalize the system, we rescale the population sizes with a reference total tumor size, $\bar{X}$, and rescale the time with the net growth rate of susceptible cells in the

absence of treatment:

$$x = \frac{X}{\bar{X}}, \quad y = \frac{Y}{\bar{X}}, \quad t = \tau|\Gamma_x^{off}|$$

We rescale all the division, death, and net growth rates with the rate $|\Gamma_x^{on}|$ and denote them by the corresponding lower-case symbol:

$$\gamma_x^{on} = \frac{\Gamma_x^{on}}{|\Gamma_x^{on}|} = 1, \quad \gamma_x^{off} = \frac{\Gamma_x^{off}}{|\Gamma_x^{on}|}, \quad \gamma_y^{on} = \frac{\Gamma_y^{on}}{|\Gamma_x^{on}|}, \quad \gamma_y^{off} = \frac{\Gamma_y^{off}}{|\Gamma_x^{on}|}.$$

and similar with the division and death rates. This rescaling procedure yields the following dimensionless equations:

$$\text{No treatment:} \quad \begin{cases} \dot{x} = & l_x^{\text{off}}(1-\mu)x - d_x^{\text{off}}x, \\ \dot{y} = & l_x^{\text{off}}\mu x + (l_y^{\text{off}} - d_y^{\text{off}})y, \end{cases} \tag{1.4}$$

$$\text{In treatment:} \quad \begin{cases} \dot{x} = & l_x^{\text{on}}(1-\mu)x - d_x^{\text{on}}x, \\ \dot{y} = & l_x^{\text{on}}\mu x + (l_y^{\text{on}} - d_y^{\text{on}})y. \end{cases} \tag{1.5}$$

In most of this chapter, we neglect $\mu$ for simplicity, and then expand the description to nonzero mutation rates. Taking $\mu = 0$, the system can be written as

$$\text{No treatment:} \quad \begin{cases} \dot{x} = & \gamma_x^{\text{off}}x, \\ \dot{y} = & \gamma_y^{\text{off}}y, \end{cases} \tag{1.6}$$

$$\text{In treatment:} \quad \begin{cases} \dot{x} = & \gamma_x^{\text{on}}x, \\ \dot{y} = & \gamma_y^{\text{on}}y. \end{cases} \tag{1.7}$$

where

$$\gamma_x^{\text{off}} = l_x^{\text{off}} - d_x^{\text{off}}, \quad \gamma_y^{\text{off}} = l_y^{\text{off}} - d_y^{\text{off}}, \quad \gamma_x^{\text{on}} = l_x^{\text{on}} - d_x^{\text{on}}, \quad \gamma_y^{\text{on}} = l_y^{\text{on}} - d_y^{\text{on}}. \tag{1.8}$$

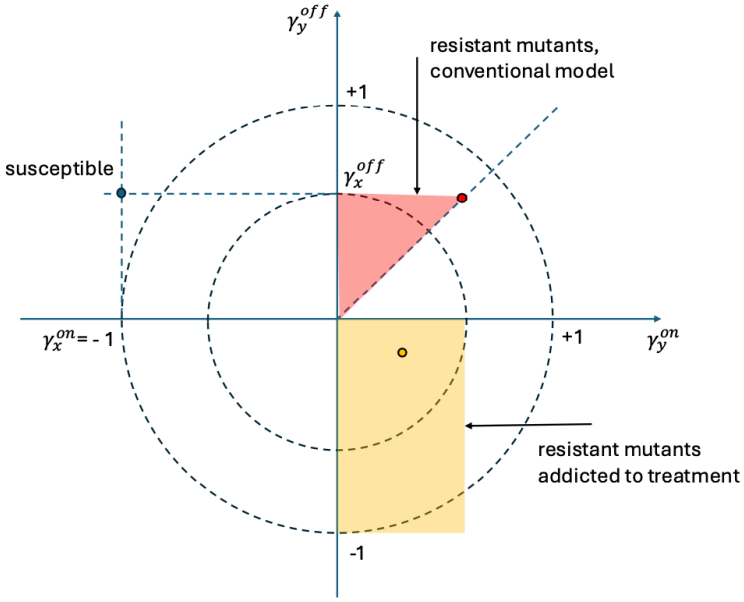


Figure 1.1: A schematic showing the parameter regions for two models of resistant cells, the conventional model (red) and the addicted to treatment model (yellow). The parameter pair for the susceptible cells, $(\gamma_x^{on}, \gamma_x^{off})$, is shown by a dot in the 2nd quadrant. The two dots in the yellow and red regions correspond to the example parameter sets listed in equations (2.3).

There are two model structures for the resistant mutants one could explore.

- Conventional model assumes that in the absence of treatment, resistant mutants grow slower than susceptible cells, and in the presence of treatment they continue to grow at the same or somewhat reduced rate. This means that

$$0 < \gamma_y^{on} \leq \gamma_y^{off} \leq \gamma_x^{off}.$$

- In the model of mutants “addicted" to treatment, we assume that in the absence of treatments, resistant mutants die off (but slower than susceptible cells killed by treatment), and

in the presence of treatment, resistant mutants grow (but slower than susceptible cells in the absence of treatment). These assumptions are equivalent to the following inequalities:

$$-1 < \gamma_y^{off} < 0, \quad 0 < \gamma_y^{on} < \gamma_x^{off}.$$

These inequalities are depicted as regions in the $(\gamma_y^{on}, \gamma_y^{off})$ space in figure 1.1. Throughout most of this chapter we focus on the second case, that is, mutants addicted to treatment. Conventional mutants are studied in Appendix 1.5.

While the theory described in the chapter is quite general and only restricts some inequalities on the parameter values, numerical illustrations are shown for specific values. See Section 2.4 that describes applications to a real parameter value set.

### 1.2.2 Statement of the minimization problem

For convenience, let us define the following constants

$$a = \gamma_x^{\text{off}},\ b = \gamma_x^{\text{on}} - \gamma_x^{\text{off}},\ c = \gamma_y^{\text{off}},\ d = \gamma_y^{\text{on}} - \gamma_y^{\text{off}}$$

and rewrite the control system combining equations (1.6) and (1.7)

$$\begin{cases} x'(t) &= (a + bu(t))x \\ y'(t) &= (c + du(t))y \\ x(0) &= x_0,\ y(0) = y_0;\ x_0,\ y_0 > 0 \end{cases} \tag{1.9}$$

where $x_0$ and $y_0$ are the initial conditions and $u(t)$ is a control function such that

$$u(t) \in U = \{0, 1\},$$

with $U$ the control set. In system, (1.9) $u(t) = 1$ and $u(t) = 0$ correspond to treatment and no treatment, respectively. We will consider the control system (1.9) on the time interval $[0, T]$ where $T$ is the total treatment time. The set of admissible control functions $\Lambda_T$ is the set of Lebesgue measurable functions $u(t)$ such that $u(t) \in U$ for almost all $t \in [0, T]$.

We wish to minimize the following objective functional

$$J(u) = \int_0^T (x(t) + y(t))\ dt + \beta(x(T) + y(T)) + \alpha \int_0^T |u(t)|\ dt \tag{1.10}$$

over the set of admissible control functions $\Lambda_T$, where $\beta$ and $\alpha$ are positive weight coefficients. The first term on the right hand side represents the cumulative tumor load, the second term is the tumor load at the final time, $T$, and the last term represents the side-effects that might be associated with treatment and are dose-dependent; we will refer to the coefficient $\alpha$ as treatment cost weight. Now, the control set $U$ is not convex which can mean that there is no solution to the problem. We will treat the problem as in [16] and seek to minimize (1.10) over the set of admissible controls $\Psi_T$. This is the set of Lebesgue measurable functions $u(t)$ satisfying $u(t) \in [0, 1]$ for almost all $t$. In the analysis that follows, we will refer to the special ray

$$\Pi = \left\{(x, y) : x > 0, y > 0, y = -\frac{b}{d}x\right\} \tag{1.11}$$

and singular control

$$u_{\text{sing}} = -\frac{a - c}{b - d} \tag{1.12}$$

whose derivation can be found in [16] and are central in the discussion of a singular regimen. We reiterate Lemma 1 as established in [16] before proceeding to the discussion of the relaxed minimization problem.

**Lemma 1.** For any admissible control $u(t)$, the components of the solution $(x(t), y(t))$ of system (1.9) are absolutely continuous functions that are defined on the interval $[0, T]$ and are subject to inequalities:

$$0 < x(t) < x_{\max} < x_0 e^{aT},\ 0 < y(t) < y_{\max} < y_0 e^{dT}, t \in [0, T]. \tag{1.13}$$

## 1.3 Mutants addicted to treatment

In this section we focus on mutants that are "addicted to treatment". In terms of system (1.6), this means that

$$\gamma_x^{\text{off}} > 0, \quad \gamma_x^{\text{on}} < 0, \quad \gamma_y^{\text{off}} < 0, \quad \gamma_y^{\text{on}} > 0. \tag{1.14}$$

While the susceptible cells grow in the absence of treatment and have a net negative growth rate under treatment, the mutants are the opposite: they grow under treatment, and have a negative net growth rate in its absence. Inequalities (1.14) are equivalent to the following conditions for the coefficients in system (1.9):

$$a > 0,\ b < 0,\ c < 0,\ d > 0. \tag{1.15}$$

We seek to minimize the objective functional, (1.10), subject to $0 \leq u(t) \leq 1$ and system (1.9). We will refer to the resulting optimal trajectory as $u_\star(t)$.

While the formulation of the relaxed problem is applicable for resistant mutants of a very general type, sections 1.3.1 and 1.4.1 solve the problem and discuss the resulting trajectories specifically for addicted mutants. Solutions for other types of mutants are discussed in Appendix 1.5.

### 1.3.1 Behavior of the optimal trajectory

We are interested in the dynamics of system (1.9) in the situation where resistant cells are addicted to treatment. To this end, we will explore the dynamics of the optimal trajectory $(x_\star(t), y_\star(t))$ during the different key phases of treatment that correspond to the optimal control $u_\star(t)$. On an interval where we do not treat, we have $u_\star(t) = 0$ and the equations of system (1.9) become

$$\begin{cases} x'_\star(t) &= ax_\star(t) \\ y'_\star(t) &= cy_\star(t). \end{cases} \tag{1.16}$$

As time, $t$, increases the susceptible cells, $x$, increase and the resistant cells, $y$, decrease. On the other hand, on an interval on which treatment is administered, we have $u_\star(t) = 1$. This means that the equations in (1.9) become

$$\begin{cases} x'_\star(t) &= (a+b)x_\star(t) \\ y'_\star(t) &= (c+d)y_\star(t). \end{cases} \tag{1.17}$$

In this case, as time increases, $x$ decreases while $y$ increases. On the singular control, $u_\star(t) = u_{\text{sing}}$ the equations (1.9) take on the form

$$\begin{cases} x'_\star(t) &= -(ad - bc)/(b - d)x_\star(t) \\ y'_\star(t) &= -(ad - bc)/(b - d)y_\star(t). \end{cases} \tag{1.18}$$

Since $(b - d) < 0$ from (1.15), for $ad - bc < 0$, the steady state $(0, 0)$ is stable and it is unstable when $ad - bc > 0$.

Some important properties of the optimal trajectory are described in Appendix 1.4. Here we briefly summarize them as follows. Given any treatment schedule and initial condition $(x_\star(0), y_\star(0))$, if the optimal trajectory intersects the special ray, $\Pi$, then the optimal trajectory can not ascend above the special ray. This property of the optimal trajectory is important as it eliminates the possibility of oscillatory behavior near the special ray. Additionally, the final switching time, with its associated optimal trajectory, is unique and has some dependence on $\alpha$, meaning that if the final switching time is different, then the optimal trajectory must be associated to a different value of $\alpha$.

## 1.4 Pontryagin's maximum principle and properties of the optimal trajectory for mutants addicted to treatment

We calculate the adjoint system, Hamiltonian, and the important switching function used to determine the optimal solution. To find the optimal control $u_\star(t)$ which minimizes 1.10, we apply the Pontryagin maximum principle. This yields the following Hamiltonian and adjoint

equations. We write the Hamiltonian:

$$H(x, y, u, \psi_1, \psi_2) = (a + bu)x\psi_1 + (c + du)y\psi_2 - [x + y + \alpha u], \tag{1.19}$$

where $\psi_1$ and $\psi_2$ are the adjoint variables. Then, according to the Pontryagin maximum principle, for the optimal control $u_\star(t)$ and the corresponding optimal solution $(x_\star(t), y_\star(t))$ for the system there exists an adjoint function $\psi_\star(t) = (\psi_1^\star(t), \psi_2^\star(t))$, such that

- it is an absolutely continuous solution of the adjoint system

$$\begin{cases} \psi_1^{\star\prime} & = -(a + bu_\star)\psi_1 + 1 \\ \psi_2^{\star\prime} & = -(c + du_\star)\psi_2 + 1 \\ \psi_1^\star(T) & = -\beta,\ \psi_2^\star(T) = -\beta \end{cases} \tag{1.20}$$

- control $u_\star(t)$ maximizes the Hamiltonian $H(x_\star(t), y_\star(t), u, \psi_1^\star(t), \psi_2^\star(t))$ in the variable $u \in [0, 1]$ for almost all $t \in [0, T]$, and therefore it satisfies the following relationship:

$$u_\star(t) = \begin{cases} 1, & \text{if } L(t) > 0, \\ \text{any } u \in [0, 1], & \text{if } L(t) = 0 \\ 0, & \text{if } L(t) < 0 \end{cases} \tag{1.21}$$

where

$$L(t) = H'_u = bx_\star(t)\psi_1^\star(t) + dy_\star(t)\psi_2^\star(t) - \alpha \tag{1.22}$$

is the switching function. We transform the adjoint system for simplification. We set $\phi_1^\star(t) = x_\star(t)\psi_1^\star(t),\ \phi_2^\star(t) = y_\star(t)\psi_2^\star(t)$. This produces the new adjoint system:

$$\begin{cases} \phi_1^{\star\prime} &= x_\star \\ \phi_2^{\star\prime} &= y_\star \\ \phi_1^\star &= -\beta x_\star(T),\ \phi_2^\star = -\beta y_\star(T) \end{cases} \tag{1.23}$$

This gives the new switching function

$$L(t) = b\phi_1(t) + d\phi_2(t) - \alpha \tag{1.24}$$

with derivatives

$$L'(t) = bx_\star(t) + dy_\star(t) \tag{1.25}$$

$$L''(t) = b(a + bu_\star(t))x_\star(t) + d(c + du_\star(t))y_\star(t) \tag{1.26}$$

Solving the system 1.23 and substituting into 1.24 we acquire the following form of the switching function

$$L(t) = -\beta(bx_\star(T) + dy_\star(T)) - \int_t^T (bx_\star(s) + dy_\star(s))\ ds - \alpha \tag{1.27}$$

Our analysis of the switching function and its derivatives $(1.24 - 1.26)$ will determine the optimal control by taking advantage of their relation in (1.21).

Let's now investigate some properties of the optimal trajectory. The lemmas of this appendix provide some analysis of the switching function (1.27) for use in the solution of the relaxed minimization problem in section (1.4.1). The analysis develops some insight into

the positioning of the optimal trajectory as it relates to the special ray. Lemmas 4 and 6 assume that the initial condition, $(x_\star(0), y_\star(0))$, is below the special. In this setting, Lemma 4 says that if the treatment schedule begins with the absence of application of treatment, the schedule will not have a switching time. Lemma 6 shows that for a treatment schedule that begins with application of treatment, the optimal trajectory remains on or below the special ray for the entirety of the schedule. Lemmas 5 and 7 don't depend on the position of the initial condition. Lemma 5 asserts that once the optimal trajectory has intersected the special ray it is no longer possible for the trajectory to terminate above the ray. Lemma 7 shows that the optimal trajectory can not terminate on the special ray since $\alpha > 0$. Lemma 8 demonstrates that the final switching time is unique given the weight factors $\beta$ and $\alpha$. Lemmas 9, 10 , and 11 demonstrate some properties of the optimal trajectory when the initial condition is above the special ray. Lemma 9 states that the optimal trajectory can not be associated to a treatment protocol that begins with a phase of treatment. Lemma 10 shows that the optimal trajectory must remain on or below the special ray after a point of intersection.

**Lemma 2.** If on the interval $\Delta \subset [0, T]$ the optimal control $u_\star(t)$ has the value 0 (which means that $L(t) < 0$), then $L''(t) < 0$. If on this interval the control $u_\star(t)$ takes the value 1 (this implies that $L(t) > 0$), then $L''(t) > 0$.

**Proof of Lemma 2.** If $u_\star(t) = 0$, then

$$L''(t) = abx_\star(t) + cdy_\star(t) < 0$$

since $ab < 0$ and $cd < 0$. If $u_\star(t) = 1$, then

$$L''(t) = b(a+b)x_\star(t) + d(c+d)y_\star(t) > 0$$

since $b(a+b), d(c+d) > 0$. $\square$

**Lemma 3.** If $L(t_1) = L(t_2) = 0$ and $t_1 < t_2$, then $L(t) = 0$ for all $t \in [t_1, t_2]$.

**Proof of Lemma 3.** Consider $t_1 < t_2$ such that $L(t_1) = L(t_2) = 0$ and assume, without loss of generality, that $L(t)$ is positive on $(t_1, t_2)$. We apply the Extreme Value Theorem and $L(t)$ achieves it's maximum at a time $t_0 \in (t_1, t_2)$ so that $L'(t_0) = 0$ and $L''(t_0) \leq 0$. But, $L''(t) \leq 0$ contradicts Lemma 2 since $L(t) > 0$ implies that $L''(t) > 0$ on $(t_1, t_2)$. $\square$

For the following Lemmas we make the assumption $\alpha > 0$.

**Lemma 4.** Assume that the initial condition, $(x_\star(0), y_\star(0))$, is below the special ray. If we begin with no treatment, $u_\star(0) = 0$ (i.e. $L(0) < 0$), then $L(t) < 0$ and $bx_\star(t) + dy_\star(t) < 0$ for all $t \in [0, T]$.

**Proof of Lemma 4.** Assume that $u_\star(0) = 0$. This means that $L(0) < 0$. Since the optimal trajectory $(x_\star(t), y_\star(t))$ begins below the special ray, there is an interval $I$ adjacent to zero such that $L'(t) < 0$. Thus, $L(t)$ is decreasing on $I$. This means that $L(t) < 0$ on $I$. Now, for $L(t)$ to be positive in $[0, T]$ $I$ there must be a time $t_1 > 0$ such that $L(t_1) = 0$. This implies that $L(t)$ achieves a minimum at some time $t_2 \in (0, t_1)$.We note that $L(t) < 0$ on $(0, t_1)$. At $t_2$ we have $L'(t_2) = 0$ and $L''(t_2) \geq 0$. But this contradicts Lemma 2, where we must have

$L''(t) < 0$ whenever $L(t) < 0$. So, $L(t) = 0$ cannot occur on the interval $[0, T]$. This means that we have $u = 0$ on the entirety of $[0, T]$. So, the optimal trajectory moves away from the special ray and remains below. Therefore, $bx_\star(t) + dy_\star(t) < 0$ for all $t \in [0, T]$. $\square$

**Lemma 5.** If the optimal trajectory intersects the special ray, then it is not possible that the optimal trajectory terminates above the special ray.

**Proof of Lemma 5.** Let's suppose to the contrary that the optimal trajectory does terminate above the special ray. Let $t_c$ be the last time that the optimal trajectory intersects the special ray. So, on the interval $I = (t_c, T]$, the optimal trajectory is entirely above the special ray. Then on, $I$, we have $L(t) < 0$. But, $L(t) < 0$ on $I$ means that $u = 0$ on $I$ and the system (1.9) implies that the optimal trajectory moves down and below the special ray, giving us a contradiction. $\square$

**Lemma 6.** Assume that the initial condition, $(x_\star(0), y_\star(0))$, is below the special ray. If we begin with treatment, $u(0) = 1$ (i.e. $L(0) > 0$), then $bx_\star(t) + dy_\star(t) \leq 0$ for all $t \in [0, T]$.

**Proof of Lemma 6.** There are three distinct scenarios in which $L(t)$ may have its first zero. We will consider each as a separate case. Let $t_0$ be the first time that $L(t)$ attains a zero.

**Case 1.** Suppose $(x_\star(t_0), y_\star(t_0))$ is above the special ray. This means that the optimal trajectory crosses the special ray before $L(t) = 0$. But, we have $L(0) > 0$ and $L'(0) = bx_\star(0) + dy_\star(0) < 0$. Now, $L'(t) < 0$ while the optimal trajectory remains below the special ray. The moment, $t = t_{on}$, that the optimal trajectory crosses the special ray we have

$L'(t_{on}) = 0$. We cross the special ray at exactly one point since otherwise we would have $u_\star(t) = u_{\text{sing}} \neq 1$ implying $L(t) = 0$ on some interval after $t_{on}$, contradicting that $L(t) > 0$. Now, since $L(t)$ attains its first zero at $t_0$ we have that $L(t) > 0$ on $[0, t_0)$. In particular, $L(t) > 0$ and $L'(t) > 0$ on $(t_{on}, t_0)$. So, $L(t)$ is increasing on $(t_{on}, t_0)$. This means that $L(t_0) = 0$ is unachievable.

**Case 2.** Suppose $(x_\star(t_0), y_\star(t_0))$ is below the special ray. This means that $L'(t_0) = bx_\star(t_0) + dy_\star(t_0) < 0$. We claim that it is necessary that $L(t) < 0$ for all $t \in (t_0, T]$. Suppose not, then there is another point $t_1 \in (t_0, T)$ such that $L(t) = 0$ and $L(t) < 0$ on $(t_0, t_1)$. This means that $L(t)$ attains a minimum in $(t_0, t_1)$ and thus $L''(t) \geq 0$, but this contradicts the sign of $L''(t)$ as required by Lemma 2. So, $L(t) < 0$ on $(t_0, T]$. This means that the optimal trajectory moves away from the special ray according to 1.16. So, the optimal trajectory is below the special ray $(bx_\star(t) + dy_\star(t))$ on $[0, T]$.

**Case 3.** Suppose $(x_\star(t_0), y_\star(t_0))$ is on the special ray. It suffices to show that the optimal trajectory does not go above the special ray. Let's suppose that this happens. By Lemma 5, we know that $bx_\star(T) + dy_\star(T) \leq 0$. To satisfy this, the optimal trajectory must cross the special ray again. So, there is an interval $I = [t_1, t_2]$, with $L(t_1) > 0$, where the optimal trajectory is above the special ray at $t_1$ and returns at $t_2$. Now, crossing the special ray at $t_2$ requires that the control switches to $u = 0$ on the interval $I$ so that $L(t) \leq 0$ on some interval contained in $I$. This means that $L(t) = 0$ for some $t_3 \in [t_1, t_2]$. However, since $L(t_0) = L(t_3) = 0$, Lemma 3 implies that $L(t) = 0$ on $(t_0, t_3)$. This contradicts the positivity of $L(t)$ on any portion of the interval $(t_1, t_3)$. Thus, $bx_\star(t) + dy_\star(t) \leq 0$ on $[0, T]$.

Lastly, let's suppose that $L(t)$ never attains a zero. We know that $L(0) > 0$ and by Lemma 5 that $bx_\star(T) + dy_\star(T) \leq 0$. We claim that the optimal trajectory may not cross the special

ray. Otherwise, the dynamics in (1.16) and (1.17) would dictate that $L(t) < 0$ above the special ray in order to satisfy $bx_\star(T) + dy_\star(T) \leq 0$ (i.e. termination of the optimal trajectory below the special ray). $\square$

**Lemma 7.** It is not possible that the optimal trajectory terminates on the special ray.

**Proof of Lemma 7:** We show that $bx_\star(T) + dy_\star(T) = 0$ is not possible. Let's assume to the contrary that this does occur. Then $L(T) = -\alpha < 0$. This means that there is an interval $I$ adjacent to $T$ such that $L(t) < 0$ on $I$. So, $u = 0$ on $I$. But this would mean that the optimal trajectory is above the special ray on $I$ in order to satisfy $bx_\star(T) + dy_\star(T) = 0$. Let $t_{on}$ be the most recent time at which the optimal trajectory crossed the special ray. So, the optimal trajectory is above the special ray on $(t_{on}, T)$. According to (1.27), we have $L(t) < 0$ so that $u_\star < 0$ on $(t_{on}, T)$. But the dynamics of 1.16 imply that $x'_\star(t) > 0$ and $y'_\star(t) < 0$. So, for $\epsilon > 0$, small enough, we have that $x_\star(t)$ increasing and $y_\star(t)$ is decreasing on $(t_{on}, t_{on} + \epsilon)$. Then, since $b < 0$ and $d > 0$, $bx_\star(t) + dy_\star(t) < 0$ which means that the optimal trajectory is below the special ray on $(t_{on}, t_{on} + \epsilon)$, contradicting the positioning of the optimal trajectory on $(t_{on}, T)$. $\square$

**Lemma 8:** Let $\left(x_\star^{(1)}(t), y_\star^{(1)}(t)\right), \left(x_\star^{(2)}(t), y_\star^{(2)}(t)\right)$ be two optimal trajectories, with respective $\alpha$ values $\alpha^{(1)}$ and $\alpha^{(2)}$, satisfying the equations yielded by Pontryagin's Maximum Principle. Let $L_1(t)$ and $L_2(t)$ be the switching functions associated with $\left(x_\star^{(1)}(t), y_\star^{(1)}(t)\right)$ and $\left(x_\star^{(2)}(t), y_\star^{(2)}(t)\right)$, respectively. Let $t_1$ and $t_2$ be the two times at which $\left(x_\star^{(1)}(t), y_\star^{(1)}(t)\right)$ and $\left(x_\star^{(2)}(t), y_\star^{(2)}(t)\right)$ begin their final phase of non-treatment. If $t_1 < t_2$ and at least one of $t_1, t_2 > 0$, then $\alpha^{(1)} > \alpha^{(2)}$.

**Proof of Lemma 8:** Suppose that $t_1 < t_2$ and that $t_2 > 0$. Given the dynamics of the constraint system (1.16), we have that $x_\star^{(1)}(t) > x_\star^{(2)}(t)$ and $y_\star^{(1)}(t) < y_\star^{(2)}(t)$ on $[t_1, T]$. This implies that $bx_\star^{(1)}(t) < bx_\star^{(2)}(t)$ and $dy_\star^{(1)}(t) < dy_\star^{(2)}(t)$ on $[t_1, T]$. Then

$$-\left(bx_\star^{(1)}(t) + dy_\star^{(1)}(t)\right) > -\left(bx_\star^{(2)}(t) + dy_\star^{(2)}(t)\right), \ t \in [t_1, T] \tag{1.28}$$

The switching functions are

$$L_1(t) = -\beta(bx_\star^{(1)}(T) + dy_\star^{(1)}(T)) - \int_t^T (bx_\star^{(1)}(s) + dy_\star^{(1)}(s))\ ds - \alpha^{(1)}$$

$$L_2(t) = -\beta(x_\star^{(2)}(T) + dy_\star^{(2)}(T)) - \int_t^T (bx_\star^{(2)}(s) + dy_\star^{(2)}(s))\ ds - \alpha^{(2)}.$$

At $t = t_2$, we have

$$L_1(t_2) = -\beta(bx_\star^{(1)}(T) + dy_\star^{(1)}(T)) - \int_{t_2}^T (bx_\star^{(1)}(s) + dy_\star^{(1)}(s))\ ds - \alpha^{(1)}$$

$$L_2(t_2) = -\beta(bx_\star^{(2)}(T) + dy_\star^{(2)}(T)) - \int_{t_2}^T (bx_\star^{(2)}(s) + dy_\star^{(2)}(s))\ ds - \alpha^{(2)}.$$

We also know that $L_1(t_2) < 0$ and $L_2(t_2) = 0$. So, $L_1(t_2) < L_2(t_2)$. Then we have, due to (1.28),

$$\begin{aligned} & L_1(t_2) + \alpha^{(1)} \geq L_2(t_2) + \alpha^{(2)} \\ \Rightarrow\ & \alpha^{(2)} - \alpha^{(1)} \geq L_2(t_2) - L_1(t_2) > 0 \\ \Rightarrow\ & \alpha^{(1)} > \alpha^{(2)} \end{aligned}$$

So, $\alpha^{(1)} \neq \alpha^{(2)}$ as desired. $\square$

**Lemma 9** If the initial condition, $(x_\star(0), y_\star(0))$, is above the special ray, then the treatment protocol can not begin with a phase of treatment.

**Proof of Lemma 9:** Suppose that the treatment protocol does begin with treatment so that $u(0) = 1$ and $L(0) > 0$. Now, if the optimal trajectory remains above the special ray on the entire interval $[0, T]$ then 1.27 would imply that $L(t) < 0$ on $[0, T]$, contradicting $L(0) > 0$. The other possibility is that the optimal trajectory travels on or below the special ray for some interval of time. This means that there is an interval $I$ adjacent to zero such that the optimal trajectory is above the special ray on $I$. Suppose that the optimal trajectory does reach the special ray at, the earliest time, $t_c$. Then the optimal trajectory is above the special ray on $[0, t_c] \supseteq I$. But then, $L'(t) > 0$ so that $L(t)$ is increasing on $[0, t_c]$. Since $L(0) > 0$, this means that $L(t)$ is positive on $[0, t_c]$ and according to (1.9) we have $x$ decreasing and $y$ increasing. This is a contradiction as this means that the optimal trajectory is moving away from the special ray. □

**Lemma 10.**: If the initial condition, $(x_\star(0), y_\star(0))$, is above the special ray and intersects the special ray at $t = t_c$, then $bx_\star(t) + dy_\star(t) \leq 0$ for all $t \in [t_c, T]$.

**Proof of Lemma 10:** By Lemma 4.1, $L(0) < 0$ so that $u(0) = 0$. Let $t_0$ be the first time that $L(t)$ attains a zero. Let's assume, without loss of generality, to the contrary that the optimal trajectory is above the special ray on an open interval $I \subset [t_c, T]$. Since Lemma 5 tells us that the optimal trajectory can not end above the special ray, we can assume that the optimal trajectory intersects the special ray at the endpoints of the closure of $I$. Now since the optimal trajectory started above the special ray, $L(t)$ must change sign from negative to positive on $I$ in order for the trajectory to depart from the special ray and return. So, there is some time $t_1 \in I$ where $L(t_1) = 0$. It follows that $L(t) < 0$ on some portion of $[0, t_c]$,

then $L(t) > 0$ on some portion of $[t_c, t_1]$ and then $L(t) < 0$ on some interval after $t_1$ and contained in $I$. By the intermediate value theorem, $L(t)$ has two zeros $t_0$ and $t_0'$ with $t_0' \in I$. But by Lemma 3, we know that $L(t) = 0$ on $[t_0, t_0']$ contradicting the positivity of $L(t)$ on any portion of $I$. $\square$

**Lemma 11.** If the initial condition, $(x_\star(0), y_\star(0))$, is above the special ray, then $L(t)$ may only attain zero when the optimal trajectory is coincident with the special ray.

**Proof of Lemma 11:** Suppose that $L(t)$ attains a zero at $t = t_0$ when the optimal trajectory is below the special ray. We may assume that $t_0$ is the first instance of zero after the most recent crossing time $t_c$. Now, in order for the optimal trajectory to cross at $t_c$ and extend below the special ray, we must have $L(t) < 0$ on $[t_c, t_0]$. But then $L'(t) < 0$ so that $L(t)$ is decreasing on $[t_c, t_0]$ contradicting $L(t_0) = 0$. Now, let's assume that $t = t_0$ occurs when the optimal trajectory is above the special ray. We can assume that this occurs before the trajectory crosses the special ray since otherwise $L(t)$ would have to attain a zero on or below the special ray. We know by 1.27 that if the optimal trajectory does not cross the special ray then $L(t) < 0$ on $[0, T]$ and so does not attain a zero. What's left to consider is the situation when the optimal trajectory crosses the special ray. This can only happen when $L(t)$ attains another zero at $t_1 > t_0$ and the optimal trajectory is above the special ray on $[t_0, t_1]$. But then by Lemma 3, $L(t) = 0$ on $[t_0, t_1]$ contradicting that the optimal trajectory is above the special ray. Thus, $L(t)$ may only attain a zero when the optimal trajectory is coincident with the special ray. $\square$

### 1.4.1 Solving the relaxed minimization problem

From the initial conditions of the adjoint system, we have the following equalities,

$$L(T) = -(bx_\star(T) + dy_\star(T)) - \alpha \tag{1.29}$$

$$L'(T) = (bx_\star(T) + dy_\star(T)) \tag{1.30}$$

$$L(t) = -\beta(bx_\star(T) + dy_\star(T)) - \int_t^T (bx_\star(s) + dy_\star(s))\ ds - \alpha \tag{1.31}$$

In what follows we will refer to the four important times:

$$t_0 = \min\{t \in [0,T] : L(t) = 0\}, \tag{1.32}$$

$$t_{on} = \frac{1}{(a+b)-(c+d)} \ln\left(\frac{-dy_0}{bx_0}\right), \tag{1.33}$$

$$\tilde{t}_{on} = \frac{1}{a-c} \ln\left(\frac{-dy_0}{bx_0}\right), \tag{1.34}$$

$$t_{off} = \max\{t \in [0,T] : L(t) = 0\}. \tag{1.35}$$

The quantity $t_0$ represents the time at which a singular regimen begins or treatment is removed. If the optimal trajectory is to intersect the special ray this will occur at the time $t_{on}$ if the initial condition, $(x_\star(0), y_\star(0))$, is below the special ray, or at the time $\tilde{t}_{on}$ if $(x_\star(0), y_\star(0))$ is above the special ray. The times $t_{on}$ and $\tilde{t}_{on}$ also coincide with the time that the singular regimen begins, if it exists. The time $t_{off}$ is the time at which the final phase of the treatment schedule begins and treatment is removed.

**Solution:** We will determine the behavior of the optimal trajectory as $\alpha \in (0, \infty)$ varies. We see that, in the scenario where treatment is never administered (i.e. $u = 0$ on $[0,T]$), if $\alpha \geq -\beta(bx_\star(T) + dy_\star(T))$, then $L(t) < 0$ on $[0,T]$ for all other possible treatment regimens.

This means that the optimal treatment regimen is in fact $u_\star = 0$ on $[0, T]$. With this in mind, we may assume that

$$\alpha < -\beta(bx_\star(T) + dy_\star(T)) \text{ when } u = 0 \text{ on } [0, T]. \tag{1.36}$$

In what follows, we will describe the behavior of the optimal trajectory in the two distinct scenarios, where the optimal trajectory begins either below or above the special ray.

Let us first show that if the initial condition $(x_\star(0), y_\star(0))$ is below the special ray, then the optimal treatment regimen must begin with a phase of treatment under the assumption (1.36). Let's assume to the contrary that we do begin with a phase of no treatment. We would be restricted to $u = 0$ on $[0, T]$ since $L'(t) \leq 0$ on $[0, T]$. But, (1.36) implies that $L(T) > 0$ so that $L(t) > 0$ (i.e. $u = 1$) on some interval adjacent to $T$, contradicting $u = 0$ on $[0, T]$. So the treatment regimen must begin with a phase of treatment under (1.36). For an analogous statement in the case where the optimal trajectory starts above the special ray, please see Lemma 9 of this chapter 1.4.

Below, we consider two separate cases, either $T$ is so small that the optimal trajectory does not reach the special ray or $T$ is large enough that the optimal trajectory has ample time to intersect the special ray.

**Case A:** Suppose that $T$ is so small that the optimal trajectory does not reach the special ray. The solution then depends on the initial condition of the optimal trajectory. We consider the following two subcases.

**Subcase A1.** We will assume, first, that $(x_\star(0), y_\star(0))$ is below the special ray and show that $\alpha$ determines whether or not there is a switching time $t_{off}$ when $u_\star$ switches from $u_\star = 1$

to $u_\star = 0$ for the remaining treatment time. In case A, the optimal trajectory cannot reach the special ray, due to lack of time, so the trajectory remains below the special ray implying that $L'(t) < 0$ on $[0, T]$. So, $L(t)$ is a decreasing function on $[0, T]$. We also know that $L(0) > 0$ since we begin with a phase of treatment. The magnitude of $\alpha$ determines whether we have $L(T) \geq \alpha$ or $L(T) < \alpha$ under the aforementioned assumption that we begin with a phase of treatment. If $L(T) \geq \alpha$, then a switching time does not exist and we must treat on $[0, T]$. The inequality $L(T) < \alpha$ is incompatible with treatment on the entire interval $[0, T]$. This implies that there is a switching time $t_{off}$ where $u$ switches to zero and together with the fact that $L(t)$ is decreasing on $[0, T]$ we have that $t_{off}$ is the unique switching time. This yields the optimal control $u_\star = 1$ on $[0, t_{off})$ and $u_\star = 0$ on $(t_{off}, T]$.

**Subcase A2.** Let's suppose that $(x_\star(0), y_\star(0))$ is above the special ray. By Lemma 9, the optimal treatment must begin with cessation of treatment. Now, there is an analogous time $\tilde{t}_{on}$ that is required for the trajectory to first intersect the special ray in the absence of treatment. Since $T < \tilde{t}_{on}$, then the optimal trajectory remains above the special ray, so 1.27 yields $L(t) < 0$ on $[0, T]$ and the optimal control is $u_\star = 0$ on $[0, T]$.

**Case B:** Suppose that $T$ is large enough that the optimal trajectory intersects the special ray. We will show, again, that $\alpha$ determines the behavior of the optimal trajectory. We will consider the subcases where the initial condition $(x_\star(0), y_\star(0))$ is either below, above, or on the special ray.

**Subcase B1.** If $(x_\star(0), y_\star(0))$ is below the special ray, it suffices to consider the situation where $\alpha$ is small enough (i.e. doesn't satisfy (1.36)) so that the treatment regimen must begin with a phase of treatment. With $T > t_{on}$ and Lemma 6 we conclude that there is a time $t_0$ such that $L(t_0) = 0$. Additionally, by Lemma 6, $L'(t) = -\beta(bx_\star(t) + dy_\star(t)) \leq 0$ on $t \in [0, T]$ which implies that $L(t) \leq 0$ for $t \in [t_0, T]$. We claim that there is a time $t_{off} \geq t_0$ such that $L(t) < 0$ on $(t_{off}, T]$. We will show that $L(t) = 0$ on $[t_0, T]$ is impossible. Suppose

to the contrary that this is the case. This would imply that

$$0 = L'(t) = \beta(bx_\star(t) + dy_\star(t)) \text{ on } [t_0, T],$$

so that the optimal trajectory coincides with the special ray on this interval and, in particular, at $t = T$, contradicting the conclusion of Lemma 7. So, the time $t_{off}$ exists and $L(t) < 0$ on $(t_{off}, T]$. There are two distinct possibilities for the optimal control. One possibility is that $t_0 = t_{off}$ so that we treat until $t = t_{off}$ and make a single switch to no treatment for the remaining time, so that

$$u_\star^{(1)}(t) = \begin{cases} 1, & \text{if } 0 \le t < t_{off}, \\ 0, & \text{if } t_{off} < t \le T. \end{cases}$$

The other possibility is that $t_0 = t_{on}$ and we treat until time $t = t_{on}$, then switch to a singular regimen until $t_{off}$ where we make the final switch to a phase of no treatment yielding the following optimal control

$$u_\star^{(2)}(t) = \begin{cases} 1, & \text{if } 0 \le t < t_{on}, \\ -\dfrac{a-c}{b-d}, & \text{if } t_{on} \le t \le t_{off} \\ 0, & \text{if } t_{off} < t \le T. \end{cases}$$

Each $u_\star^{(i)}$ satisfies (1.21) and is the control corresponding to $\alpha_\star$ where

$$\alpha_\star = -\beta(bx_\star(T) + dy_\star(T)) - \int_{t_{off}}^{T} (bx_\star(s) + dy_\star(s))\, ds.$$

Then $u_\star^{(i)}(t)$ corresponds to the switching function

$$L(t) = -\beta(bx_\star(T) + dy_\star(T)) - \int_t^T (bx_\star(s) + dy_\star(s))\, ds - \alpha_\star$$

By Lemma 8, each $t_{off}$ induces a control for a unique $\alpha_\star$. Therefore, $t_{off}$ induces the optimal control $u_\star^{(i)}(t)$ for $i \in \{1, 2\}$ that minimizes (1.10) with $\alpha = \alpha_\star$.

**Subcase B2.** Now let's assume that $(x_\star(0), y_\star(0))$ begins above the special ray. Recall that $T$ is so large that the optimal trajectory must intersect the special ray. By Lemma 9, we know that the treatment protocol must begin with cessation of treatment. When (1.36) is not satisfied, we claim that induced is a singular regimen followed by cessation of treatment. By Lemmas 10 and 11, we know that $L(t)$ attains its first zero when the optimal trajectory first intersects the special ray at $\tilde{t}_{on}$ and does not return above the special ray. We note that the time $t_{off}$ exists, as proved in Subcase B1. Lemma 8 implies uniqueness of the time $t_{off}$ which then implies that the control must take on a singular regimen and will deviate from the special ray at $t_{off}$. So, we have, the optimal control,

$$u_\star(t) = \begin{cases} 0, & \text{if } 0 \le t < \tilde{t}_{on}, \\ -\dfrac{a-c}{b-d}, & \text{if } \tilde{t}_{on} \le t \le t_{off} \\ 0, & \text{if } t_{off} < t \le T. \end{cases}$$

corresponding to the switching function

$$L(t) = -\beta(bx_\star(T) + dy_\star(T)) - \int_t^T (bx_\star(s) + dy_\star(s))\, ds - \alpha_\star$$

that minimizes (1.10).

**Subcase B3.** Suppose that $(x_\star(0), y_\star(0))$ is on the special ray and that (1.36) is not satisfied.

We know that, by Lemma 7, that the optimal trajectory does not terminate on the special ray. We also know that the time $t_{off}$ exists. By Lemma 5, the optimal trajectory does not go above the special ray. This implies that $t_0 = 0$ and gives us the resulting optimal control

$$u_\star(t) = \begin{cases} -\dfrac{a-c}{b-d}, & \text{if } 0 \le t \le t_{off} \\ 0, & \text{if } t_{off} < t \le T. \end{cases}$$

corresponding to the switching function

$$L(t) = -\beta(bx_\star(T) + dy_\star(T)) - \int_t^T (bx_\star(s) + dy_\star(s))\, ds - \alpha_\star$$

that minimizes (1.10). □

## 1.5 Extensions to other resistant phenotypes

The model of 1.2.2 and many of the methods of Appendix 1.4 are applicable with any set of parameters $a, b, c$, and $d$ such that $ad - bc \neq 0$. In particular, we may consider the relaxed minimization problem with conventional mutants where $0 < \gamma_y^{on} \le \gamma_y^{off} \le \gamma_x^{off}$ so that $c > 0$ and $d < 0$. With these sign changes, we may apply the Pontryagin Maximum Principle to yield the adjoint equations (1.20) and the following familiar switching function

$$L(t) = -\beta(bx_\star(T) + dy_\star(T)) - \int_t^T (bx_\star(\tau) + dy_\star(\tau))\, d\tau - \alpha. \tag{1.37}$$

We can consider the usual special ray $\Pi = \{(x, y)^T : x > 0, y > 0, y = -\frac{b}{d}x\}$. Due to the sign change of $d$, $\Pi \cap \{(x, y) : x > 0, y > 0\} = \emptyset$ and the optimal trajectory, $(x_\star(t), y_\star(t))$, is contained in the first quadrant for $t \in [0, T]$. Let's first consider the case where $\alpha = 0$. The optimal trajectory is above the special ray. So, $y_\star(t) > -\frac{b}{d}x_\star(t)$. Thus, $bx_\star(t) + dy_\star(t) < 0$

for all $t \in [0, T]$. We can make two conclusions. First, a singular regimen is not possible because such a regimen would imply that the optimal trajectory coincides with the special ray. The second implication is that $L(t) > 0$ for all $t \in [0, T]$ and for all $\beta$. Therefore, $u_\star(t) = 1$ for all $t \in [0, T]$.

Now let's consider the situation where $\alpha > 0$. Let

$$\bar{L(T)} = -\beta(bx_\star(T) + dy_\star(T))$$

be the value determined by treating on $[0, T]$. Then for $\alpha \in [0, \bar{L(T)}]$, the optimal control is $u_\star(t) = 1$ on $[0, T]$. As $\alpha$ increases, induced is a switching time $t_{off}$ that begins a phase of non-treatment. This gives the optimal control $u_\star(t) = 1$ on $[0, t_{off}]$ and $u(t) = 0$ on $(t_{off}, T]$.

In Figure 1.2 we plot the optimal control and optimal trajectory for two different values of $\alpha$ that demonstrate the two distinct types of solutions. For the plots in 1.2 we use the following parameter values (see the dot in the red region in figure 1.1):

$$\gamma_x^{\text{on}} = -1,\ \gamma_x^{\text{off}} = 0.25,\ \gamma_y^{\text{on}} = 0.2,\ \gamma_y^{\text{off}} = 0.25,\ \beta = 1,\ T = 30. \tag{1.38}$$

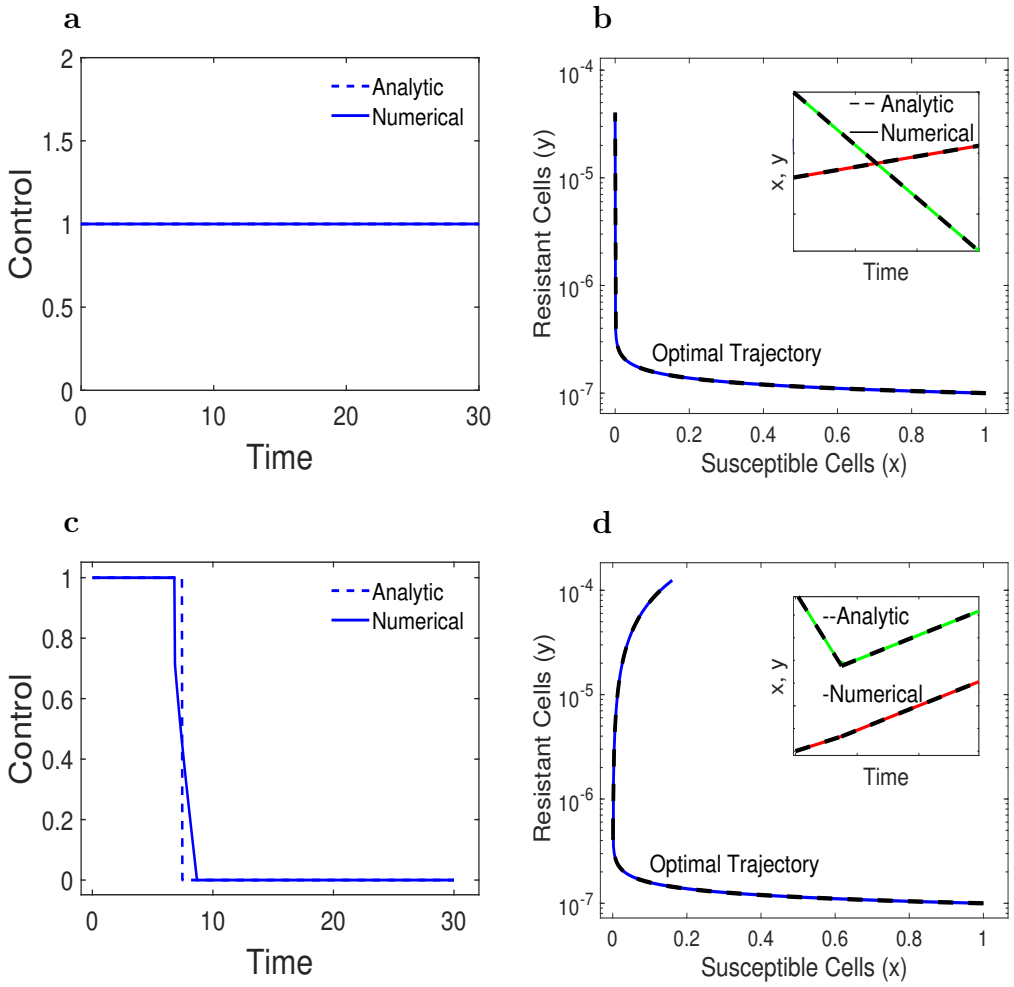


Figure 1.2: Values of $\alpha$ are given by $\alpha = 10^{-7}$ for $(a, b)$ and $\alpha = 1$ for $(c, d)$. Panels (a, c) show the optimal control, and panels (b, d) show the corresponding trajectories $x(t)$ and $y(t)$. The analytical solutions (dashed lines) are overlaid on the numerical solutions (solid lines).

In this figure, are two examples of the control problem representing two qualitatively distinct types of treatment schedules. In the first case, the control is constant with $u_\star = 1$ on the entire treatment interval and, in the other, there is a switching time $t_{off}$ where treatment is removed for the remaining treatment schedule. Corresponding to these two optimal controls are the optimal trajectories and included time series plots. In 1.2a, we observe the optimal control that is constant on $[0, T]$ with corresponding optimal trajectory in 1.2b. The treatment cost weight parameter $\alpha$ is small enough that there is no switch in the treatment

regimen. Increasing to $\alpha = 100$, there is a switching time $t_{off}$ in the control, 1.2c, at which point treatment is no longer applied and both susceptible and resistant cell populations grow.

## 1.6 Conclusions

In this chapter, we formulated the treatment of a heterogeneous melanoma population as an optimal control problem. The tumor was divided into treatment-susceptible and treatment-addicted resistant populations whose net growth rates respond differently to the administration of therapy. Treatment intensity was represented by a bounded control, and the objective functional accounted for cumulative tumor burden, terminal tumor size, and the cost associated with treatment.

Application of Pontryagin's maximum principle reduced the analysis of the optimal control to the behavior of a switching function. The resulting trajectories are organized by a special ray in the state space, which partitions the state space into regions where the optimal treatment protocol has a different initial structure. When an optimal trajectory follows this ray, the corresponding control is singular and represents a constant intermediate treatment intensity. The properties established in this chapter restrict both the order and the number of phases that can occur in an optimal protocol.

For an initial condition below the special ray, as is relevant when the resistant population is initially small, the admissible optimal structures are particularly simple. Depending on the treatment-cost weight and the duration of the optimization horizon, the protocol may consist of full treatment throughout, full treatment followed by permanent cessation, or full treatment followed by a singular phase and then permanent cessation. In particular, once treatment has been discontinued, the necessary conditions do not permit its subsequent reintroduction. Thus, the analysis excludes repeatedly cycling treatment schedules within

the present model.

When the initial condition lies above the special ray, the protocol instead begins without treatment and may subsequently enter a singular phase before returning to no treatment. This distinction demonstrates that the qualitative structure of a treatment protocol depends on the initial composition of the tumor, and not only on its total size. The analysis also shows how the final switching time is related to the treatment-cost weight and establishes the uniqueness properties required to classify the candidate optimal trajectories.

These conclusions describe the qualitative structure imposed by the model, but they do not by themselves determine the switching times for a given parameter set or quantify the resulting tumor burden. Moreover, the occurrence and duration of a singular phase depend on the model parameters, the initial condition, and the weights in the objective functional. The next chapter therefore turns to numerical computation. A sequential quadratic Hamiltonian method is introduced and analyzed, and the theoretically predicted treatment structures are examined across a range of parameter values. Experimental melanoma data are then used to estimate the model parameters and investigate the treatment schedules produced in a biologically informed setting.

# Chapter 2

# Numerical Computation and Data-Informed Application

## 2.1 Motivation and computational objectives

The previous chapter established an analytical framework for determining the possible structure of optimal treatment protocols in a bilinear model of melanoma. Through the Pontryagin maximum principle and an analysis of the switching function, the treatment schedules were shown to consist of a limited sequence of phases: full treatment, possibly an intermediate singular regimen, and eventual treatment cessation. This characterization rules out more complicated switching patterns under the assumptions of the model. It does not, however, directly determine which of these structures occurs for a particular parameter set, when the transitions between phases take place, or how sensitive the resulting protocol is to the relative cost assigned to treatment.

These questions require a numerical treatment of the optimal control problem. Numerical

computation is especially important in the presence of singular control, but in all cases, although the analytical results provide explicit relationships among the state variables, switching times, and model parameters, solving these relationships for a specific biological system may still require numerical methods.

In this chapter, we implement a sequential quadratic Hamiltonian method to compute treatment protocols for the relaxed control problem. The method updates the control by solving a sequence of regularized Hamiltonian optimization problems. We establish convergence properties of the resulting scheme and compare its output with the treatment structures predicted by the analysis of the previous chapter. This comparison serves both to validate the numerical implementation and to illustrate how the qualitative form of the treatment schedule changes with the parameters of the objective functional.

We first investigate the model using representative parameter values. Particular attention is given to the treatment-cost weight, which controls the balance between reducing tumor burden and limiting treatment exposure (this represents the role of side-effects in treatment). By varying this weight and the initial susceptible and resistant populations, we identify parameter regimes corresponding to continuous full treatment, treatment followed by permanent cessation, and protocols containing a singular intermediate-dose phase. We also examine how the final switching time depends on the weight of side-effects and how robust this dependence is to mutation between the two populations.

The final part of the chapter connects the model to experimental observations of vemurafenib-sensitive and vemurafenib-resistant melanoma cells. Growth and decay rates are estimated from published data under treatment and no-treatment conditions. These estimates are then incorporated into the optimal control problem to produce data-informed treatment protocols. The purpose of this application is not to propose an immediately implementable clinical schedule, but to evaluate the mathematical predictions of the model in a biologically relevant parameter regime.

## 2.2 Numerical methods

In this section, we describe the numerical method, in the framework of PMP, for solving the optimal control problem (1.10), with the functional $J(u)$ given in (1.10), and the Hamiltonian function $H(x, y, u, \psi_1, \psi_2)$ given in (1.19), where $\psi_1, \psi_2$ solve the adjoint equations (1.20).

To implement the optimal control numerically, we use a sequential quadratic Hamiltonian (SQH) method. The guiding principle behind these methods is the minimization of the Hamiltonian function associated with a given optimal control problem. The starting point is the consideration of the augmented Hamiltonian given below:

$$\begin{aligned} H_\epsilon(x(t), y(t), u(t), \tilde{u}(t), \psi_1(t), \psi_2(t)) =& H(x(t), y(t), u(t), \psi_1(t), \psi_2(t)) \\ &+\epsilon[u(t) - \tilde{u}(t)]^2. \end{aligned} \tag{2.1}$$

Here, $\epsilon > 0$ denotes a penalization parameter that is adaptively updated at each iteration of the SQH procedure. Specifically, $\epsilon$ is increased when the functional $J$ fails to exhibit a sufficient decrease, and decreased when an adequate reduction in $J$ is achieved. The quantity $\tilde{u}$ represents the previous iterates of the control variable $u$. The quadratic term $\epsilon(u - \tilde{u})^2$ serves to keep the pointwise minimizer of $H_\epsilon$, and thus the update to $u$ is close to its preceding value $\tilde{u}$, particularly when $\epsilon$ is large. During each minimization sweep over the computational grid, the state variables $x, y$ and the adjoint variables $\psi_1$, $\psi_2$ from the previous iteration are employed and are only refreshed after a successful minimization step, ensuring a fast and stable implementation. The SQH algorithm is summarized below: [SQH method for the optimal control $u$]

- Input: initial approx. $u^0$, max. number of iterations $k_{max}$, tolerance $\kappa > 0$, $\epsilon > 0$, $\lambda > 1$, $\eta > 0$, and $\zeta \in (0, 1)$; set $\tau > \kappa$, $k := 0$.

- Compute the solution $(x^0, y^0)$ to the ODE system (1.9) with control $u = u^0$.

- while ($k < k_{max}$ && $\tau > \kappa$ ) do

  1) Compute the solution $(p^k, q^k)$ to the adjoint problem (1.20) with $(x, y) = (x^k, y^k)$.

  2) Determine $u$ such that the following optimization problem is satisfied

$$H_\epsilon\left(x^k, y^k, u(t), u^k, \psi_1^k, \psi_2^k\right) = \min_{v\in[0,1]} H_\epsilon\left(x^k, y^k, v, u^k, \psi_1^k, \psi_2^k\right),$$

  at almost all $t \in [0, T]$.

3) Compute the solution $(x, y)$ to the ODE system (1.9) with control $u$.

4) Compute $\tau := \|u - u^k\|^2_{L^2(0,T)}$.

5) If $J(u) - J\left(u^k\right) > -\eta\,\tau$, then increase $\epsilon$ with $\epsilon = \lambda\,\epsilon$ and go to Step 2, else if $J(u) - J\left(u^k\right) \leq -\eta\,\tau$, then decrease $\epsilon$ with $\epsilon = \zeta\,\epsilon$, set $u^{k+1} = u$, $(x^{k+1}, y^{k+1}) = (x, y)$, compute the adjoint $(\psi_1^{k+1}, \psi_2^{k+1})$ and continue.

6) Set $k := k + 1$.

- end while

We now state and prove some convergence properties of the SQH Algorithm 2.2. Let $u^k$ and $u^{k+1}$ be generated by the SQH scheme, $k \in \mathbb{N}_0$. Then, for the $\epsilon > 0$ chosen by Algorithm 2.2, we have:

$$J(u^{k+1}) - J(u^k) < -\epsilon\|u^{k+1} - u^k\|^2_{L^2(0,T)}.$$

*Proof.* We have the following HP function

$$H(x, y, u, \psi_1, \psi_2) = (a + bu)x\psi_1 + (c + du)y\psi_2 - [x + y + \alpha u].$$

In Step 2 of Algorithm 2.2, we have that for almost all $(x,t) \in \Sigma$ it holds:

$$H_\epsilon\left(x^k, y^k, u^{k+1}, u^k, \psi_1^k, \psi_2^k\right) \leq H_\epsilon\left(x^k, y^k, v, u^k, \psi_1^k, \psi_2^k\right),$$

for all $v \in [0,1]$. Therefore

$$H_\epsilon\left(x^k, y^k, u^{k+1}, u^k, \psi_1^k, \psi_2^k\right) \leq H_\epsilon\left(x^k, y^k, u^k, u^k, \psi_1^k, \psi_2^k\right) = H\left(x^k, y^k, u^k, \psi_1^k, \psi_2^k\right). \quad (2.2)$$

Hence, we obtain the inequality

$$H\left(x^k, y^k, u^{k+1}, \psi_1^k, \psi_2^k\right) + \epsilon\left|u^{k+1} - u^k\right|^2 \leq H\left(x^k, y^k, u^k, \psi_1^k, \psi_2^k\right).$$

We define $\delta x = x^{k+1} - x^k$, $\delta y = y^{k+1} - y^k$ and $\delta u = u^{k+1} - u^k$. We have

$$\begin{aligned}
J(u^{k+1}) - J(u^k) &= \beta(\delta x(T) + \delta y(T)) + \int_0^T (\delta x(s) + \delta y(s))\, dt + \alpha \int_0^T \delta u\, dt \\
&= \int_0^T (a + b\delta u) x^k \psi_1^k + (c + d\delta u) y^k \psi_2^k - [x^k + y^k + \alpha\delta u] \\
&= \int_0^T \left(H\left(x^k, y^k, u^{k+1}, \psi_1^k, \psi_2^k\right) - H\left(x^k, y^k, u^k, \psi_1^k, \psi_2^k\right)\right) dt \\
&= -\epsilon \int_0^T |u^{k+1}(t) - u^k(t)|^2\, dt \\
&= -\epsilon \left\|u^{k+1} - u^k\right\|^2_{L^2(0,T)},
\end{aligned}$$

where the last step follows from (2.2). □

As a consequence, we have the following convergence result.

**Theorem 2.1.** *Let the sequence $u^k$ be generated by the SQH scheme. Then the sequence of cost functional values $J(u^k)$ monotonically decreases with*

$$\lim_{k\to\infty} \left[J(u^{k+1}) - J(u^k)\right] = 0,$$

*and*

$$\lim_{k\to\infty} \|u^{k+1} - u^k\|_{L^2(0,T)} = 0.$$

*Proof.* The first statement follows from the fact that $(J(u^k))_k$ is monotonic decreasing in $\mathbb{R}$ due to Lemma 2.2, and so is Cauchy and convergent. The second statement follows from the fact that

$$\|u^{k+1} - u^k\|^2_{L^2(0,T)} \leq \frac{1}{\epsilon}\left[J(u^k) - J(u^{k+1})\right].$$

Thus,

$$\sum_{k=0}^{N} \|u^{k+1} - u^k\|^2_{L^2(0,T)} \leq \frac{1}{\epsilon}\left[J(u^0) - J(u^{N+1})\right].$$

As $N \to \infty$, the sequence of partial sums of the series $\sum_{k=0}^{\infty} \|u^{k+1} - u^k\|^2_{L^2(0,T)}$ is convergent and, thus, $\lim_{k\to\infty} \|u^{k+1} - u^k\|_{L^2(0,T)} = 0$. □

Theorem 2.1 guarantees that Algorithm 2.2 is well defined for $\kappa > 0$. Hence, there is an iteration number $k_0 \in \mathbb{N}$ such that $\|u^{k_0+1} - u^{k_0}\|_{L^2(0,T)} \leq \kappa$ and therefore the SQH algorithm stops in finitely many steps.

## 2.3 Numerical results

Here we describe the solution to the optimal control problem (1.36) and provide numerical illustrations.

For the plots in this section, we use the following parameter values (they are marked by a

dot in the yellow region of figure 1.1):

$$\gamma_x^{\text{on}} = -1,\ \gamma_x^{\text{off}} = 0.25,\ \gamma_y^{\text{on}} = 0.125,\ \gamma_y^{\text{off}} = -0.0625,\ \beta = 1,\ T = 30 \tag{2.3}$$

Suppose that the reference tumor size is $\bar{X} = 10^{12}$, and assume that treatment in the absence of resistance and under the maximal dose takes 2.5 years. This corresponds to $\Gamma_x^{on} \approx -0.92\text{mon}^{-1}$. Therefore, the choice of $\gamma_x^{off}$ in (2.3) corresponds to $\Gamma_x^{off} \approx 0.23\text{mon}^{-1}$, which means that in the absence of treatment, susceptible cells grow to size $\bar{X}$ in 10 years. Further, the choice of parameters for resistant cells in (2.3) corresponds to $\Gamma_y^{on} \approx 0.115\text{mon}^{-1}$ and $\Gamma_y^{off} \approx 0.0575\text{mon}^{-1}$. The choice of $T = 30$ corresponds to treatment of approximately 33.6 months.

The equation for the special ray for the parameters in (2.3) is given by

$$y = -\frac{b}{d}x = -\frac{20}{3}x,$$

and $u_{\text{sing}} \approx 0.217$.

**The shape of the optimal control.** We start with the assumption that the optimal trajectory begins below the special ray. Then there is a unique time $t_{on}$ when the optimal trajectory first intersects the special ray. Now, let $T > t_{on}$ be the total time of the treatment regimen. When $\alpha = 0$, the optimal control is $u_\star = 1$ on $[0, t_{off}]$ and $u_\star = u_{\text{sing}}$ on $(t_{on}, T]$. The optimal trajectory $(x_\star(t), y_\star(t))$ moves to the special ray on $[0, t_{on}]$ and coincides with the ray on $(t_{on}, T]$. As $\alpha$ increases, a new switching time, $t_{off}$, is induced where we have a phase of non-treatment. This means that $u_\star = 0$ on $(t_{off}, T]$. So, the optimal trajectory departs from the special ray at $t_{off}$ and descends below. Overall, for $\alpha$ small enough, the solution is as follows. The optimal control is $u_\star = 1$ on $[0, t_{on}]$, $u_\star = u_{\text{sing}}$ on $(t_{on}, t_{off}]$, and $u_\star = 0$ on $(t_{off}, T]$. The optimal trajectory approaches the special ray on $[0, t_{on}]$, coincides with the

special ray on $(t_{on}, t_{off}]$, and falls below the special ray on $(t_{off}, T]$. In these scenarios, $\alpha$ is small enough that the optimal control takes on a singular regimen. If $\alpha$ is large enough, then the time $t_{off}$ happens before the control can achieve a singular regimen (singular regimens can only occur on the special ray). So, the optimal solution is as follows. We have $u_\star = 1$ on $[0, t_{off}]$ and $u = 0$ on $(t_{off}, T]$. The optimal trajectory approaches the special ray on $[0, t_{off}]$ and moves away from the special ray on $(t_{off}, T]$. The trajectory may intersect the special ray at a single point in the case that $t_{on} = t_{off}$. The time $t_{off}$ decreases as $\alpha$ increases and if $\alpha$ is large enough then $t_{off} = 0$ and no treatment is applied on all of $[0, T]$.

Let us now consider the case where the optimal trajectory begins above the special ray. For $\alpha = 0$, the optimal trajectory $(x_\star(t), y_\star(t))$ moves to the special ray on $[0, \tilde{t}_{on}]$ and coincides with the ray on $(\tilde{t}_{on}, T]$. This yields the optimal control $u = 0$ on $[0, \tilde{t}_{on}]$ and $u_\star = u_{\text{sing}}$ on $(\tilde{t}_{on}, T]$. As $\alpha$ increases, the switching time $t_{off}$ is induced where we have a final phase of non-treatment so that the optimal trajectory departs from the special ray and descends below. So, for small enough $\alpha$ the solution is as follows. The optimal control is $u_\star = 1$ on $[0, \tilde{t}_{on}]$, $u_\star = u_{\text{sing}}$ on $[\tilde{t}_{on}, t_{off}]$ and $u_\star = 0$ on $(t_{off}, T]$. Now for large enough $\alpha$, $t_{off} = 0$ and the optimal control is $u_\star = 0$ on $[t_{off}, T]$. Finally, let's consider the optimal control when the optimal trajectory begins on the special ray. When $\alpha = 0$, the optimal control is $u_\star = u_{\text{sing}}$ on $[0, T]$. For small enough $\alpha$, the optimal control is $u_\star = u_{\text{sing}}$ on $[0, t_{off}]$ and $u_\star = 0$ on $(t_{off}, T]$. If $\alpha$ is large enough then the optimal control is $u_\star = 0$ on $[0, T]$.

While explicit formulas exist for the times $t_{on}$ and $\tilde{t}_{on}$ (see equations (1.33) and (1.34), respectively), parameter dependence of $t_{off}$ can be defined implicitly. To see this, we substitute the optimal trajectory $(x_\star(t), y_\star(t))$ into equation (1.27) recognizing that $L(t_{off}) = 0$. In particular, in the case when the optimal control takes on the singular control value, we

obtain

$$\begin{aligned} L(t_{off}) = &- \beta b x_0 e^{b(1-u_s)t_{on}} e^{bu_s t_{off}} e^{aT} \\ &- \beta d y_0 e^{d(1-u_s)t_{on}} e^{du_s t_{off}} e^{cT} \\ &- \frac{bx_0}{a} e^{b(1-u_s)t_{on}} e^{(a+bu_s)t_{off}} (e^{a(T-t_{off})} - 1) \\ &- \frac{dy_0}{c} e^{d(1-u_s)t_{on}} e^{(c+du_s)t_{off}} (e^{c(T-t_{off})} - 1) - \alpha = 0. \end{aligned} \tag{2.4}$$

If the optimal control does not take on the singular control value, we have

$$\begin{aligned} L(t_{off}) = &- \beta \left( b x_0 e^{bt_{off}} e^{aT} + d y_0 e^{dt_{off}} e^{cT} \right) \\ &- \frac{bx_0}{a} e^{(a+b)t_{off}} (e^{a(T-t_{off})} - 1) \\ &- \frac{dy_0}{c} e^{(c+d)t_{off}} (e^{c(T-t_{off})} - 1) - \alpha = 0. \end{aligned} \tag{2.5}$$

Importantly, equations (2.4) and (2.5) are valid when the initial condition is below the special ray and when $t_{off} > 0$. If the initial condition is above the special ray, then the relationship between $t_{off}$ and other parameters is given by

$$\begin{aligned} L(t_{off}) = &- \beta \left( b x_0 e^{bu_s(t_{off} - \tilde{t}_{on})} e^{aT} + d y_0 e^{du_s(t_{off} - \tilde{t}_{on})} e^{cT} \right) \\ &- \frac{bx_0}{a} e^{bu_s(t_{off} - \tilde{t}_{on})} (e^{aT} - e^{at_{off}}) \\ &- \frac{dy_0}{c} e^{du_s(t_{off} - \tilde{t}_{on})} (e^{cT} - e^{ct_{off}}) - \alpha = 0. \end{aligned} \tag{2.6}$$

**Dependence on $\alpha$, the relative contribution of side-effects:** We assume once again that the initial condition is above the special ray. Examples of plots of the control and the trajectories for different values of $\alpha$ are given in figures 2.1 and 2.2. The figures compare the analytic and numerical solutions to the relaxed minimization problem for increasing $\alpha$. For each $\alpha$, we corroborate the analytic solution of the control with the SQH method of section 4

and present the induced optimal trajectory and included time series plots of the susceptible and resistant cell populations.

The plots in figures 2.1 and 2.2 use parameter values given in (2.3), while we vary $\alpha$ in equation (1.10). The initial condition in the ODEs 1.9 is taken $x_0 = 1, y_0 = 10^{-7}$, which corresponds to scaling of the tumor size with respect to its initial cell population.

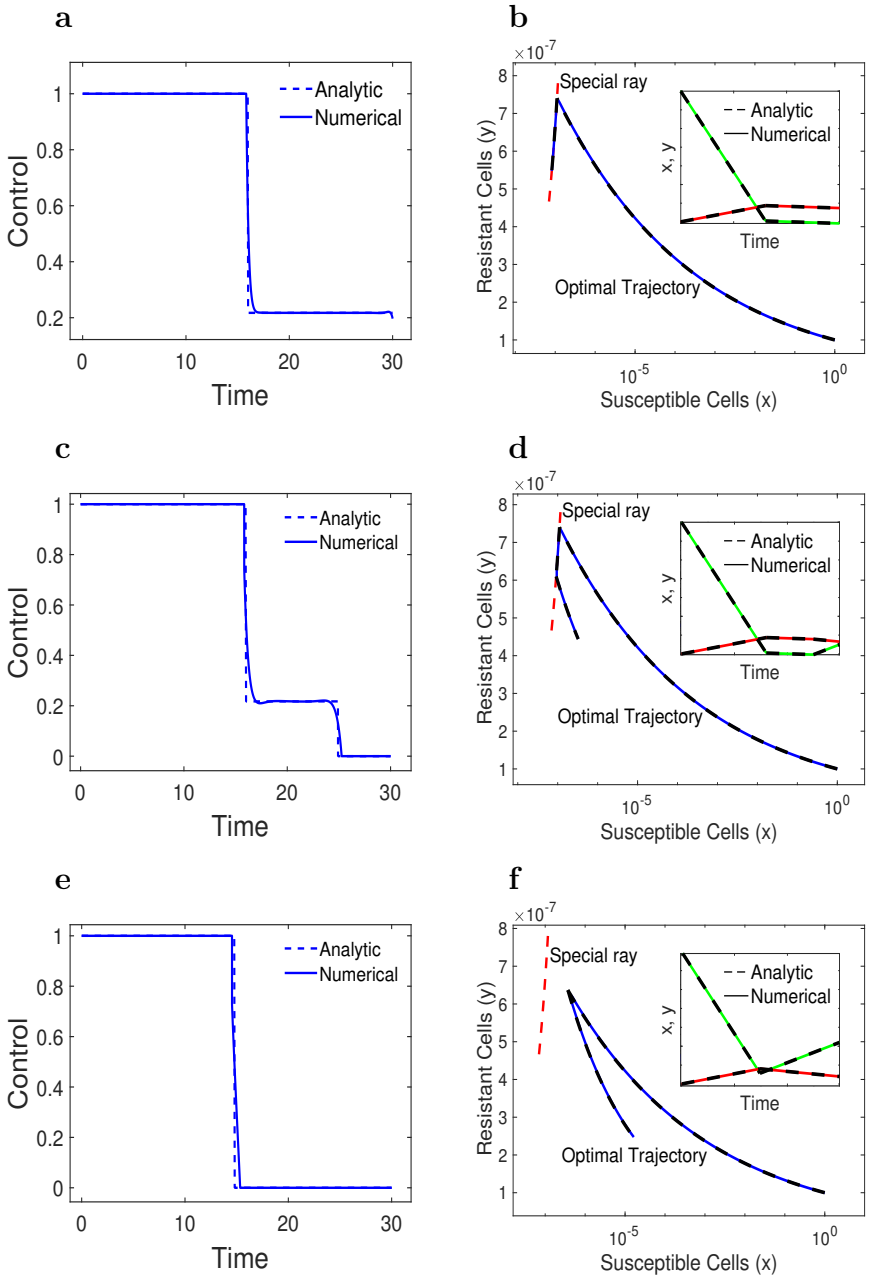


Figure 2.1: Three qualitatively distinct types of treatment schedules for different values of $\alpha$. The values of $\alpha$ are given by $\alpha = 0$ for (a,b), $\alpha = 10^{-6}$ for (c,d), $\alpha = 10^{-4}$ for (e,f). Panels (a,c,e) show the optimal control, and panels (b,d,f) show the corresponding trajectories $x(t)$ and $y(t)$. The analytical solutions (dashed lines) are overlaid on the numerical solutions (solid lines). All units are dimensionless; the initial condition in the ODEs is taken $x(0) = 1, y(0) = 10^{-7}$, and the rest of the parameters are given by (2.3).

The plots of the optimal controls demonstrate that as $\alpha$ increases, the switching time $t_{off}$

emerges and decreases. At the same time, in the plots of the optimal trajectory, the trajectory spends less time coincident with the special ray until eventually no intersection is made. Specifically, in Figure 2.1a, we observe that the control is a bang-bang control that induces the optimal trajectory shown in 2.1b. With the treatment cost weight parameter $\alpha = 0$, the optimal control determines a treatment regimen where treatment is applied on the interval $[0, t_{on}]$ until the optimal trajectory reaches the special ray at which point there is a switch to the singular control on $(t_{on}, T]$. When $\alpha$ increases to $10^{-6}$, sub-figure 2.1c demonstrates the need to remove treatment at the time $t_{off}$ so that the optimal trajectory in 2.1d departs from the special ray and treatment is no longer applied on $(t_{off}, T]$. As $\alpha$ increases further, the time $t_{off}$ decreases further and the singular regimen vanishes from the treatment schedule. Indeed, at $\alpha = 10^{-4}$ we see in 2.1e that the optimal control is a bang-bang control consisting of treatment on $[0, t_{on}]$ and removal of treatment on $(t_{off}, T]$. Figure 2.1f shows that the optimal trajectory must stay below the special ray on $[0, T]$.

In figure 2.2 we continue to increase $\alpha$ and observe its impact on the optimal control and corresponding optimal trajectory. As $\alpha$ increases, we can observe in figure 2.2a that the final switching time $t_{off}$ decreases and in 2.2b, $\alpha$ is so large that $x(T) > x(0)$. Then for $\alpha$ large enough, in our case $\alpha = 2 \cdot 10^4$, we see in 2.2c that $t_{off} = 0$ and no treatment is administered on the interval $[0, T]$. The corresponding optimal trajectory is plotted in 2.2d. Note that the panels of figures 2.1 and 2.2 explore a very large range of value of $\alpha$. This was necessary to visualize a range of behaviors of the optimal trajectory. The reason for this is a very weak dependence of the optimal trajectory on $\alpha$. The value of time $t_{on}$ is independent of $\alpha$ (see equation (1.33)), and $t_{off}$ depends on $\alpha$ logarithmically, as can be seen from equations (2.4) and (2.5). This weak dependence of $\alpha$ means that in order to determine an optimal treatment strategy, a rough estimate of $\alpha$ (treatment cost weight parameter) is sufficient.

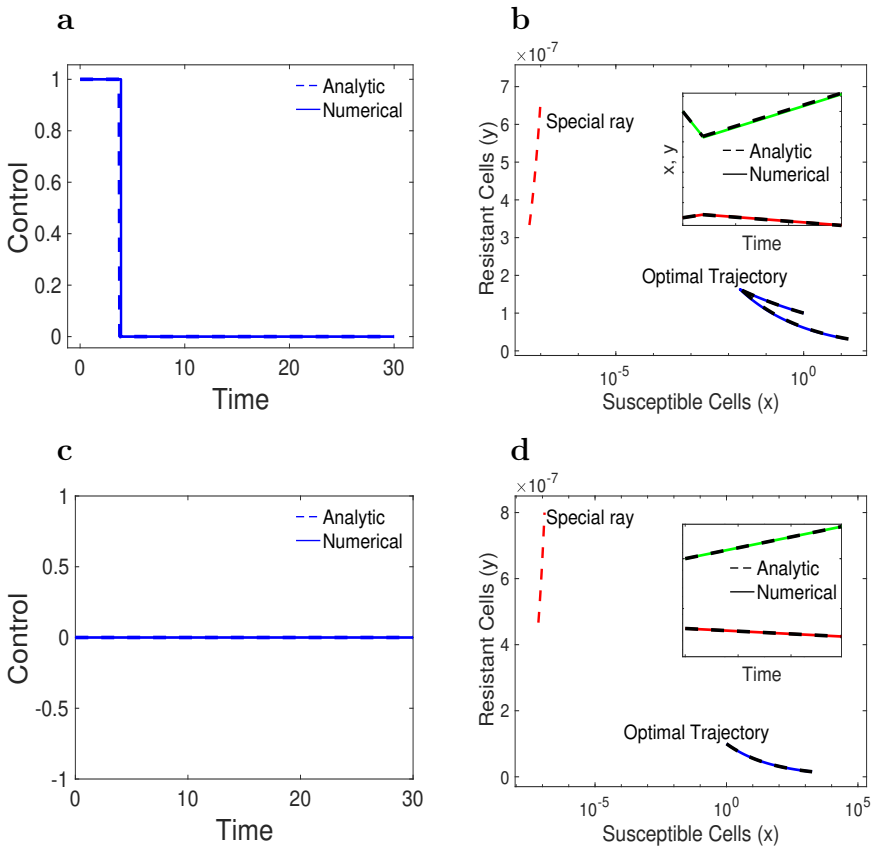


Figure 2.2: Similar to figure 2.1, values of $\alpha$ are given by $\alpha = 100$ for $(a, b)$ and $\alpha = 2 \cdot 10^4$ for $(c, d)$. Panels (a,c) show the optimal control, and panels (b,d) show the corresponding trajectories $x(t)$ and $y(t)$. The analytical solutions (dashed lines) are overlaid on the numerical solutions (solid lines).

**Dependence of the initial conditions:** Next, we have investigated the behavior of the solutions of the optimal control problem, where we start with different initial populations, both below and above the special ray. Figure 2.3 presents a case study that shows some general trends. Here, we sampled $\log_{10} x_0 \in [-9, 1]$ with the step-size of 0.5, and $y_0 \in [10^{-7}, 10^{-6}]$ with the step-size of $5 \times 10^{-8}$. This resulted in $20 \times 20$ sets of initial conditions, which is the grid shown in the central panel of figure 2.3.

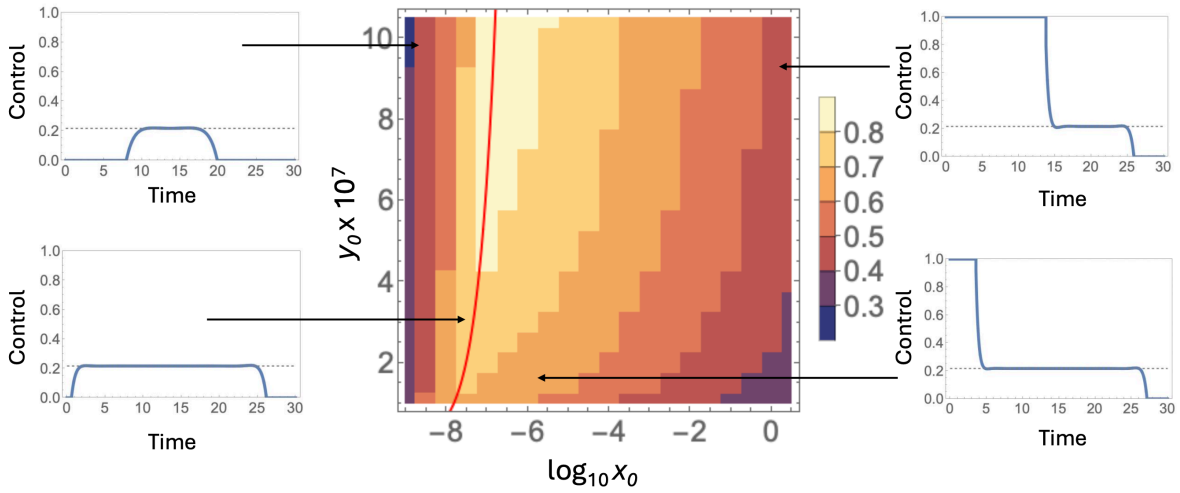


Figure 2.3: Dependence of the optimal control on the initial conditions of the problem. The heat plot in the center shows the fraction of time the optimal control is singular (equation (1.12)), see the color legend on the right. The horizontal axis is $\log_{10} x_0$ and the vertical axis is $y_0 \times 10^7$. The red curve corresponds to the special ray, equation (1.11). Four examples of optimal control solutions corresponding to different initial conditions are shown on the left and on the right, and are connected to their initial condition by an arrow. In these graphs, $u(t)$ is plotted as a function of time, and sing is shown by the dotted line. Here $\alpha = 5.5 \times 10^{-7}$, and the rest of the parameters are given by (2.3).

For each set of the initial conditions, the optimal control was obtained. For the parameters chosen, the shape of the optimal control was of one of the two types:

- For the initial conditions above the special ray (that is, above the red curve in the central panel of figure 2.3), optimal control started at $u = 0$, then switched to singular control (1.12), and finally returned to $u = 0$, see two examples on the left.

- For the initial conditions below the special ray (that is, below the red curve in the central panel of figure 2.3), optimal control started at $u = 1$, then switched to singular control, and finally dropped to $u = 0$, see two examples on the right.

In other words, it always contained a nontrivial portion at $u_{\text{sing}}$. The fraction of time spent at $u_{\text{sing}}$ is shown by the color on the heat plot. We observe that the proportion of time spent on singular control is higher when the initial condition is closer to the special ray. For more examples of the optimal control solutions, see Appendix A.1.

*De novo* **mutations:** So far, we have assumed that $\mu = 0$ in equation (1.4), and studied the co-dynamics of susceptible and resistant cell types under different treatment strategies. Now we include *de novo* mutations to explore how they might affect the solution for the optimal control. Figure 2.4 presents a set of typical outcomes in the presence of mutations. We applied the SQH algorithm to the original model (1.4) with $\mu \geq 0$ and plotted the resulting optimal control in panel (a), as well as the corresponding trajectories in the $(x, y)$ plane in panel (b).

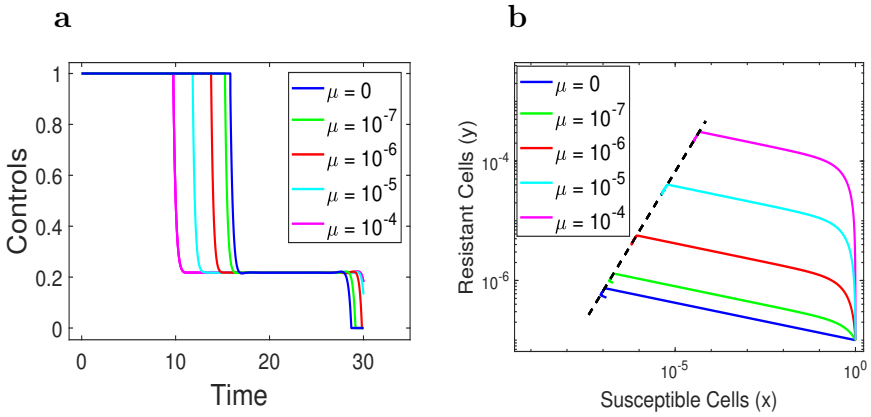


Figure 2.4: Panel (a), shows the optimal controls for $\mu = 0, 10^{-7}, 10^{-6}, 10^{-5}$, and $10^{-4}$. Panel (b) shows the corresponding optimal trajectories along with the special ray plotted in dashed black. The following parameters were used: $l_x^{on} = 1, d_x^{on} = 2, l_x^{off} = 1, d_x^{off} = 0.75$, and the rest of the parameter values are given in (2.3). All units are dimensionless.

We observe that the overall structure of the optimal control for $\mu = 0$ remains the same, that is, the optimal control solution begins at $u = 1$, the switches to $u_{\text{sing}}$, before dropping to $u = 0$. As $\mu$ increases, we can see in figure (2.4a) that the time $t_{off}$ increases while the time $t_{on}$ decreases, resulting in an overall longer time with $u = u_{\text{sing}}$. In figure (2.4b) we observe that as before, under a nonzero mutation rate, the singular control value occurs exactly when the optimal trajectory is coincident with the special ray.

## 2.4 Real data applications

We now demonstrate the versatility of our model for application to real data. For this purpose, we consider the work in [34], where a primary human patient–derived xenograft (PDX) melanoma model harboring the BRAFT1799A mutation, was continuously treated with vemurafenib in immunocompromised mice. The chapter recapitulates the development of drug-resistant melanoma under drug exposure conditions comparable to those observed in patients. Moreover, serial biopsy sampling from a single tumor was used, facilitating the investigation of multiple clonally distinct resistance mechanisms arising within the original tumor. We considered two different cases, one with the susceptible tumor and another with the resistant tumor, and fit our model (1.9) to the extracted data, see Appendix A.2 for details.

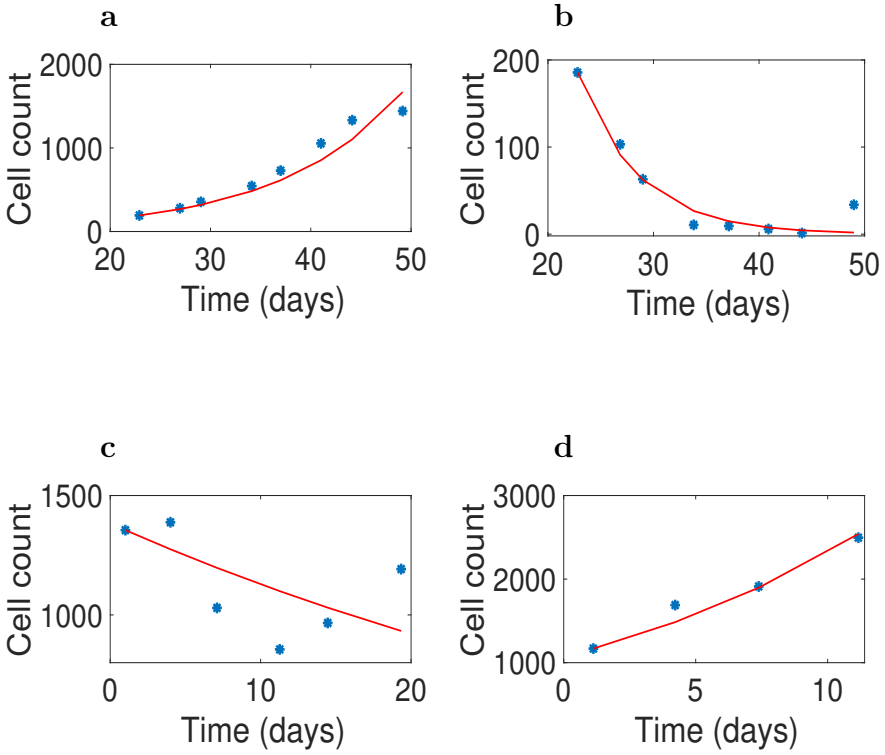


Figure 2.5: Plots for the fits of the susceptible and resistant tumors in absence and presence of treatment. The first row corresponds to the susceptible tumor plots while the second row corresponds to the resistant tumor plots. The data from [34] are shown by dots and the best fits by lines. (a) Susceptible tumor cells in the absence of treatment (Figure 1(a) from [34]). (b) Susceptible tumor cells under maximum Vemurafenib treatment (45 mg/kg) (Figure 1(a) from [34]). (c) Resistant tumor cells in the absence of treatment (Figure 4(b) from [34]). (d) Resistant tumor cells under maximum Vemurafenib treatment (45 mg/kg) (Figure 4(b) from [34]).

For our first test case, we consider the Figure 1(a) from [34], where the mice bearing subcutaneous HMEX1906 tumours were dosed with vehicle and with 45mg/kg vemurafenib twice daily. This corresponds to evolution of the susceptible population dynamics as given by the first equation of our model (1.9), given by $x' = (a + bu)x$. The resulting fits are given in the first row of Figure 2.5. Figure 2.5a shows the fit of our susceptible population model in the absence of treatment, while Figure 2.5b shows the fit in presence of the maximum treatment. From the first fit, we obtain the value of the parameter $a$ since $u = 0$. From the second fit,

we obtain the value of $a + b$, setting $u = 1$ to represent the maximum dose, which gives us the parameter $b$. The corresponding parameter values of $a, b$ are provided in Table 2.1.

For our next test case, we considered Figure 4(b) from [34], where the Vemurafenib-resistant tumours were implanted in the mice and were dosed with vehicle and with 45mg/kg vemurafenib twice daily. Due to the absence of any susceptible tumor in this experiment, this corresponds to the second equation of our model (1.9), given by $y' = (c+d)y$. The resulting fits are given in Figure 2.5. Figure 2.5c shows the fit of our resistant population model in the absence of treatment, while Figure 2.5d shows the fit in presence of the maximum treatment. From the first fit, we obtain the value of the parameter $c$ since $u = 0$. From the second fit, we obtain the value of $c + d$, setting $u = 1$ to represent the maximum dose, which gives us the parameter $d$. The corresponding parameter values of $c, d$ are provided in Table 2.1.

Table 2.1: Parameter values for the real data test case.

| Parameter | Value | units |
|---|---|---|
| $a$ | 0.0822 | $\text{day}^{-1}$ |
| $b$ | -0.2583 | $\text{day}^{-1}$ |
| $c$ | -0.022 | $\text{day}^{-1}$ |
| $d$ | 0.097 | $\text{day}^{-1}$ |

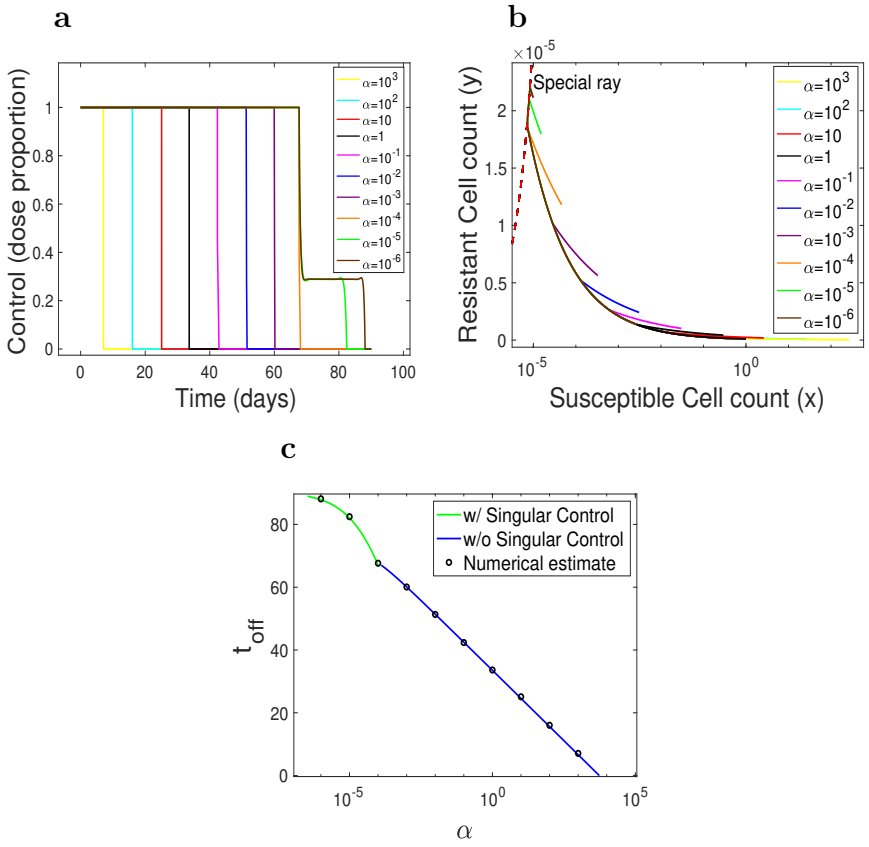


Figure 2.6: Real data application. Panel (a) shows plots of the optimal control as $\alpha$ is varied from $10^{-6}$ to $10^{3}$ and panel (b) shows the associated optimal trajectories. In panel (c), the plotted circles indicate the $(\alpha, t_{off})$ pairs corresponding to the plots in panel (a). The green and blue curves are produced by evaluating equations (2.4) and (2.5), respectively, for $t_{off}$ in the range of $10^{-6}$ to $10^{3}$. These plots were produced using the parameters in Table 2.1. The initial condition was $(x_\star(0), y_\star(0)) = (1, 10^{-7})$.

In figure 2.6, we present the optimal solution to (1.10) for the parameter values in Table 2.1 where $\alpha = 10^n$ for each integer $-6 \leq n \leq 3$. In figure (2.6a) we can see that for smaller values of $\alpha$, namely $\alpha = 10^{-6}$ and $10^{-5}$, the treatment demands that the maximum dose (corresponding to $90mg/kg$ per day) is applied, followed by a switch to an intermediate dose, $u_{\text{sing}} = -\frac{a-c}{b-d}$ (this corresponds to approximately $26mg/kg$ per day), and ending with cessation of treatment for the remaining treatment time. For larger values of $\alpha$, the maximum dose is applied for some time followed by a complete cessation of treatment for the remaining treatment time. In figure 2.6b, the optimal trajectories corresponding to $\alpha = 10^{-6}$ and $10^{-5}$

approach the special ray on the maximum treatment dose, then run coincident with the special ray at the intermediate dose value, followed by complete cessation of treatment and departure from the special ray. For the values of $\alpha$ between and including $10^{-4}$ and 10, the optimal trajectory approaches the special ray at the maximum dose and then, without intersecting the ray, is redirected away from the ray. For the even larger values of $\alpha$, $10^2$ and $10^3$, the maximum does is applied for less than 20 days resulting in optimal trajectories that are redirected away from the special ray such that the final susceptible cell count surpasses its starting number.

Finally, in figure 2.6c, we present a systematic study of the dependence on the time $t_{off}$ (the time when the optimal control becomes zero) on $\alpha$. The dots in the figure correspond to the numerical solution of the optimization problem, while the lines are solutions of equations (2.4) and (2.5). This figure illustrates the very weak (logarithmic) dependence of the switching time on the value of $\alpha$.

## 2.5 Discussion

The results of this study provide a rigorous numerical characterization of optimal treatment schedules in a bilinear control model of melanoma containing treatment-addicted resistant mutants and a penalty term representing cumulative drug toxicity or side-effects.

These findings emerge from a numerical sequential quadratic Hamiltonian method, whose convergence confirms the analytic structure and efficiently computes solutions even in the presence of non-differentiable terms such as the absolute-value penalty for drug usage. Notably, the numerical method provides stable verification of the analytic regime-structure over a wide range of $\alpha$.

Connecting these findings with paper [1], the randomized phase II clinical trial that com-

pared continuous vs. intermittent BRAF/MEK inhibition, there is striking conceptual agreement. That trial showed that continuous therapy outperformed fixed intermittent regimens, contradicting earlier mouse experiments that had motivated cycling schedules. The optimal-control solutions here provide a mechanistic/mathematical explanation for why purely cyclic (on–off–on–off...) treatment patterns are not optimal when the goal is to minimize tumor burden and treatment cost/side-effects over a fixed horizon. In our model, once treatment is turned off, any return to therapeutic dosing would push the trajectory upward relative to the special ray and would violate the necessary conditions of optimality; the switching function cannot support such oscillations. Instead, the best strategy always begins with treatment, possibly transitions to a reduced steady level (the singular control), and then ends with cessation, never returning to treatment, and never showing a periodic pattern. This structure mirrors the clinical finding that preplanned cycling provides no benefit: if treatment is helping, it should continue; if toxicity makes reduction necessary, the optimal reduction resembles a singular steady regimen rather than a switch to zero followed by future resumption. Thus, although earlier mouse studies implied a benefit for cycling, the analytic structure of the optimal control problem and the clinical data both suggest that monotone, not cyclic, treatment strategies are superior when side effects and realistic competitive dynamics of addicted mutants are considered.

It is also useful to contrast the singular-control strategy identified here with the adaptive therapy (AT) paradigm (see e.g. [11, 38]). Classical AT protocols typically employ threshold-based treatment rules that alternate between periods of treatment and treatment withdrawal, generating repeated on–off cycles. By contrast, the optimal controls derived in this study arise from minimizing an explicit objective functional and often contain a singular arc corresponding to a sustained intermediate dose. In this sense, the singular-control phase may be viewed as a form of adaptive dose modulation rather than adaptive on–off cycling. Within the present model, the resulting treatment schedules are globally optimal among all admissible controls, including cyclic schedules, and therefore repeated treatment re-challenge does

not improve the objective functional. This conclusion, however, is tied to the assumptions of the underlying bilinear model and should not be interpreted as a general argument against adaptive therapy in clinical practice.

The model studied here is deliberately minimal: it lacks competition terms between sensitive and resistant cells, tumor–microenvironment interactions, immune dynamics, or spatial aspects of tumor growth. In particular, the sensitive and resistant populations do not interact directly, whereas ecological competition is often regarded as a key mechanism underlying adaptive therapy strategies. Consequently, the absence of optimal cycling strategies in our analysis should be interpreted as a property of the present model and objective functional rather than a general statement about clinical treatment design. Mechanisms omitted here may create selective advantages for intermittent dosing or drug re-challenge schedules. For example, immune-mediated effects can introduce delayed responses and feedbacks, while cumulative or non-linear toxicity costs may favor treatment holidays that are not captured by the current formulation. Prior work, such as [18, 31], has shown that competition and phenotypic switching profoundly influence optimal schedules, and incorporating these features into the current optimal-control framework could alter the switching structure and potentially favor treatment schedules more closely resembling adaptive therapy. Further, spatial dynamics has been shown to be an important factor in optimizing adaptive therapy [33]. Growth of solid tumors in 3D differs from the exponential or logistic growth laws predicted by simple ODE models, and (resistant) mutant accumulation follows scaling laws that differ in spatial systems compared to well-mixed, mean-field dynamics [37, 19]. Factors such as spatial heterogeneity in drug concentration [10] or tumor microenvironment [28] have also been identified as important determinants of treatment strategy design. There is a growing number of papers developing mathematical techniques to study optimal control in spatial tumor growth and treatment models [4, 23, 8, 20, 5].

Our numerical tools are versatile and can handle a wide range of model modifications. The

most immediate next step is to include cell competition, which is organically connected with the geometry of tumor growth. As discussed, e.g., in [25], solid tumors engage in growth dynamics that may be very different from the exponential (like in the present chapter) or logistic (as in many other studies) growth. “Surface growth" dynamics have been studied in [37], where the population expands as a power law of time. Mutant accumulation in such spatially growing colonies obeys scaling laws that depend on whether the mutants are disadvantageous, neutral, or advantageous [19]. The competition dynamics in such spatial situations may be different from those described by generalized Lotka-Volterra equations, and could affect the solution of optimization problem in non-trivial ways. This is subject to ongoing research.

# Chapter 3

# Spatial Stochastic Modeling of Melanoma

## 3.1 Motivation and objectives

The melanoma models developed in the previous chapters describe susceptible and resistant cells through their total population sizes. This population-level description makes the optimal control problem analytically and numerically tractable, but it does not represent the physical arrangement of cells within a tumor. In particular, the two populations do not compete explicitly for space: the growth rate of one population is unaffected by the locations or abundance of cells belonging to the other population.

Spatial competition can change the response of a heterogeneous tumor to treatment. When treatment kills susceptible cells, the resulting space may become available to resistant cells. Treatment can therefore affect the resistant population indirectly, in addition to the direct effect on its intrinsic division and death rates. The magnitude of this effect depends on the spatial organization of the tumor, including the locations of resistant cells relative to susceptible cells and to the boundary of the growing population.

To investigate these effects, we develop a two-dimensional spatial stochastic model containing treatment-susceptible and treatment-addicted resistant melanoma cells. Individual cells occupy sites on a lattice and undergo division and death events whose rates depend on whether treatment is present. Cell division also depends on the availability of a neighboring empty site, introducing direct competition for physical space. The models in this chapter and in the next chapter use the same event-driven lattice framework (to be adapted to biofilm dynamics in Chapter 4).

The current chapter has two objectives. The first is to determine how spatial competition influences the response of susceptible and resistant melanoma populations to treatment. In particular, we examine whether stronger suppression of the susceptible population can indirectly facilitate resistant-cell expansion by releasing spatial constraints. The second objective is to determine whether the aggregate output of the spatial model can be represented by a lower-dimensional system of ordinary differential equations.

To address the second objective, we generate time-series data under a range of intermittent treatment schedules and use sparse identification of nonlinear dynamics to identify effective growth laws. The resulting differential equations are not intended to replace the agent-based model. Rather, they provide a reduced description that incorporates some of the population-level effects of spatial crowding while remaining more amenable to analysis and optimization.

## 3.2 Spatial Stochastic Model

While melanoma lives in three dimensions, the tumor generally is flatter in one dimension than in the other two. For this reason, in our agent-based model, we make a simplification by modeling melanoma as a two dimensional entity. We are interested in exploring the

co-dynamics between the susceptible and resistant cells located together at the cancer site. These two types of cancer cells comprise the two types of agents in the model to be simulated. Using assumptions of the previous two chapters, here the resistant cells thrive in the presence of treatment but decay in the absence of treatment and we say that these types of cancer cells are “addicted" to treatment. The susceptible cells exhibit the opposite behavior in that they experience net growth in the absence of treatment and decay in the presence of treatment.

We model the melanoma tumor by using a square 2D grid of $N \times N$ spots. Each spot can be occupied by a susceptible cell ($x$) or a resistant cell ($y$). The grid is initialized by filling the inner portion of an annulus with resistant cells and filling the outer ring of the annulus with susceptible cells. The dynamics proceed as a sequence of updates that include cell divisions and death. Cell divisions happen by picking a random neighbor out of the 4 neighbors and putting a cell of the same type in that position so long as that position is unoccupied. If the position is occupied, then cell division does not occur. Cell death occurs according to the phase of treatment or no-treatment. If the system is in a phase of treatment, growth and death rate parameters are chosen in such a way that the net growth rate, $\gamma_x^{on}$, of the susceptible cells is negative while the net growth rate of the resistant cells, $\gamma_y^{on}$, is positive. On the other hand, if the system is in a phase of no-treatment, the net growth rate, $\gamma_x^{off}$, of the susceptible cells is positive while the net growth rate of the, $\gamma_y^{off}$, resistant cells is negative.

Numerically, this agent-based model is implemented as a continuous-time spatial stochastic process on a two-dimensional lattice. Cell division and death are represented as discrete events whose occurrence times are determined by their state-dependent rates. The simulations are implemented using the stochastic simulation algorithm and the Next Reaction Method also described in Section 4.3. The same event-driven lattice framework is used in Chapter 4 to model biofilm dynamics, with the melanoma cell types and treatment-dependent events replaced by bacterial phenotypes, flow-induced removal, and phenotypic conversion.

Running simulations of our agent-based model allows us to generate spatial and temporal data that gives a description of tumor growth under the effects of intermittent treatment. Changes in the treatment regimen, i.e. modifying the treatment and no-treatment times, results in changes to the overall tumor growth and decay and eventual outcomes for both the susceptible and resistant cells. This is a spatial generalization of the assumptions used in the previous chapters. The natural question is optimizing the decay of the tumor, and both types of cancer cells, via an optimal choice of treatment scheduling.

## 3.3 Spatial Competition During Treatment

We first isolate the effect of susceptible-cell removal on resistant-cell expansion. A number of simulations were performed using the same initial condition and the same resistant-cell parameters, while the net growth rate of susceptible cells during treatment was varied. Here we present results for the following values of $\gamma_x$:

$$\gamma_x^{\text{on}} \in \{-0.1, -0.2, -0.3, -0.4, -0.5, -0.6\}.$$

Because the resistant-cell parameters are held fixed, differences in the resistant population can be attributed to changes in the surrounding susceptible population and the resulting availability of space. Figure 3.1 presents typical examples of the spatial picture of the simulated cancer growth/decay for the two cases $\gamma_x^{on} = 0.3$ and $\gamma_x^{on} = 0.6$.

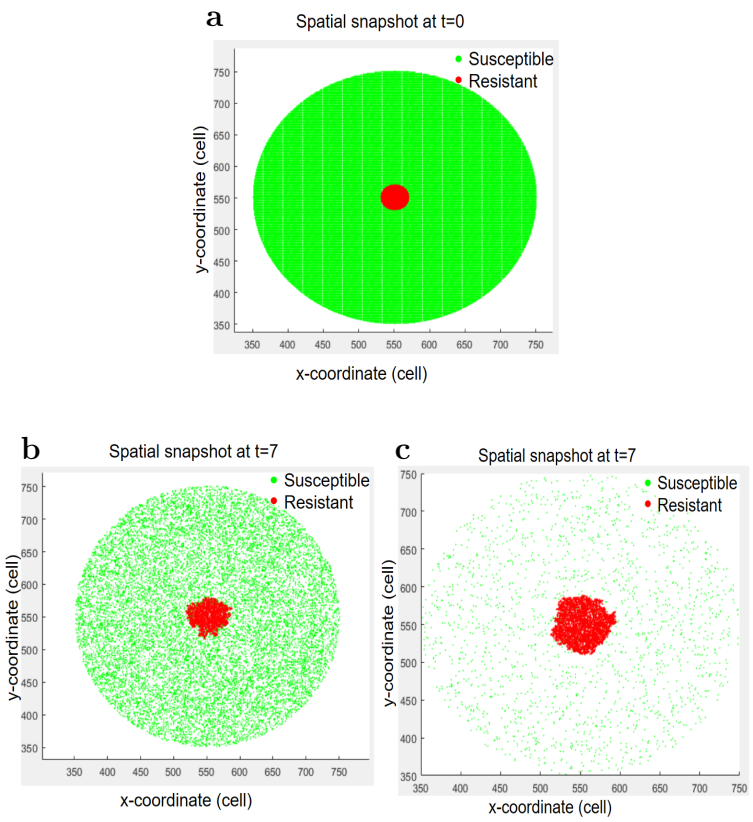


Figure 3.1: Panel (a) shows the spatial initial condition used in different simulations. Panel (b) shows the result of simulating treatment for 7 units of time with susceptible cell net growth rate $\gamma_x^{on} = -0.3$. Panel (c) shows the result of simulating treatment for 7 units of time with susceptible cell net growth rate $\gamma_x^{on} = -0.6$.

We observe that although the net growth rate, $\gamma_y^{on}$, of the resistant cells is unchanged between simulations, yet the colony of resistant cells is visibly larger in figure 3.1c compared to figure 3.1b. This observation suggests that reducing the net growth rate of susceptible cells from $\gamma_x^{on} = -0.3$ to $\gamma_x^{on} = -0.6$ frees up physical space around the resistant cell colony, space that is then available for cell division. It turns out that, in fact, the resistant cell colony is larger as evidenced by the following time series data of the number of susceptible and resistant cells.

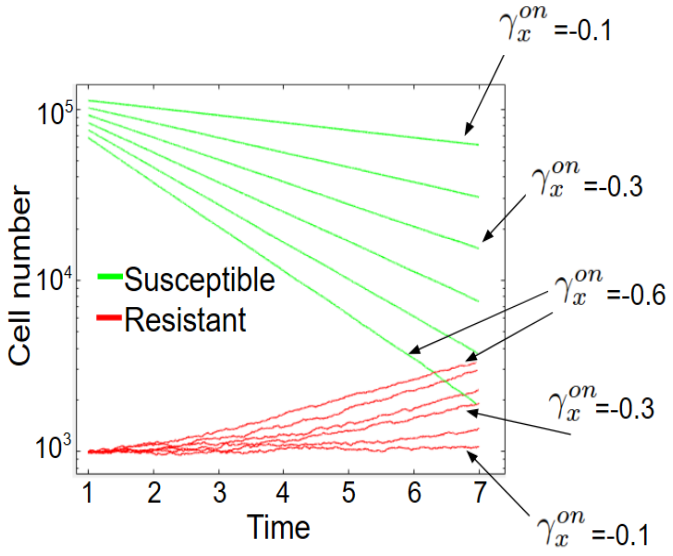


Figure 3.2: The susceptible and resistant cell number is plotted against time in the 6 separate simulations of decreasing susceptible cell net growth rate.

In figure 3.2, we can see that as the susceptible cell net growth rate $\gamma_x^{on}$ decreases, the overall susceptible cell number at time $t = 7$ decreases. At the same time, the number of resistant cells, at the end of the simulation, increases. This supports the idea that the spatial element contributes to the increased growth of resistant cells. Figures 3.1 and 3.2 aim to demonstrate some of the complexity of the system and to draw a connection between the spatial component and the growth and decay of the cells involved.

Increasing the magnitude of the susceptible-cell decay rate produces the expected reduction in susceptible-cell abundance. At the same time, however, the resistant population becomes larger, even though its intrinsic division and death rates are unchanged. This increase is a spatial effect. Faster removal of susceptible cells creates empty sites near the resistant colony, increasing the number of successful resistant-cell division events.

The result illustrates a form of competitive release. Treatment acts directly against susceptible cells but indirectly benefits resistant cells by reducing competition for space. This

coupling is absent from the bilinear ODE model of Chapters 1 and 2, in which the two populations grow independently once the treatment schedule is specified. In the spatial model, the effective growth of either population depends not only on its intrinsic rates but also on the abundance and spatial arrangement of the competing population.

## 3.4 Aggregate Dynamics Under Intermittent Treatment

We next consider simulations containing repeated treatment and no-treatment phases. These simulations produce aggregate time series for the susceptible and resistant populations while retaining the local spatial interactions of the agent-based model. Figure 3.3 shows a representative schedule consisting of ten time units of treatment followed by ten time units without treatment.

The two populations respond oppositely to each change in treatment status. Susceptible cells decline during treatment and expand after treatment is withdrawn, while resistant cells expand during treatment and decline in its absence. The trajectories are determined not only by the corresponding net growth rates. The available space and the spatial configuration inherited from the preceding phase also affect the aggregate response. Consequently, the same phenotype can exhibit different effective growth behavior in different phases of a treatment schedule.

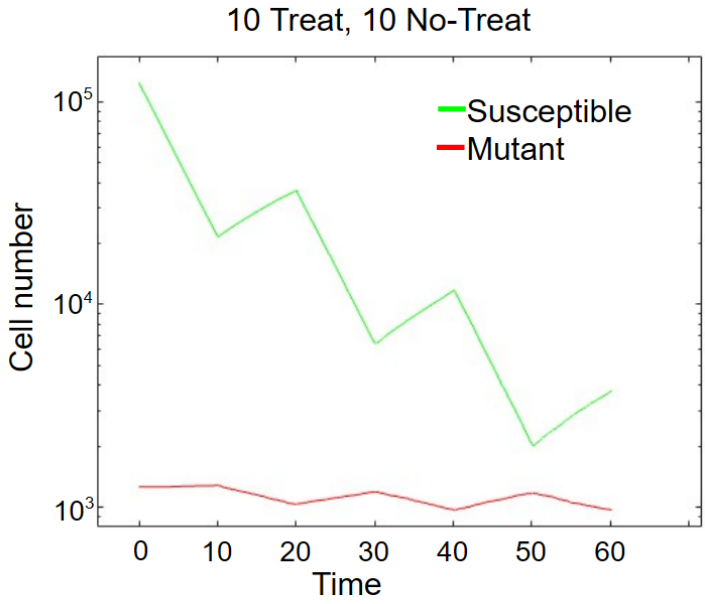


Figure 3.3: This plot presents the susceptible cell and resistant cell number for a simulated treatment consisting of 10 units of treat time and 10 units of rest time.

## 3.5 Identification of Effective Growth Laws

We want to implement a series of treatment and no-treatment phases within a single run of the agent-based model with the goal of fitting 4 ODEs. We seek one ODE that describe the decay of susceptible cells in treatment, another ODE that describes the decay of resistant cells in no-treatment over each interval of respective treat/no-treat. Further, we seek one ODE which describes the growth of susceptible cells in no-treatment and a final ODE which describes the growth of resistant cells in treatment for each phase of treat/no-treat.

In order to find a proper fit to the data, we first tried ODEs of the form $\dot{u} = \gamma u$. It turns out that we can find $\gamma_x^{on}$ to fit $\dot{x} = \gamma_x^{on} x$, $x(t_i) = x_i$ to each interval of treatment, where the initial condition is the amount of susceptible cells, $x_i$, at the start of the treatment phase at time $t_i$ and $i$ ranges from 1 to the total number of no-treatment phases. Similarly, we can find $\gamma_y^{off}$ to fit $\dot{y} = \gamma_y^{on} y$, $y(r_i) = y_i$ to each interval of treatment, where the initial condition

is the amount of resistant cells, $y_i$, at the start of the no-treatment phase at time $r_i$. Now, the difficulty is that linear ODEs do not give a good fit to the spatial data for susceptible and resistant cells during their phases of net growth in no-treatment and treatment, respectively.

To help us search for a better fit we employed a method of sparse identification of Nonlinear Dynamical systems (SINDy) via the Python package PySINDy. We took the time-series data of susceptible and resistant cell counts as input and ran the PySINDy algorithm to output the fitted dynamics as well as plots of the fit, data, and residual percentage error. Initially, we attempted to fit an ODE of the form

$$U' = AU - BU^2 + CU^3 + D\Gamma\text{"}0221AU,\ U(t_0) = U_0 \tag{3.1}$$

to capture, primarily, logistic growth and surface growth. We found an acceptable fit, that is, a set of $\{A, B, C, D\}$ for susceptible cells during no-treatment phases, and a separate set of coefficients for resistant cells during treatment phases.

Next, we wanted to investigate if this model can be reduced to provide a simplified description. We removed pairs of terms from equation (3.1) thus leaving only two terms, to determine what minimal model gives the smallest error. With some experimentation we found that we could simplify the form of the ODE to one of only logistic growth to fit the growth of both susceptible and resistant cells. In what follows, we will present the fits obtained from a logistic growth model,

$$U' = AU - BU^2,\ U(t_0) = U_0,\ A, B > 0.$$

The coefficient $A$ represents the effective low-density growth rate, whereas $B$ summarizes the reduction in successful population growth due to spatial crowding. These are effective parameters of the aggregate model; they need not coincide with the intrinsic division and

death rates assigned to individual cells.

Let us give more detail on how a fit was produced. We will focus on a treatment schedule which treats for $n$ units of time and no-treatment for $4n$ units of time. We apply the SINDy algorithm to each of the following treat-no-treat schedules, 1-4, 3-12, 5-20, 8-32, and 10-40 to produce four sets of coefficients $A$ and $B$. We then take two averages, one from the sets of $A$ values, $A_{ave}$, and the other from the set of $B$ values, $B_{ave}$, to form a new model fit

$$x' = A_{ave}x - B_{ave}x^2,\ x(t_0) = x_0,$$

where $x$ is the susceptible cell count. This provides a single fit for the $n$ treat, $4n$ no-treat schedule. Similarly, we have found a fit for resistant cell growth during treat phases in the $n$ treat, $4n$ no-treat schedule.

In our simulation, we used a $1100 \times 1100$ grid and initialized an annulus with 5,025 resistant cells surrounded by 277,672 susceptible cells. During no-treatment, the susceptible cells have net growth rate $\gamma_x^{off} = 0.077$ and in treatment the resistant cells have net growth rate $\gamma_y^{on} = 0.1744$. Figure 3.4 demonstrates model fits for susceptible cells with logistic growth model

$$x' = 0.0562x - 0.0392x^2,\ x(t_i) = x_i,$$

where $i$ ranges from 1 to the total number of intervals of no-treatment in the simulation.

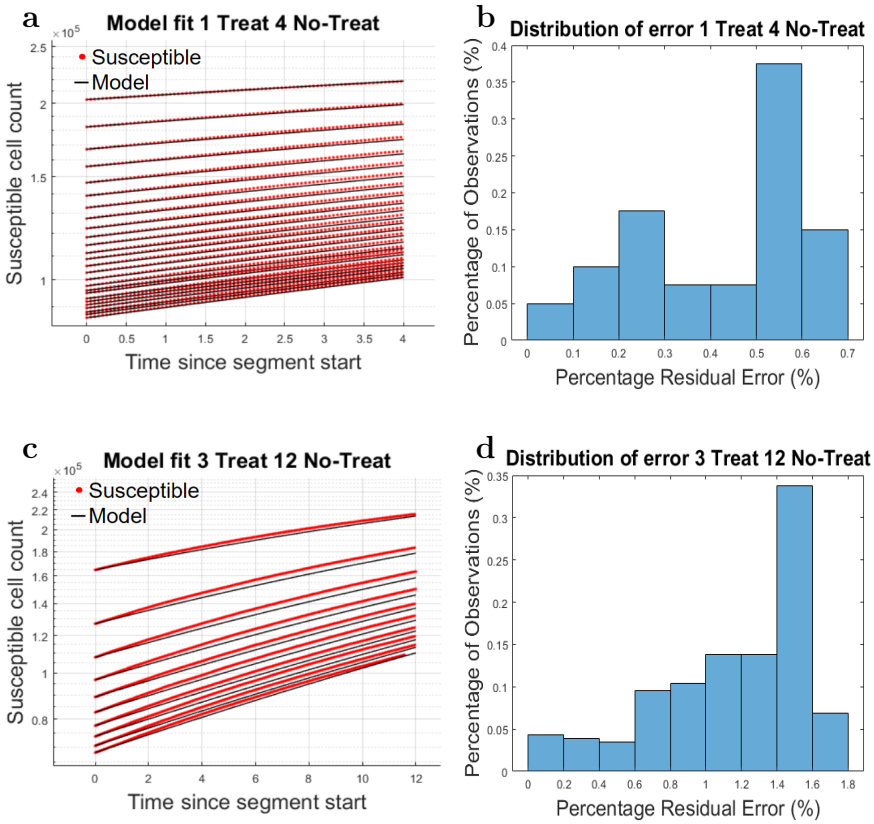


Figure 3.4: Panel a shows the susceptible cell count from the simulation with model fitted to data from the 4 time unit no-treatment phases. Panel b shows the distribution of error produced by the fit. Panels c and d present the same type of information for the treatment schedule 3 treat, 12 no-treat. All data shown is from the no-treatment phases of the treatment regimen.

Figure 3.5 demonstrates the model fitted to the resistant cell data during the treatment phases of the treatment regimens 1-4, 3-12, 5-20, 8-32, 10-40. As before the SINDy algorithm was used to find dynamics for the resistant cell data in each treatment regimen and the coefficients in the dynamics were averaged to get

$$y' = 0.0714y - 0.0586y^2, \; y(t_i) = y_i,$$

where $i$ ranges from 1 to the total number of intervals of treatment and $y$ is the resistant cell count.

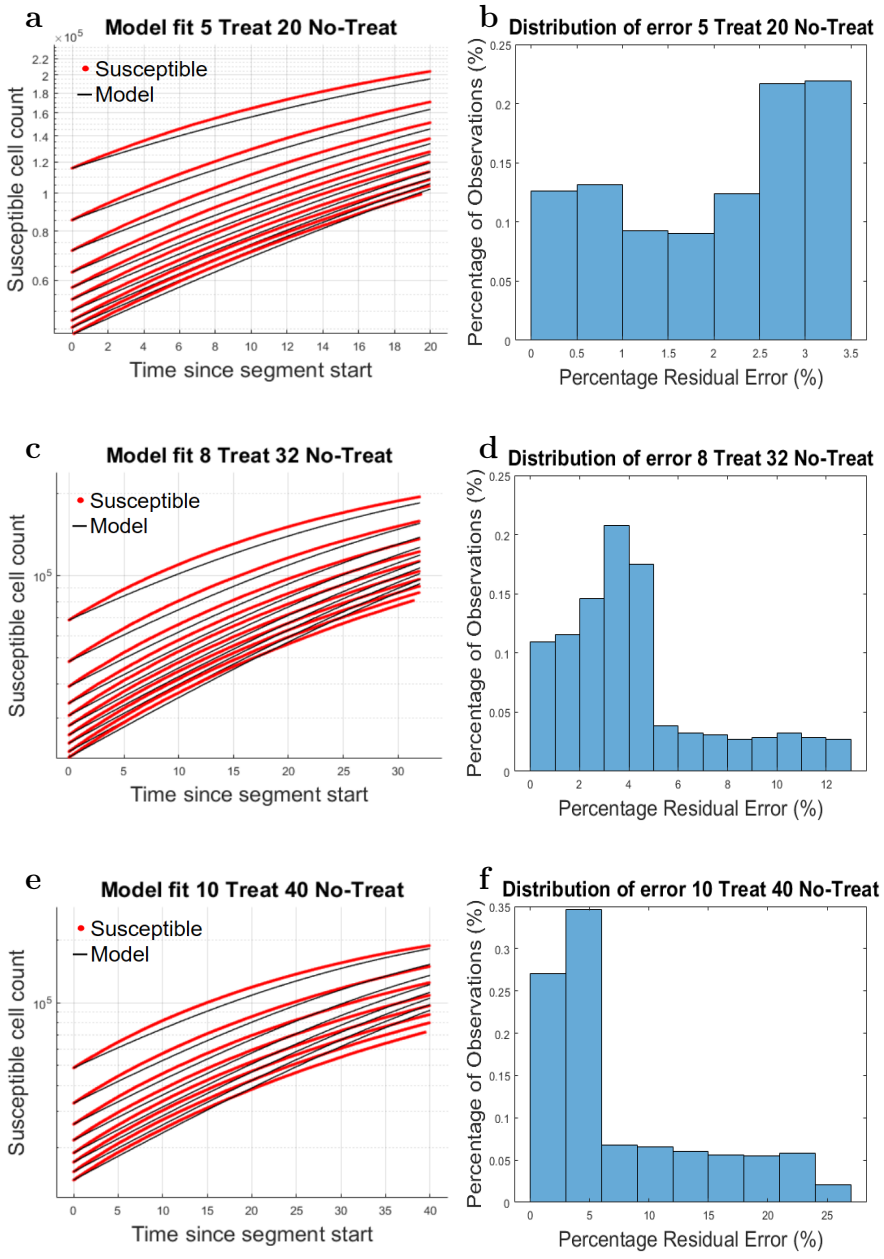


Figure 3.5: Panel a shows the susceptible cell count from the simulation with model fitted for treatment schedule 5 treat, 20 no-treat. Panel b shows the distribution of error produced by the fit. Panels c and d present the same type of information for the treatment schedule 8 treat, 32 no-treat. Panels e and f present the same type of information for the treatment schedule 10 treat, 40 no-treat. All data shown is from the no-treatment phases of the treatment regimen.

In figures 3.4 and 3.5 we can see that the model fits to the data reasonably well. The model fits better at lower overall treat, no-treat time as we can see in the error histograms in 3.4b,

3.4d, and even 3.5b. The error increases more drastically and the model fit is less accurate as we can see from the error histograms in figures 3.5d and 3.5f.

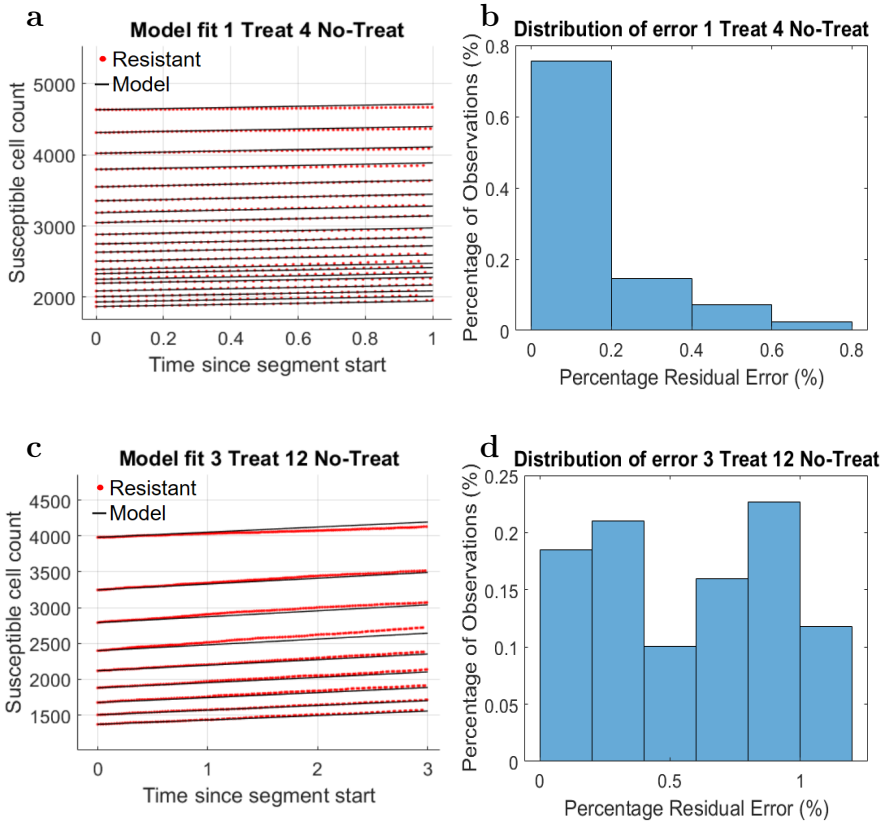


Figure 3.6: Panel a shows the resistant cell count from the simulation with model fitted to the data during the 1 time unit treatment phases. Panel b shows the distribution of error produced by the fit. Panels c and d present the same type of information for the treatment schedule 3 treat, 12 no-treat. All data shown is from the treatment phases of the treatment regimen.

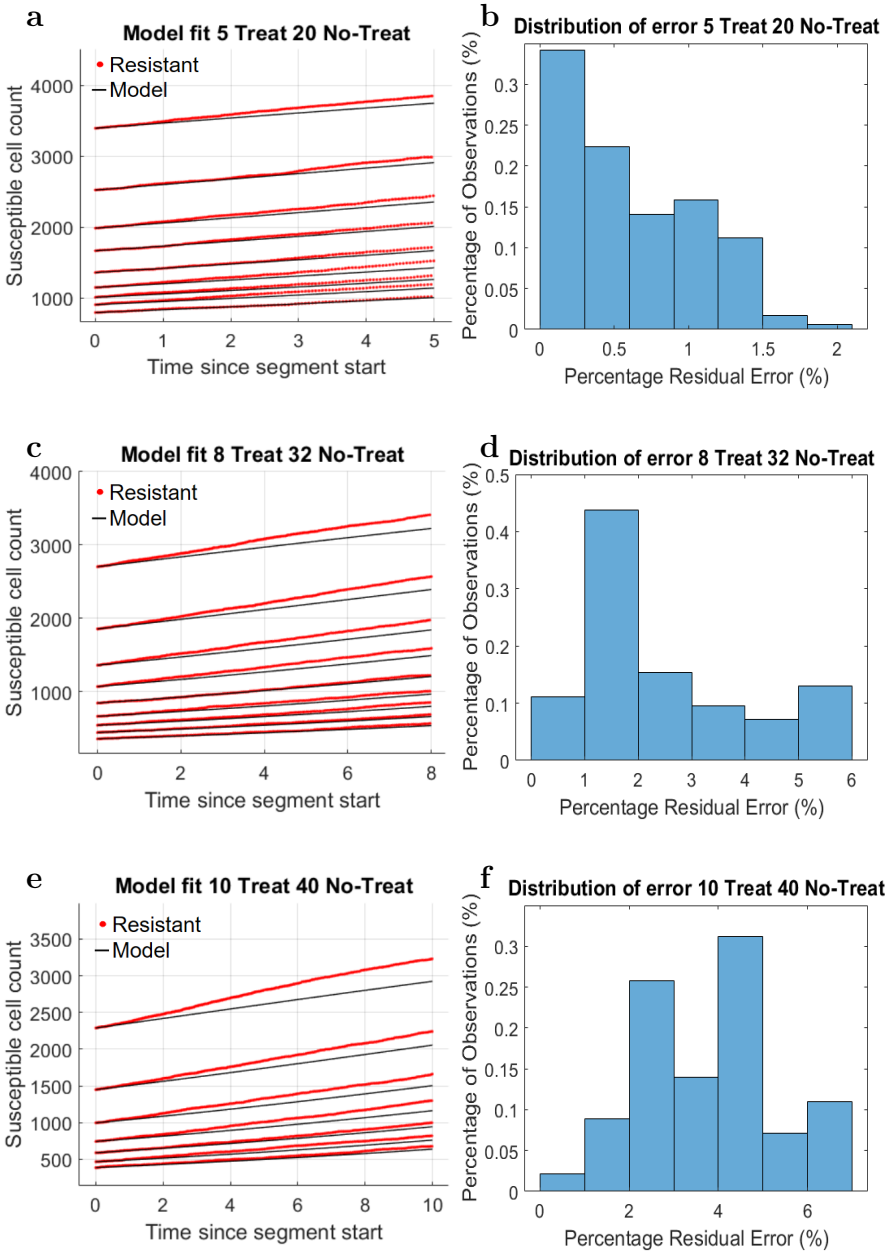


Figure 3.7: Panel a shows the susceptible cell count from the simulation with model fitted for treatment schedule 5 treat, 20 no-treat. Panel b shows the distribution of error produced by the fit. Panels c and d present the same type of information for the treatment schedule 8 treat, 32 no-treat. Panels e and f present the same type of information for the treatment schedule 10 treat, 40 no-treat. All data shown is from the treatment phases of the treatment regimen.

Figures 3.6 and 3.7 further quantifies the model fits of the simulated data. The fitted logistic model provides a stable approximation of resistant-cell growth during treatment. In

comparison with the susceptible fits, the residual errors increase less rapidly as the duration of the treatment phase grows. Thus, within the schedules examined, resistant-cell growth is represented more consistently by a single reduced logistic law.

To summarize, sparse regression consistently indicated that the aggregate growth phases could be represented by a logistic equation without retaining the higher order and square-root terms in the candidate library.

## 3.6 Reduced-Model Results

The logistic model captures the principal curvature of susceptible-cell growth during treatment withdrawal across the schedules considered. The fit is most accurate for shorter treatment and no-treatment phases. For longer phases, the residual errors increase, indicating that a single pair of aggregate coefficients does not capture every effect of the evolving spatial configuration.

The fits show that although aggregate nonlinear growth laws can be recovered from a spatial stochastic model, the accuracy of the reduction depends on the treatment schedule and cell type. The fitted quadratic term acts as an effective representation of spatial limitation: it converts local competition for empty lattice sites into a density-dependent correction at the population level. The reduced equations therefore retain information about spatial crowding without explicitly tracking the position of every cell.

While the logistic growth model provided reasonable fit under the kinetic parameter (that is, the rates $\gamma_i^{on}$ and $\gamma_i^{off}$, $i \in \{x, y\}$) investigated here, a more detailed study must be performed to explore other regions of the parameter space. At this time we cannot exclude the possibility that other laws, such as e.g. the surface growth model, may be more appropriate under different parameter combinations.

More generally, SINDy was able to yield reasonably good model dynamics with the important feature that we have found a fit that works for treatment regimens of the type $n$ treat, $4n$ no-treat. Further research is required to generate a deterministic ODE description for spatial tumor growth in a setting where treat and no-treat times are chosen more flexibly. This will eventually allow us to apply optimal control methods and seek an optimal treatment regimen.

## 3.7 Conclusions

This chapter examined melanoma treatment dynamics in a spatial setting where susceptible and resistant cells compete for empty sites. Our simulations demonstrated that the population-level consequences of spatial growth can be represented, to a useful approximation, by reduced ordinary differential equations. Exponential equations describe declining populations but do not adequately capture expansion under increasing spatial crowding. Sparse identification instead selects logistic growth laws for susceptible cells during treatment withdrawal and resistant cells during treatment. The linear terms describe effective low-density growth, while the quadratic terms incorporate the reduction in successful division caused by limited space.

The reduced models reproduce the principal trends of the agent-based time series over several intermittent treatment schedules. Their accuracy depends on cell type and phase duration, with resistant-cell growth exhibiting more consistent fits across the schedules examined. This dependence reflects the fact that aggregate dynamics retain information about the spatial configuration inherited from earlier phases of treatment.

The chapter thus establishes a connection between two modeling scales. The agent-based model describes local division, death, and competition for space, whereas the identified

differential equations provide a compact description of the resulting population trajectories. This connection offers a way to incorporate spatially generated growth effects into models that remain amenable to mathematical analysis and treatment optimization.

# Chapter 4

# Mathematical modeling of spatial stochastic dynamics of biofilms

## 4.1 Biological motivation and modeling objectives: biofilms of Vibrio cholerae

Biofilms are surface-associated communities of bacteria embedded in an extracellular matrix. This matrix contributes to the mechanical stability of the community, protects cells from environmental stresses, and promotes attachment to surfaces. Biofilm formation helps bacteria survive environmental stresses and persist in diverse habitats. From a human perspective, however, biofilms can contribute to persistent infections, increase tolerance to antimicrobial treatment, and obstruct or damage industrial systems. Understanding how cells interact with the matrix is therefore important both for explaining the biofilm life cycle and for developing methods to disrupt harmful biofilms.

Vibrio Cholerae can form biofilms by organizing into an extracellular matrix through Vibrio

polysaccharide (VPS) and matrix proteins, the secreted “sticky" substance that allows the biofilm to adhere to surfaces. Cells that actively produce VPS can become associated with the matrix and with neighboring cells through matrix proteins. As nutrient availability decreases, however, cells may reduce VPS production and lose the VPS coating on their surfaces. Experimental results indicate that these cells no longer interact attractively with the existing matrix. Instead, they are excluded from VPS-rich regions and become more susceptible to removal by flow. Thus, VPS appears to play two related roles: it supports adhesion and community stability during biofilm growth, while also enabling the recognition and eventual dispersal of cells that no longer produce matrix.

Thus, V. Cholerae cells can be categorized as producers (P), and non-producers (nP) based on there phenotype and behavior within the biofilm. The P-cells produce VPS and adhere more strongly to the surface and to neighboring producers, and nP-cells do not incur the metabolic cost of VPS production but are more readily removed by flow. The two states represent reversible phenotypes: producer cells can switch to the non-producer state, and non-producers can resume production. Because nutrient availability decreases as the population becomes crowded, these switching rates are assumed to depend on cell density.

This system contains two competing fitness advantages. Non-producers divide more rapidly because they do not bear the cost of VPS production. Producers, by contrast, experience lower removal rates under flow because of their stronger adhesion. The relative success of the two phenotypes therefore depends on the physical environment. In the absence of flow, the reproductive advantage of non-producers can dominate. Under flow, their weak adhesion becomes costly, while producers are preferentially retained in stable spatial clusters.

To investigate these interactions, we formulate a spatial stochastic model on a two-dimensional lattice. Individual cells undergo division, death, flow-induced removal, and phenotype switching. The removal rate of a producer decreases when it is surrounded by other producers, representing the additional protection provided by local adhesion and cluster formation.

The spatial formulation is essential because this protection depends on cellular neighborhoods and cannot be represented by the total numbers of producers and non-producers.

The model is designed to address three questions. First, how do the reproductive cost of VPS production and its adhesive benefit determine the composition of the biofilm? Second, how does the onset of flow alter the balance between producer and non-producer phenotypes? Third, how does flow-induced removal affect phenotype switching indirectly by changing population density?

The simulations described below examine how the onset of flow changes the composition and spatial organization of the biofilm. In particular, they show how differential removal and density-dependent phenotypic switching combine to determine which cells remain attached to the surface. These results provide a mathematical explanation for how phenotypic switching and environmental flow can coordinate biofilm dispersal and persistence.

## 4.2 Model assumptions and parameterization

We model the population of V. Cholerae cells by using a square 2D grid of $N \times N$ spots. Each spot can be empty, or occupied with a bacterium, which can be a producer (P) or a non-producer cell (nP). The dynamics proceed as a sequence of updates that include cell divisions, death/wash out events, and conversion between the two types. Variables used in the model are summarized in Table 4.1.

**System size and cell density.** Cells are spaced roughly with distance of $1 - 3\mu m$ apart. The fact that the vertical system size (several layers of cells, or 20-30$\mu m$) is much smaller than the horizontal dimension, motivates the 2D approximation used here, which integrates out the dynamics along the vertical axis.

| Notation | Parameter definition | Unit |
|---|---|---|
| $r_P$ | Division rate of P | $\text{hrs}^{-1}$ |
| $r_{nP}$ | Division rate of nP | $\text{hrs}^{-1}$ |
| $s$ | Relative advantage from being nP | 1 |
| $d_P, d_{nP}$ | Removal rates for P and nP | $\text{hrs}^{-1}$ |
| $d$ | Basic death rates for both P and nP | $\text{hrs}^{-1}$ |
| $d_f$ | Wash-out rate due to flow, no sticking | $\text{hrs}^{-1}$ |
| $A$ | Reduction in wash-out rate for a single P cell | 1 |
| $\nu$ | The normalization coefficient in the model of adhesion | 1 |
| $\mu_{P\to nP}$ | Conversion rate from P to nP | $\text{hrs}^{-1}$ |
| $\mu_{nP\to P}$ | Conversion rate from nP to P | $\text{hrs}^{-1}$ |
| $\mu$ | Max conversion rate from P to nP and from nP to P | $\text{hrs}^{-1}$ |
| $\beta$ | Exponents in the expressions for the conversion rates | 1 |
| $N$ | Linear grid size | $2\mu m$ |

Table 4.1: Definitions of parameters, their notations and units.

**Cell division.** Division happens by picking a random neighbor out of the 4 neighbors, and placing a cell of the same type in that spot, unless that spot is already occupied. If it is occupied, reproduction does not happen. P cells produce a substance that allows them to stick to the surface and to each other. This substance is costly, meaning that the division rate of P is reduced due to the fact that they use resources to produce the substance. Denote by $r_P$ and $r_{nP}$ the division rates of P and nP cells, respectively. The reproductive fitness advantage of nP cells is denoted by $s$, such that $r_{nP} = (1 + s)R_P$. The measured cell replication every 30 mins translates into the division rate of $r_P = \ln 2/0.5 = 1.39$ $\text{hrs}^{-1}$. We further assume that the relative advantage of nP cells is $s = 0.2$.

**Cell death and washing out.** We assume that all bacteria die at a certain rate, denoted by $d$. In the absence of a flow, the two types have the same death rate, which is small compared to the division rate ($d = 0.05r_P$). In the presence of a flow, we assume nP cells are removed /washed out from the system at a constant rate, which is flow-dependent. The wash-out rate of P cells depends on how many P neighbors they have. The more P neighbors surround a P cell, the less its wash-out rate becomes. Note that P cells are assumed to have

a smaller wash-out rate than nP cells, even in the absence of P neighbors. The total removal rate consists of the death rate $d$ and a wash-out rate (in the presence of a flow):

$$d_{nP} = \begin{cases} d, & \text{no flow} \\ d + d_f, & \text{with flow} \end{cases}, \quad d_P = \begin{cases} d, & \text{no flow} \\ d + A\, d_f \left(1 - \frac{k}{\nu}\right), & \text{with flow} \end{cases} \tag{4.1}$$

Here $k$ is the number of P cells among the neighbors of a given P cell, and coefficient $\nu = 5$ is used to characterize the force of adhesion force of P cells to each other. The constant $0 < A < 1$ indicates that even if a P cell has no P neighbors, its wash-out rate ($Ad_f$) is still smaller than that of a nP cell: $Ad_f < d_f$.

**The process of conversion.** The conversion P$\rightarrow$ nP and nP$\rightarrow$P happens in a density-dependent manner. It is assumed that the amount of nutrients available is negatively correlated with the total number of cells. A lack of nutrients promotes the P$\rightarrow$ nP conversion, which is modeled as

$$\mu_{P\rightarrow nP} = \mu \left(\frac{k'}{N^2}\right)^{\beta}, \tag{4.2}$$

where $k'$ is the total number of cells (P or L) in the system, $k'/N^2$ is their density, and $\beta = 2$. To estimate $\mu$, we assume that the maximum conversion rate is such that 90% of cells are converted within 4 hours. This results in $\mu = |\ln 0.1|/4 = 0.58\text{hrs}^{-1}$. In our simulations we used values of $\mu$ ranging from $0.1\text{hrs}^{-1}$ to $0.58\text{hrs}^{-1}$. The conversion nP$\rightarrow$ P increases with the amount of available nutrients:

$$\mu_{nP\rightarrow P} = \bar{\mu} \left(1 - \frac{k'}{N^2}\right)^{\beta}. \tag{4.3}$$

## 4.3 Stochastic simulation method

The biofilm simulations use the same event-driven numerical framework as the spatial melanoma simulations in Chapter 3. The distinction between the two applications lies in the biological states and event propensities, rather than in the underlying simulation algorithm.

In the images shown in this chapter, the grid size was $50 \times 50$, with periodic boundary conditions. The grid was initialized by placing a P cell with probability 0.15 and an nP cell with probability 0.085.

We implement the model on a rectangular 2-dimensional matrix, such that at any given time each spot in the matrix can be empty or occupied by a single cell of any type. The distance between individuals in the population is determined by the distance between the spots they occupy in the matrix using the $\ell^1$ distance (also known as the Manhattan distance).

The simulations are carried out based on the stochastic simulation algorithm (SSA) [13]. More precisely, given a state of the system $\mathbf{x}$, every possible event $\mu$ has a given propensity $a_\mu(\mathbf{x})$ based on the model's rates. The time at which the next reaction $\mu$ will occur is exponentially distributed with intensity $a_\mu(\mathbf{x})$. Thus if $r_\mu$ is a uniform random number in $[0, 1)$, setting the time for the next $\mu$-reaction as $\tau_\mu = -\log(r_\mu)/a_\mu(\mathbf{x})$ will generate random times with the correct probability distribution. This algorithm produces exact solutions to the master equation resulting from the specified rates. In particular, this means that each simulated outcome (one run of the algorithm) occurs with the correct statistical frequency.

We code the algorithm using indexed priority queues implemented with binary heaps to find the next stochastic event. We do so using the so called Next Reaction Method [12]. We also used indexed priority queues to update the effects of nearby neighbors on a given individual. As a result the simulations are fast and amenable to large populations, with an overall algorithmic complexity of $\mathcal{O}(n \log n)$, where $n$ is the number of individuals.

Each simulation begins with a randomly initialized mixture of producer and non-producer cells and is first allowed to evolve in the absence of flow. This initial phase permits the population to approach the high-density state in which non-producers are favored by their higher division rate and by density-dependent conversion from (P) to (nP). Flow is then introduced at time $t = t_f$, and the subsequent changes in producer abundance, non-producer abundance, total density, and spatial organization are recorded. Results are summarized over many independent realizations of the stochastic process. This protocol is designed to reproduce the experimental comparison between an established biofilm before washing and its response after exposure to flow.

## 4.4 Producer and Non-Producer Dynamics Under Flow

We now use the spatial stochastic model to examine how an established biofilm responds to the onset of flow. The system was first simulated in the absence of flow until time $t = 70$ hours. Flow was then introduced, while all other parameter values were held fixed. This creates a controlled change in the environment: differences between the populations before and after $t = 70$ can be attributed to the flow-dependent removal mechanisms defined above and to the subsequent effects of cell removal on population density and phenotypic conversion.

### 4.4.1 Flow reverses the composition of the population

Before flow was introduced, non-producer cells constituted the majority of the population. This pre-flow state provides an important baseline. In an environment without flow, the adhesive benefit of VPS production does not affect cell removal, while the metabolic cost of production reduces the division rate of producer cells. In addition, the high population

density reached during biofilm growth favors conversion toward the non-producer phenotype. The pre-flow population is therefore dominated by cells that do not produce VPS.

The composition of the population changed markedly after flow was introduced. Non-producer cells were rapidly depleted, whereas a larger fraction of the producer population remained on the lattice. As a result, the population changed from a dense community dominated by non-producers to a smaller attached community dominated by producers. The spatial configurations in Figure 4.1 illustrate this transition.

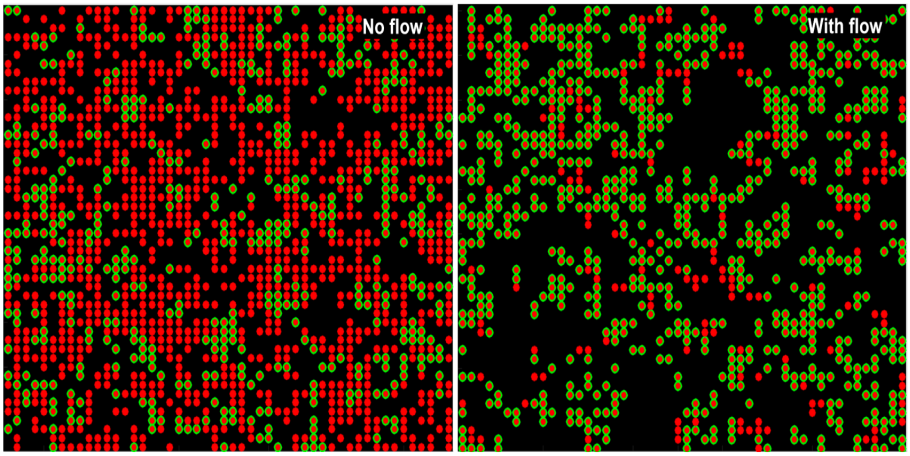


Figure 4.1: Representative snapshots of results from a spatial stochastic model on a 2D grid including producer (P in red with green outline) and non-producer (nP in red) cells, before (left) and after (right) turning on flow.

The spatial structure of the surviving population is as important as the change in total cell number. Producer cells do not persist independently of their surroundings: their protection from removal increases when they have producer neighbors. Consequently, flow preferentially eliminates cells in poorly protected regions while preserving groups of adjacent producers. The result is a form of spatial selection in which the imposed environment acts not only on cell phenotype but also on local cellular organization. Producer-rich neighborhoods are retained, whereas cells lacking access to the adhesive protection associated with VPS production are more readily removed.

The corresponding changes in population size are summarized in Figure 4.2.

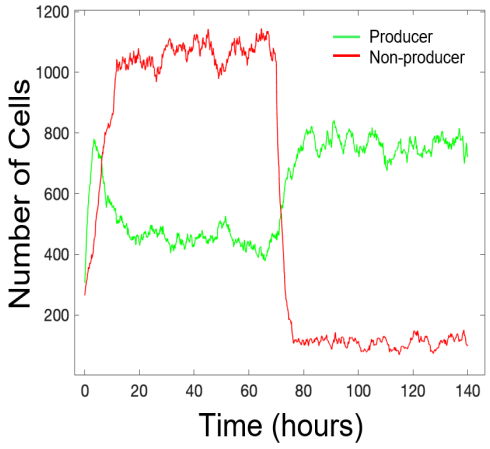


Figure 4.2: Quantification of the number of P and nP cells before and after turning on flow (at t = 70 hours).

Before flow, non-producers outnumber producers. After flow is introduced, the non-producer population decreases sharply and producers become the majority among the cells that remain attached. Flow therefore does not simply reduce the biomass uniformly. It selectively changes the phenotypic composition of the population.

### 4.4.2 Feedback between cell removal and phenotypic conversion

The reversal in population composition is not produced by differential removal alone. Flow initially acts directly by removing non-producer cells more rapidly than producer cells. This loss of cells lowers the total population density, which in turn changes the relative rates of conversion between the two phenotypes.

At the high densities attained before flow, conversion from the producer state to the non-producer state predominates. Following the loss of cells under flow, conversion toward the producer phenotype becomes more favorable. The environmental perturbation therefore produces a two-stage response. First, flow selectively removes the less adhesive phenotype.

Second, the resulting reduction in density shifts phenotypic conversion toward the producer state. Newly generated producers are then less susceptible to removal, particularly when they occur near other producers.

This feedback stabilizes the attached population. Without the dependence of phenotypic conversion on density, continued switching toward the non-producer state could progressively replace protected producer cells with cells that are readily washed away. In the present model, however, the decrease in density caused by flow also promotes restoration of the producer phenotype. The response is therefore self-reinforcing: selective removal lowers density, lower density favors producers, and producer-rich neighborhoods are more resistant to subsequent removal.

The outcome depends on how strongly the conversion rates respond to density. For example, when $\mu = 0.58\,\mathrm{hr}^{-1}$ and $\beta = 0.1$, the population collapses in the presence of flow. For this small value of $\beta$, conversion toward the non-producer phenotype remains substantial over a broad range of densities, weakening the recovery of the protected producer population. When $\beta = 2$, the reduction in density produces a stronger shift toward the producer phenotype, and nonzero steady-state populations of both phenotypes can be maintained. The value $\beta = 2$ was therefore used in the simulations shown above. These examples demonstrate the importance of coupling environmental removal to a sufficiently density-sensitive phenotypic response.

### 4.4.3 Biological interpretation

The simulations show how the cost and benefit of VPS production depend on the physical environment. Before flow, production is costly but provides little advantage against removal, and non-producers become abundant. Under flow, the same adhesive properties become beneficial, reversing the relative success of the two phenotypes. Neither producer nor non-

producer status is therefore universally advantageous. The favored phenotype is determined by the interaction between cellular behavior and environmental conditions.

The results also suggest that VPS can act as a mechanism of selective retention. Producer cells contribute to the matrix and are preferentially maintained within the attached community, whereas cells that cease production become more susceptible to removal. In this sense, the matrix does more than provide mechanical support. It distinguishes between cells according to their current phenotypic state and determines which cells remain associated with the biofilm under flow.

This selective retention provides a possible connection between biofilm persistence and dispersal. Producer-rich clusters preserve an attached population that can continue maintaining the biofilm. Non-producer cells, by contrast, are more readily released into the surrounding environment and may act as a dispersing population capable of reaching new locations. Phenotypic switching allows these two outcomes to arise within a single population without requiring permanent genetic differences between the cells. A cell can leave the protected producer state when conditions favor non-production and can potentially resume production when environmental conditions or population density change.

The model therefore offers a population-level interpretation of how an established biofilm can simultaneously maintain a surface-associated community and generate cells that are susceptible to dispersal. It does not represent the molecular structure of VPS or the detailed mechanics of cell–matrix interactions. Instead, those effects are incorporated through phenotype- and neighborhood-dependent removal rates. The agreement between the simulated behavior and the qualitative biological observations shows that differential growth, local adhesive protection, flow-induced removal, and density-dependent phenotypic switching are sufficient to generate the observed transition from a non-producer-rich population to an attached, producer-rich population.

The present analysis assumes that cells which do not produce VPS also fail to benefit from VPS produced by neighboring cells. The next section relaxes this assumption by considering cheater cells that do not contribute to matrix production but can obtain an adhesive benefit from nearby producers.

## 4.5 Producers and cheaters

This modeling framework allow further exploration of different assumptions. We assumed above that nP cells are repelled by producer secretions. A different type of cells, which we term Cheaters (C, referring to the usual terminology of the evolutionary game theory), have the ability to utilize these secretions to gain benefit from its adhesive effect. In this modified version of the model, C will take on the role of nP, requiring only a change to the total removal rate in the presence of flow. The total removal rate of C cells is comprised of the death rate $d$ and wash-out rate (in the presence of a flow):

$$d_C = \begin{cases} d, & \text{no flow} \\ d + d_f \left(1 - \frac{k}{\nu}\right), & \text{with flow} \end{cases} \tag{4.4}$$

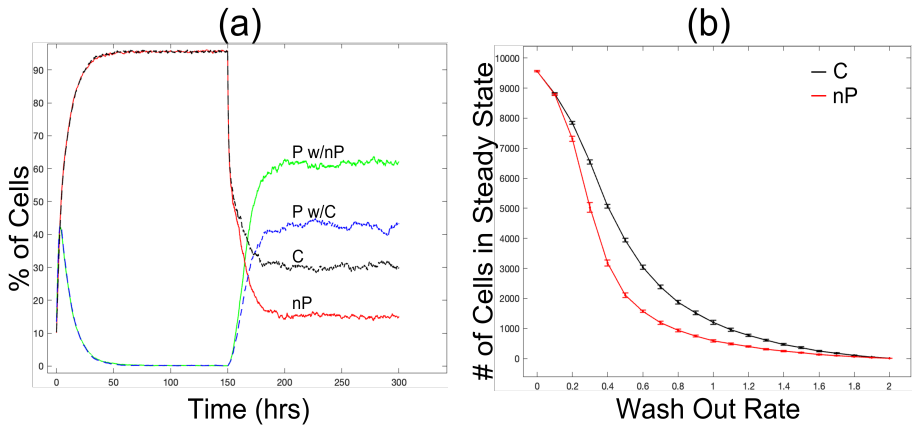


Figure 4.3: Figure (a) is a dual time-series plot of simulations of P with nP and P with C. This plot demonstrates the advantage that C has over nP in the presence of a flow since C has a reduced effect of washing out depending on the number of surrounding P neighbors. Figure (b) plots the average steady state value of C and nP as it depends on the wash out rate term $d_f$. C and nP have, on average, the same steady state value in the absence of a flow ($d_f = 0$) and in the presence of high flow ($d_f = 2$), but otherwise C shows improved staying power in the presence of a flow. Listed here are the parameter values used in figures (a) and (b): $A = .5$, $\beta = 2$, $\mu = .1$, $r_P = 1.39$, $r_{np} = 1.2 * 1.39$, $d = .05 * r_P$, $d_f = .6$ (in (a)).

Using numerical simulations, we have explored differences between the co-dynamics of P and nP cells on the on hand, and P and C cells on the other. To do this we separately ran simulatons with P and nP cells, and simulations with P and C cells. This is illustrated in figure S1. Panel (a) plots the time-series simulations of P with nP (green and black lines) and P with C (blue and black lines). The equilibria in the absence of flow are identical in both cases, and we observe that P cells comprise a small minority of the population. After the flow is applied, we can see that the balance is reversed, and now P cells are the majority. The proportion of C cells in the presence of flow (under all parameters being equal) is larger compared to the proportion of nP cells (the black line is above the red line).

This trend is investigated more systematically in panel (b) of the figure, where we plot the steady state density of nP cells and C cells for different values of $d_f$ (the flow). As mentioned above, the two solutions coincide under no flow ($d_f = 0$), and they also become identical for very large values of $d_f$, where all cells are extinct. For intermediate flow values, there are

more C cells that can be maintained in the system. The reason for this is their ability to utilize the secretions of P cells and stick to their P neighbors.

## 4.6 Biological Interpretation and Conclusions

The spatial stochastic model demonstrates how the reproductive cost of matrix production can be balanced by its protective effect under flow. Without flow, non-producer cells are favored by their higher division rate and by density-dependent switching from the producer state. The onset of flow reverses this advantage: non-producers are removed rapidly, whereas producers are preferentially retained, particularly when they occur in clusters containing multiple producer neighbors.

Flow also changes the system indirectly. The removal of non-producers lowers the population density, which shifts the predominant direction of phenotypic conversion from (P

# Bibliography


[1] A. P. Algazi, M. Othus, A. I. Daud, R. S. Lo, J. M. Mehnert, T.-G. Truong, R. Conry, K. Kendra, G. C. Doolittle, J. I. Clark, et al. Continuous versus intermittent BRAF and MEK inhibition in patients with BRAF-mutated melanoma: a randomized phase 2 trial. *Nature medicine*, 26(10):1564–1568, 2020.

[2] A. Borzì. *The Sequential Quadratic Hamiltonian Method: Solving Optimal Control Problems*. Chapman and Hall/CRC Numerical Analysis and Scientific Computing Series. Taylor & Francis, 2023.

[3] A. Borzì and S. Roy. Optimal control of the viscous wave equation via the pontryagin maximum principle. *Numerical Methods for Partial Differential Equations*, 42(3):e70103, 2026.

[4] A. Bratus, E. Fimmel, and S. Y. Kovalenko. On assessing quality of therapy in nonlinear distributed mathematical models for brain tumor growth dynamics. *Mathematical Biosciences*, 248:88–96, 2014.

[5] E. Can. Fractional spatiotemporal hahnfeldt tumor model with convergence analysis and optimal control. *Scientific Reports*, 2026.

[6] P. B. Chapman, A. Hauschild, C. Robert, J. B. Haanen, P. Ascierto, J. Larkin, R. Dummer, C. Garbe, A. Testori, M. Maio, et al. Improved survival with vemurafenib in melanoma with BRAF V600E mutation. *New England Journal of Medicine*, 364(26):2507–2516, 2011.

[7] A. Dey, A. Borzì, and S. Roy. A high contrast and resolution reconstruction algorithm in quantitative photoacoustic tomography. *Journal of Computational and Applied Mathematics*, page 116065, 2024.

[8] S. Esmaili, M. Eslahchi, and D. F. Torres. Optimal control for a nonlinear stochastic pde model of cancer growth. *Optimization*, 73(9):2745–2789, 2024.

[9] M. Fallahi-Sichani, V. Becker, B. Izar, G. J. Baker, J.-R. Lin, S. A. Boswell, P. Shah, A. Rotem, L. A. Garraway, and P. K. Sorger. Adaptive resistance of melanoma cells to RAF inhibition via reversible induction of a slowly dividing de-differentiated state. *Molecular systems biology*, 13(1):905, 2017.

[10] F. Fu, M. A. Nowak, and S. Bonhoeffer. Spatial heterogeneity in drug concentrations can facilitate the emergence of resistance to cancer therapy. *PLoS computational biology*, 11(3):e1004142, 2015.

[11] R. A. Gatenby, A. S. Silva, R. J. Gillies, and B. R. Frieden. Adaptive therapy. *Cancer research*, 69(11):4894–4903, 2009.

[12] M. A. Gibson and J. Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *J. Phys. Chem. A*, 104:1876–1889, 2000.

[13] D. T. Gillespie. Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.*, 58:35–55, 2007.

[14] M. Gluzman, J. G. Scott, and A. Vladimirsky. Optimizing adaptive cancer therapy: dynamic programming and evolutionary game theory. *Proceedings of the Royal Society B*, 287(1925):20192454, 2020.

[15] M. Holderfield, M. M. Deuker, F. McCormick, and M. McMahon. Targeting RAF kinases for cancer therapy: BRAF-mutated melanoma and beyond. *Nature Reviews Cancer*, 14(7):455–467, July 2014.

[16] E. Khailov and E. Grigorieva. Optimal melanoma treatment protocols for a bilinear control model. *Mathematics*, 11(15):3289, 2023.

[17] E. Kim, J. S. Brown, Z. Eroglu, and A. R. Anderson. Adaptive therapy for metastatic melanoma: predictions from patient calibrated mathematical models. *Cancers*, 13(4):823, 2021.

[18] E. Kim, V. W. Rebecca, K. S. Smalley, and A. R. Anderson. Phase i trials in melanoma: A framework to translate preclinical findings to the clinic. *European Journal of Cancer*, 67:213–222, 2016.

[19] N. L. Komarova, J. R. Pritchard, and D. Wodarz. Efficient mathematical methodology to determine multistep mutant burden in spatially growing cell populations. *PNAS nexus*, 4(9):pgaf271, 2025.

[20] F. Li and B. You. Optimal control for a reaction–diffusion model with tumor-immune interactions. *Communications in Nonlinear Science and Numerical Simulation*, 145:108677, 2025.

[21] G. Moriceau, W. Hugo, A. Hong, H. Shi, X. Kong, C. C. Yu, R. C. Koya, A. A. Samatar, N. Khanlou, J. Braun, et al. Tunable-combinatorial mechanisms of acquired resistance limit the efficacy of BRAF/MEK cotargeting but result in melanoma drug addiction. *Cancer cell*, 27(2):240–256, 2015.

[22] R. Nazarian, H. Shi, Q. Wang, X. Kong, R. C. Koya, H. Lee, Z. Chen, M.-K. Lee, N. Attar, H. Sazegar, et al. Melanomas acquire resistance to B-RAF (V600E) inhibition by RTK or N-RAS upregulation. *Nature*, 468(7326):973–977, 2010.

[23] A. Nowakowski and A. Krawczyk. Control of tumor growth modeled by system of pde, numerical analysis of optimality conditions. *Mathematical Methods in the Applied Sciences*, 45(16):9371–9385, 2022.

[24] L. S. Pontryagin, V. G. Boltyanskiĭ, R. V. Gamkrelidze, and E. F. Mishchenko. *The Mathematical Theory of Optimal Processes*. John Wiley & Sons, New York-London, 1962.

[25] I. A. Rodriguez-Brenes, N. L. Komarova, and D. Wodarz. Tumor growth dynamics: insights into evolutionary processes. *Trends in ecology & evolution*, 28(10):597–604, 2013.

[26] S. Roy. A new nonlinear sparse optimization framework in ultrasound-modulated optical tomography. *IEEE Transactions on Computational Imaging*, 8:1–11, 2022.

[27] S. Roy and S. Pal. A PINN-driven game-theoretic framework in limited data photoacoustic tomography. *Inverse Problems*, 2025. to appear.

[28] A. Ruiz-Martinez, C. Gong, H. Wang, R. J. Sové, H. Mi, H. Kimko, and A. S. Popel. Simulations of tumor growth and response to immunotherapy by coupling a spatial agent-based model with a whole-patient quantitative systems pharmacology model. *PLoS computational biology*, 18(7):e1010254, 2022.

[29] D. Schadendorf, A. C. J. van Akkooi, C. Berking, K. G. Griewank, R. Gutzmer, A. Hauschild, A. Stang, A. Roesch, and S. Ugurel. Melanoma. *Lancet*, 392(10151):971–984, Sept. 2018.

[30] H. Shi, G. Moriceau, X. Kong, M.-K. Lee, H. Lee, R. C. Koya, C. Ng, T. Chodon, R. A. Scolyer, K. B. Dahlman, et al. Melanoma whole-exome sequencing identifies V600E B-RAF amplification-mediated acquired B-RAF inhibitor resistance. *Nature communications*, 3(1):724, 2012.

[31] I. Smalley, E. Kim, J. Li, P. Spence, C. J. Wyatt, Z. Eroglu, V. K. Sondak, J. L. Messina, N. A. Babacan, S. S. Maria-Engler, et al. Leveraging transcriptional dynamics to improve BRAF inhibitor responses in melanoma. *EBioMedicine*, 48:178–190, 2019.

[32] J. A. Sosman, K. B. Kim, L. Schuchter, R. Gonzalez, A. C. Pavlick, J. S. Weber, G. A. McArthur, T. E. Hutson, S. J. Moschos, K. T. Flaherty, P. Hersey, R. Kefford, D. Lawrence, I. Puzanov, K. D. Lewis, R. K. Amaravadi, B. Chmielowski, H. J. Lawrence, Y. Shyr, F. Ye, J. Li, K. B. Nolop, R. J. Lee, A. K. Joe, and A. Ribas. Survival in BRAF V600-mutant advanced melanoma treated with vemurafenib. *New England Journal of Medicine*, 366(8):707–714, Feb. 2012.

[33] M. A. Strobl, J. Gallaher, J. West, M. Robertson-Tessi, P. K. Maini, and A. R. Anderson. Spatial structure impacts adaptive therapy by shaping intra-tumoral competition. *Communications medicine*, 2(1):46, 2022.

[34] M. D. Thakur, F. Salangsang, A. S. Landman, W. R. Sellers, N. K. Pryer, M. P. Levesque, R. Dummer, M. McMahon, and D. D. Stuart. Modelling vemurafenib resistance in melanoma reveals a strategy to forestall drug resistance. *Nature*, 494(7436):251–255, Feb. 2013.

[35] J. West, J. Gallaher, M. Strobl, M. Robertson-Tessi, and A. R. Anderson. The fundamentals of evolutionary therapy in cancer. *Cancer Systems Biology and Translational Mathematical Oncology*, 2023.

[36] J. West, L. You, J. Zhang, R. A. Gatenby, J. S. Brown, P. K. Newton, and A. R. Anderson. Towards multidrug adaptive therapy. *Cancer research*, 80(7):1578–1589, 2020.

[37] D. Wodarz and N. L. Komarova. Mutant evolution in spatially structured and fragmented expanding populations. *Genetics*, 216(1):191–203, 2020.

[38] J. Zhang, J. J. Cunningham, J. S. Brown, and R. A. Gatenby. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nature communications*, 8(1):1816, 2017.

# Appendix A

# Appendix Title

## A.1 Additional numerical illustrations

Here we present figure A.1, which is a companion plot for figure 2.3. While $20 \times 20$ initial conditions went into creating the heat plot in figure 2.3, here we show $5 \times 5$ of them sampling $\log_{10} x_0 \in [-9, 1]$ with the step-size of 2, and $y_0 \in [10^{-7}, 10^{-6}]$ with the step-size of $2 \times 10^{-7}$. The initial conditions are indicated above each plot, and the yellow shading is used to highlight cases with initial conditions above the special ray.

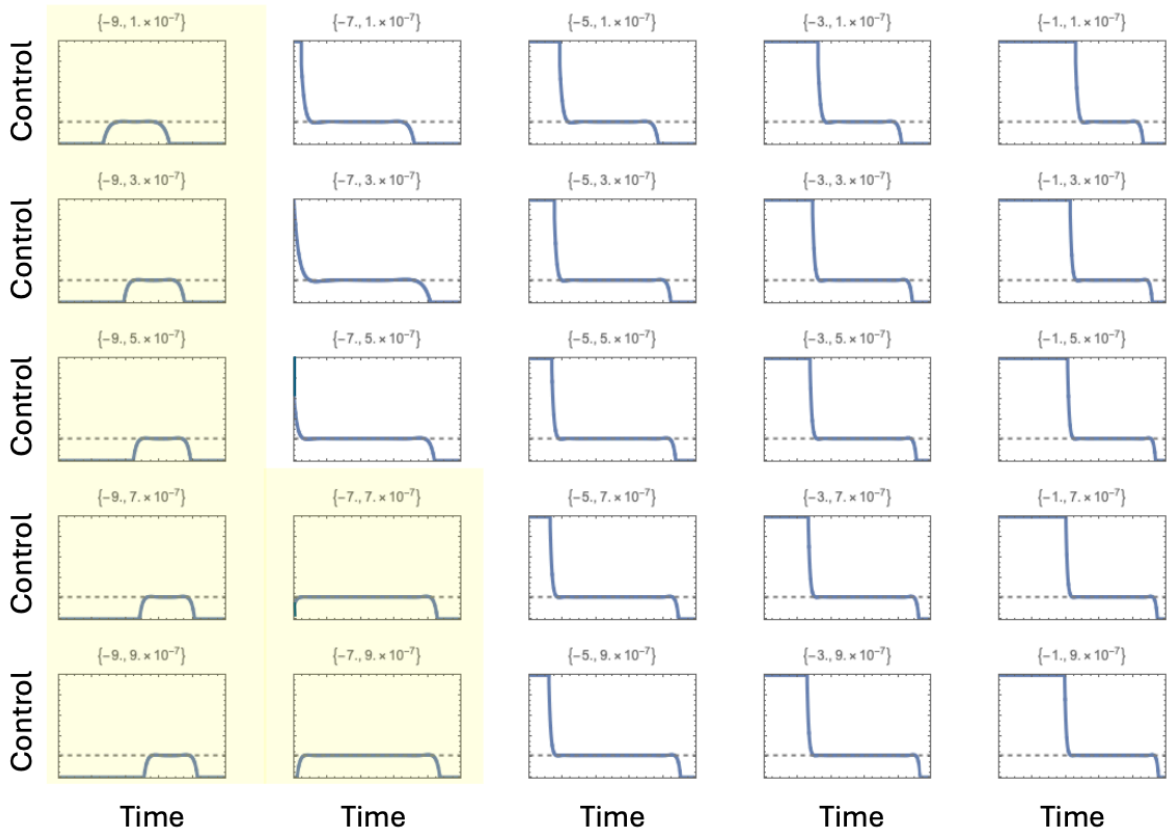


Figure A.1: Dependence of the optimal control on the initial conditions of the problem. A subset of optimal control solutions used in figure 2.3 is presented. Each graph shows the optimal control for a specific set of initial conditions, which are specified above the plot as $\{\log_{10} x_0, y_0\}$. The horizontal axis for each graph (time) has a range of $[0, 30]$, and the vertical axis (control) has the range of $[0, 1]$ (arbitrary units). The shaded plots correspond to the initial conditions above the special ray, equation (1.11). Here $\alpha = 5.5 \times 10^{-7}$, and the rest of the parameters are given by (2.3).

All the simulations were carried out in MATLAB on a Windows 11 64-bit laptop with 3.2 GHz processing speed and 32 GB RAM. The typical number of iterations taken for the SQH algorithm to converge were 1175 with a CPU time of 112.55 seconds.

## A.2 Details of the data fitting procedure

The data presented in figure 2.5 were fitted to our model, ((1.9)). This was formulated as the following constrained least-squares optimization problem:

$$\begin{aligned} &\min_{\boldsymbol{\theta}\in T_{ad}} f(C,\boldsymbol{\theta}) := \frac{1}{2}\sum_{j=1}^{N_p}(C(t_j)-C_j)^2, \\ &\text{subject to } C'(t) = (A+Bu)C(t),\ C(0)=C_0. \end{aligned} \tag{A.1}$$

Here, $\boldsymbol{\theta} = (A, B)$ is the parameter vector estimated using the minimization problem (A.1) in the admissible set of parameters $T_{ad}$. Due to the dataset types, as discussed later, it is possible to fit the susceptible and the resistant tumor dynamics at times $t_j$, given in (1.9), separately to the data. Thus, when $C_j$ represents the susceptible population data, $C$ stands for the variable $x$ in (1.9), $A = a,\ B = b$, and the admissible set $T_{ad}$ is chosen as $[0.001, 1] \times [-0.3, -0.01]$. On the other hand, when $C_j$ represents the resistant population, $C$ stands for the variable $y$ in (1.9), $A = c,\ B = d$, and the admissible set $T_{ad}$ is chosen as $[-0.1, -0.01] \times [-0.0001, 0.1]$. The parameter estimation problem (A.1) was solved using a multistart algorithm in MATLAB that executed repeated runs of the interior-point optimization algorithm with different choices of initial guesses to eventually obtain a close approximation to the global minima of (A.1).

The fit in panel (c) of figure 2.5 (that is, the dynamics of resistant cells in the absence of treatment) is visibly poorer than that for the rest of the cases. This is because the experimental data from [34] used here do not show straight exponential decay. The authors of Ref. [34] report that vemurafenib-resistant melanoma cells experience a fitness disadvantage in the absence of drug based on a combination of in vitro and in vivo observations. This conclusion is based on the overall pattern across several experiments rather than on a reported

statistical test demonstrating a significant negative growth rate in their Fig. 4b alone. In particular, in vitro, resistant 45V-RT cells showed maximal proliferation at low concentrations of vemurafenib and reduced proliferation when the drug was withdrawn, accompanied by morphological changes and elevated ERK signaling (Fig. 3a–c). In vivo, resistant tumors implanted into mice grew less efficiently when animals were treated with vehicle rather than vemurafenib (Fig. 4a). Most importantly, when established resistant tumors were switched from vemurafenib to vehicle, tumor volume initially decreased over approximately 10 days (Fig. 4b), while ERK signaling increased (Fig. 4c). The authors interpret this transient regression, together with the reduced tumor establishment in Fig. 4a and the drug-dependent proliferation observed in Fig. 3, as evidence that resistant cells tend to decline in the absence of drug. We used the data in their Figure 4(b) to obtain an order-of-magnitude estimate of the decay rate.