



City Research Online

City St George's, University of London

Citation: Coyler, B. A. (2025). Combined Bayesian And Theoretical Analysis of Cancer Treatment Strategies and Incidence Data. (Unpublished Doctoral thesis, City St George's, University of London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/36444/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

CITY ST. GEORGE'S, UNIVERSITY OF LONDON

DOCTORAL THESIS

**Combined Bayesian And Theoretical
Analysis of Cancer Treatment Strategies and
Incidence Data**

Author:

Blair Alexander COLYER

Supervisor:

Dr. Robert NOBLE

A thesis submitted in fulfillment of the requirements

for the degree of Doctor of Philosophy

in the

Department of Mathematics

School of Science and Technology

November 6, 2025

Declaration of Authorship

I, Blair Alexander COLYER, declare that this thesis titled, “Combined Bayesian And Theoretical Analysis of Cancer Treatment Strategies and Incidence Data” and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

34 *“I am at the moment writing a lengthy indictment against our century. When my brain begins*
35 *to reel from my literary labors, I make an occasional cheese dip.”*

36 *-John Kennedy Toole, A Confederacy of Dunces*

37

CITY ST. GEORGE'S, UNIVERSITY OF LONDON

38

Abstract

39

School of Science and Technology

40

Doctor of Philosophy

41

Combined Bayesian And Theoretical Analysis of Cancer Treatment Strategies and

42

Incidence Data

43

by Blair Alexander COLYER

44 **Mathematical Modelling of Adaptive Cancer Therapy**

45 Maximum tolerated dose (MTD) chemotherapy has long been the standard of care for cancer
 46 treatment. An MTD protocol will rapidly deplete the population of treatment sensitive cells
 47 in the tumour, leaving only treatment-resistant cells behind, which will then grow without
 48 inhibition. A proposed alternative, called adaptive therapy (AT), aims to maintain a constant,
 49 high tumour burden, in a bid to leverage the evolutionary dynamics present in the tumour mi-
 50 croenvironment to the patient's advantage. Building on previous works, in collaboration with
 51 experimental biologists, we conduct a combined mathematical and experimental analysis of
 52 cancer adaptive therapy with epidermal growth factor receptor tyrosine kinase inhibitors. We
 53 present an analysis of both in-vitro cell assays and in-vivo mice experiments with PC-9 non-
 54 small cell lung cancer under treatment with erlotinib and osimertinib. From this analysis, we
 55 use a hierarchical Bayesian inference framework to parameterise an ODE model which we
 56 use to both explain phenomena and test hypotheses resulting from our data analysis. Com-
 57 bining this with analytical and numerical techniques which tell us the nature of the interaction
 58 between the cancer cells in vitro, we propose an adaptive therapy treatment protocol which
 59 should outperform a standard MTD protocol in terms of survival time. We use reproducible
 60 techniques to quantify model parameters and describe the nature of competition between
 61 these cells, informing future clinical trials based on these models.

62 **Mathematical Modelling of Bipolar Androgen Therapy**

63 Androgen deprivation therapy (ADT) has long been the standard of care for treatment of ad-
 64 vanced and metastatic prostate cancer. The aim of ADT is to deprive the hormone-dependent
 65 cancer cells of testosterone, inhibiting their proliferative capabilities or killing them outright.
 66 However, the cancer will eventually adapt to its new environment, reducing the efficacy of
 67 androgen ablation. When the disease progresses despite continued ADT, we refer to it as
 68 castrate-resistant prostate cancer (CRPC).

69 The proposed method for Bipolar Androgen therapy, or BAT, involves cycling between
 70 near-castrate levels of testosterone seen during ADT and supraphysiologic testosterone lev-
 71 els achieved via an intramuscular testosterone injection. This injection of large amounts of
 72 testosterone has been shown to inhibit the growth of castrate-resistant prostate cancer cells by

overloading their androgen receptors, and has been shown to restore sensitivity to ADT. As of yet, no mathematical models have been designed which try to describe this treatment. As such, we present an ODE model with minimal assumptions which aims to describe the BAT treatment protocols seen in the literature. By simplifying this model to a to make it more analytically tractable, we investigate the key facets of the treatment. We present a means to find the optimal treatment in a model where therapies alternate, and build upon this result to suggest a protocol for scheduling BAT which improves upon the protocol consistently seen in literature, supporting these claims with simulations.

Bayesian Modelling of Cancer Incidence Rates

Recent reports show that cancer incidence in adults under the age of 50 are increasing. Many different factors have been considered as potential drivers for these increases. It is possible, however, that processes such as screening for cancers such as breast or colorectal cancers, could be the driving force behind this increase in rate.

In order to investigate this, we adapt the hierarchical Bayesian inference workflow used for our investigations on adaptive therapy to study DEVCAN (probability of Developing or Dying of Cancer) statistics pulled from the US SEER (Surveillance, Epidemiology and End Results) data for cancer incidence. We used our inference workflow to fit two different Weibull-type models to this pan-cancer data set. The first model tests whether or not a true increase in incidence is occurring; the second fixes the incidence rate and tests whether or not the increased incidence can be explained by a decrease in detection lag time; the third tests both simultaneously to see which of the two might have a greater influence on results when both parameters governing these different factors are allowed to vary. We then test individual cancers to examine whether the trends in the pan-cancer data hold up for individual cases. We find that although both models do fit pan-cancer data, which would suggest that either hypothesis could explain the increase in incidence, the model fails to properly fit the data for individual cancers, suggesting that these Weibull type models may be too simple to capture the trends in these cases.

Altogether, we present in this thesis a methodological framework for analysing an array

101 of different cancer data sets and treatment strategies. We use statistical and theoretical mod-
102 elling to propose novel treatment strategies for both adaptive therapy and bipolar androgen
103 therapy. We demonstrate the versatility of our hierarchical bayesian modelling framework
104 by using it to parameterise models which aim to describe phenomena seen in *in vitro* and *in*
105 *vivo* experimental cancer data, as well as test hypotheses which can explain trends in cancer
106 incidence data by using our framework to parameterise simple models which validate these
107 hypotheses.

Acknowledgements

108
109 First and foremost, I would like to thank Robert Noble for his unwavering support through-
110 out the duration of my studies at City. I am continuously baffled by the freedom he has af-
111 forded me to pursue ideas, and the seemingly limitless level-headedness and patience that he
112 has demonstrated when working with me. I cannot think of an occasion to date where I have
113 come to him with an issue and he has not had a solution, and the list of opportunities he has
114 provided throughout my time spent under his guidance is as long as my arm. Special thanks
115 go to my collaborators at the IGMM, Dan, Liliana, Alexandre and Mohamad. There would
116 very literally be no thesis without the excellent experimental data that they have provided.
117 Special thanks also go to Yannick Viossat for his invaluable advice and thought-provoking
118 research discussions.

119 I also extend my gratitude to my colleagues in the office at City. It was long enough
120 ago now that I can admit that moving to London on my own with naught but a suitcase and
121 a rucksack was, shall we say, *disquieting*. Had I known that I would receive such a warm
122 welcome from everyone in the office, I would not have been so anxious. Although there are
123 too many to count, it would be remiss not to mention Ves in particular. His ability to talk so
124 easily to people terrifies me to my core, but I admit that I would have had a much tougher
125 time settling in in London had he been any different.

126 I would like to thank my Mum and Dad for their constant emotional, moral, and perhaps
127 most importantly, financial, support. I am sure that like many parents, when they imagined
128 retirement, they envisioned their fully grown deadbeat son thundering around their home
129 well into the late hours of the night, doubling their electricity bills and eating them out of
130 house and home. I feel honoured to have had the opportunity to help them realise this vision.
131 All jokes aside, they have been a constant source of encouragement and an ever-present
132 sympathetic ear, asking for little in return other than a decent cup of coffee every day. I
133 would like to thank Steph and Jas, who generously put me up and fed me when I visited
134 London in the last year of my thesis. I also thank my sisters, Jenny, Ashleigh, and Kirsty, and
135 give particular thanks to my twin sister Erin, and my sort-of-brother-in-law Efe. Our weekly
136 phone calls when I was living in London made me feel a bit less homesick than I might have.

137 Last, and also least, I thank Aron and Lachlan. They once learned that I had referred to
138 them as “my friends from uni”, to which they both took tremendous offence. I still do not
139 know why, frankly, because they are my friends, and we met at university – as far as I am
140 concerned, they had no grounds to complain. Naturally, I would never outright admit to them
141 that regardless of whether or not we are related, they are as much family to me as my own
142 flesh and blood. Nor would I admit that their consistent presence and constant support has
143 kept me sane over the past few years, because that would be very embarrassing. So many
144 thanks to my friends from uni.

Contents

146	Declaration of Authorship	i
147	Abstract	iii
148	Acknowledgements	vii
149	1 Literature Review	1
150	1.1 Adaptive Therapy	1
151	1.2 Bayesian Inference and <i>Stan</i>	16
152	1.3 Spatial Agent-Based Modelling of Tumour Evolution	32
153	1.4 Bipolar Androgen Therapy	57
154	2 Bayesian and Computational Modelling of Adaptive Therapy	74
155	2.1 Foreword	74
156	2.2 Introduction	74
157	2.3 Methods	76
158	2.4 Hypotheses and Results	90
159	2.5 Discussion	127
160	3 Using Mathematical and Bayesian models to Predict Adaptive Therapy Out-	
161	comes	133
162	3.1 Introduction	133
163	3.2 Methods	134
164	3.3 Results	136
165	3.4 Discussion	152

166	4 Mathematical Modelling of Bipolar Androgen Therapy	156
167	4.1 Introduction	156
168	4.2 Methods	157
169	4.3 Results	158
170	4.4 Discussion	176
171	5 Bayesian Modelling of Cancer Incidence Data to Investigate Increased Early	
172	Onset Cancer	181
173	5.1 Foreword	181
174	5.2 Introduction	181
175	5.3 Methods	185
176	5.4 Results	189
177	5.5 Discussion	199
178	6 Conclusion	207
179	A Appendices for Chapter 2	238
180	B Appendices for Chapter 3	255
181	C Appendices for Chapter 5	260

182 List of Abbreviations

AT	Adaptive Therapy
MTD	Maximum Tolerated Dose
NSCLC	Non-small Cell Lung Cancer
WT	Wild-type
OSM	Osimertinib
OR	Osimertinib Resistant
ER	Erlotinib Resistant
EGFR-TKI	Epidermal Growth Factor Receptor Tyrosine Kinase Inhibitor
QoL	Quality of Life
ODE	Ordinary Differential Equation
ABM	Agent-based Model
SABM	Spatial Agent-based Model
ADT	Androgen Deprivation Therapy
CADT	Continuous Androgen Deprivation Therapy
IADT	Intermittent Androgen Deprivation Therapy
BAT	Bipolar Androgen Therapy
CRPC	Castrate-Resistant Prostate Cancer
PSA	Prostate-Specific Androgen
AR	Androgen Receptor
AD	Androgen Dependent
AI	Androgen Independent
LH	Luteinising Hormone
LHRH	Luteinising-Hormone-Releasing Hormone
DHT	(5 α –)Dihydrotestosterone

DHEA(-S)	Dehydroepiandrosterone (Sulphate)
PFS	Progression-Free Survival
OS	Overall Survival

Dedicated to my Parents.

Chapter 1

Literature Review

1.1 Adaptive Therapy

Maximum tolerated dose (MTD) chemotherapy as a standard of care for cancer treatment arose from a 1970 study (1) in which acute lymphoblastic leukaemia was successfully cured in infants. MTD protocols aim to reduce the size of a tumour as much, and as quickly, as possible, whilst limiting the toxicity to the patient. Due to the toxic nature of the drugs, however, treatment will typically be to the detriment of the patient's quality of life. The full eradication of disease following MTD treatment of acute lymphoblastic leukaemia was a rare case (2). Outside of the handful of chemotherapy-curable cancers, MTD chemotherapy will often be followed by an emergence of therapeutic resistance following an initially promising decrease of tumour burden.

With a developing understanding of the biological interactions involved with the disease, Gatenby and colleagues (3) proposed an evolutionarily-informed alternative to the traditional MTD protocol. An MTD protocol will rapidly deplete the population of sensitive cells in the tumour, leaving only treatment-resistant cells behind, which will then grow without inhibition. The proposed alternative, called adaptive therapy (AT), aims to maintain a constant, high tumour burden, in a bid to leverage the evolutionary dynamics present in the tumour microenvironment to the patient's advantage. Maintaining a higher tumour burden preserves a large population of treatment-sensitive cells, which will compete with rarer, treatment-resistant cells for space and resources. Where MTD protocols are executed based on giving the patient as much treatment in as short a time window as possible, adaptive protocols instead aim for a dynamic treatment schedule, involving either adaptive alteration of drug dose

or timing of drug application, informed by ongoing treatment response. The objective of this treatment is to prolong the time until treatment-resistant cells are able to grow free of competitive inhibition by the more abundant treatment-sensitive cells, termed “competitive release” (4; 5), for as long as possible.

1.1.1 Phenotypes and Genotypes

It is important that before we discuss strategies of cancer treatment that we outline some key terminology that will be needed for later discussions.

An organism’s genotype is defined as the set of genes that the organism carries – specifically, it is the combination of alleles that that organism possesses for a specific gene. Genotype is inherited by the organism’s parent. An organism’s phenotype comprises its observable traits, and is influenced by both its genotype and its environment (6).

The traits a phenotype possess will vary as the system evolves via mutation, and these mutations can either be genetic or epigenetic. A genetic mutation is a change in the actual DNA sequence of the organism as a result of errors during DNA replication. These tend to be rare and often will not cause any major changes to an organism’s phenotype. Epigenetic mutations, on the other hand, are heritable changes which do not physically alter the actual DNA sequence but rather influence gene regulation. Genetic and epigenetic changes can be further subcategorised into three groups: driver mutations, passenger mutations, and deleterious mutations. A driver mutation, or beneficial mutation, is defined as a mutation with a selective growth advantage. A passenger mutation, or neutral mutation, is a mutation which has no advantage over its parent. A deleterious mutation is a mutation with a selective growth disadvantage.

1.1.2 Population Genetics

Population genetics depend on four processes: mutation (either genetic or epigenetic), selection (differential survival in the population due to phenotypic differences (7)), genetic drift (variation in an allele’s frequency in a population due to random sampling of cells (8)) and gene flow (transfer of genetic information between populations (9)). Cancer cells in the tumour microenvironment can be considered as evolving populations of somatic cells (that is,

cells in an organism other than gametes, germ cells, gametocytes or undifferentiated stem cells – in humans, these cells build up internal organs, skin, bones, blood and connective muscle tissues), which are subject to these four processes.

1.1.3 The Cell Cycle

Often, cancer treatment relies on the interruption of the cell cycle, and so it is important to understand the events that take place in the cell as it prepares to divide. It comprises four phases, called G1, S, G2 and M. In the G1 (gap 1) phase, the cell grows in size as it accumulates the energy and protein reserves needed for DNA synthesis. In the S (synthesis) phase, DNA replication occurs as two copies of each of the cell's chromosomes are created in the cell nucleus. In the G2 (gap 2) phase, the cell replenishes the energy and protein expended during the S phase to prepare for cell division. In the M (mitosis) phase, the cell will undergo a number of subprocesses which will organise the replicated chromosomes into the correct position and allow the original parent cell to split into two new daughter cells. The two daughter cells, if dividing, will re-enter the G1 phase and the process will repeat.

If a cell either does not need to divide, or cannot do so due to external factors, it can enter a fifth phase after the M phase, called the G0 phase, in which it will lie dormant, not making any preparation to divide. It will remain in this phase until some external factor pushes it into the G1 phase. While in this phase, we call the cells quiescent (10).

Cells can also undergo apoptosis, which is a process of genetically determined elimination of cells. This can happen for a number of reasons, and regular apoptosis is a normal part of cell turnover (in which older cells are replaced by newer, healthier cells). Cancers often exhibit irregular apoptosis or increased proliferation rates, where too few cells are removed from the population, or too many cells are produced too quickly (11). Cells can also exhibit senescence, a process in which the cell cycle is interrupted and proliferation is halted, but the cell persists instead of dying (12).

1.1.4 The Tumour Microenvironment

To consider a tumour as a mass of undesirable malignant cells is an oversimplification – in reality, tumours are akin to newly grown and harmful organs, capable of recruiting and

corrupting otherwise healthy cells surrounding them. The interactions between cells in a tumour and the cells surrounding them create what we call the tumour microenvironment (TME). Key players in the tumour microenvironment are discussed in (13).

Within the TME, we can find a variety of T and B lymphocytes. T lymphocytes such as CD8⁺ memory T cells and CD4⁺ helper T cells have been found to work towards killing tumour cells, and have been associated with good prognosis, whereas some CD4⁺ T cells have been associated with promoting tumour growth. B cell infiltration into the TME has been associated with good prognosis in some cancers, but these findings are inconsistent with animal studies – furthermore, some B cells, such as regulatory B cells, have been shown to increase tumour burden and inhibit tumour-specific immune responses. We also have natural killer (NK) and natural killer T (NKT) cells, which can indicate good prognosis in some cases, but have been found to lack their cell-killing function when found in the tumour stroma. Tumour-associated macrophages are abundant in most cancers, but are usually pro-tumourigenic, aiding in malignant cell migration, invasion, and metastases. Myeloid-derived suppressor cells are also found in the TME, but their variable phenotype can make them difficult to identify – they have, however, been found to have the potential to differentiate into tumour associated macrophages. Dendritic cells typically carry out antigen processing and presentation to promote immune response, but when present in the TME, they typically lack this function due to the inflammatory and hypoxic conditions, sometimes suppressing T cell activation. Adipocytes can also be present, and in some cancers, aid the recruitment of malignant cells via adipokine secretion and aid their growth by providing nutrients. Cancer-associated fibroblasts (CAFs) in the TME secrete growth factors, and in particular, epidermal growth factor, a protein which stimulates cell growth. They can also promote malignant cell growth, promote new vascular growth, recruit immune cells to promote inflammation, secrete Extracellular matrix-remodelling enzymes, and surround malignant cells with dense tissue that reduces the effectiveness of anti-cancer drugs.

When tumours develop, they can send signals throughout the TME which promote angiogenesis – the formation of new blood vessels – in the tumour. A number of cells present in the TME are associated with this process. Vascular endothelial growth factor (VEGF), fibroblast growth factor (FGF), platelet-derived growth factors (PDGF) and chemokines all stimulate

endothelial cells to promote the development of new vasculature in the TME. These new vessels behave abnormally – their heterogeneity, chaotic branching structure and leakiness leads to uneven oxygen, nutrient and drug distribution in the TME, which leads to hypoxic local conditions and can promote metastasis. Perivascular stromal cells provide structural support to these newly formed vessels, and low counts of these cells has been associated with increased chance of metastasis. Tumours can also initiate lymphangiogenesis (the creation of new lymph vessels) and increased lymphatic cell production. The interference in the lymphatic system by the tumour has been linked to the spread of malignant cells, TME modulation and alteration of immune response to the tumour.

The final key component of the TME is the extracellular matrix, or ECM. The ECM holds all of the surrounding cells together, and dictates the evolution and spread of cancers. How cancer cells adhere to the ECM will dictate how they grow, and how they can move in and out of the TME. The ECM contains angiogenic factors which promote new vessel development and chemokines which stimulate cell migration. CAFs promote further ECM development, which explains why it is that tumours are more rigid than healthy surrounding tissue. Malignant cells and CAFs can, however, induce ECM degradation and thus ECM remodelling, which further promotes the release of CAFs and angiogenic factors. This is a vicious loop, and leads to greater levels of irregular angiogenesis, inflammation, and eventually, metastasis.

Understanding which components of the TME aid tumour development and which ones hinder it will be key to designing future treatment strategies – using drugs which target factors that aid in tumourogenesis could prevent or even reverse tumour growth, or could help to manufacture an environment which selects for easy-to-treat phenotypes in the TME.

1.1.5 Evolutionary Principles

In the rare case that complete cure of disease is possible, a protocol with that objective (such as an MTD protocol) is clearly the optimal treatment option. But when considering treatment of disease through an evolutionary lens (14), cracks begin to form in curative strategies. In a situation where cancer cannot be cured – for example, in the case of metastatic cancers,

or other settings where cancer treatment is palliative rather than curative – an aim to cure is suboptimal, since treatment failure is all but guaranteed.

Containment strategies have been used to treat other diseases – a famous example is the treatment for HIV, in which virus replication is inhibited via the use of antiretroviral drugs, preventing further infection of healthy cells (15). The HIV genome has nine genes which encode fifteen viral proteins, and the genome contains somewhere between 9200 and 9600 nucleotides (16). Each of these genes can be targeted individually by different drugs, permitting containment. In an ideal world, a similar approach would be used to treat cancers, but the scope of the problem is much larger. In the human genome, which has around 22,000 protein-encoding genes, at least 350 of these genes exhibited recurrent somatic mutations in cancer (that is, a mutation characterised by a change in the DNA sequence) (17). Studies in mice have suggested that more than 2000 genes could contribute to cancer development (18). Furthermore, it is thought that to identify every single somatic mutation in a single cancer genome, one would need over 100,000,000,000 base pairs of the related DNA sequence (19). Given that cancer is a hugely diverse disease, an alternative approach to containment would be more practical until this specific information on each cancer is available.

A more pragmatic evolutionary strategy for cancer, then, will rely on exploiting the competition for space and resources, such as oxygen and glucose, between resistant cells and sensitive cells in the tumour bed. In (20), the authors discuss application of evolutionary principles to cancer therapy, and propose two general principles to guide our strategy:

- Resistant cell growth is subject to evolutionary forces, and so can be controlled by altering either resistant cell fitness or the fitness of competing cell populations;
- Any evolving population can only adapt to current local selection forces, and cannot anticipate future dynamics.

This second point is key to our strategy – although cells cannot predict their future conditions, a clinician, with sufficient information, can engineer these future dynamics of the environment to their advantage. Through this lens, we can see why MTD protocols for palliative treatments are a poor choice: by treating the patient with as much drug as possible, we select for resistant phenotypes and kill as many sensitive cells in the tumour as possible in order to shrink the tumour as much as possible. This creates an environment which

allows uncontrolled growth of the untreatable resistant phenotype, a phenomenon known as competitive release (4; 5). From this perspective, then, it would be better to try and engineer an environment in which sensitive cells are abundant enough to out-compete resistant cells. Leaving the tumour to grow without treatment would create the best possible environment for these sensitive cells to thrive, but this uncontrolled tumour growth would inevitably lead to patient death. As such, a middle ground must be struck, and this is the objective of adaptive therapy.

In adaptive therapy, the objective is to maintain a roughly constant tumour burden by using frequent drug applications at lower doses than one would use in an MTD protocol. A containment strategy should, in theory, force resistant cells into competition with their more abundant sensitive counterparts, whilst preventing unchecked growth of the tumour. In practice, one would begin with some low dose, and increase or decrease this dose as necessary, depending on whether the tumour has grown or shrunk. This delays the time until the emergence of resistance, and, in turn, delays time until treatment failure and death. On paper, this sounds simple enough, but in practice, a number of questions as to how one might best implement such a strategy arise – a recent review (21) outlines a number of open questions about containment strategies, touching on optimal adaptive dosing protocols, the roles of plasticity and homeostasis, the adverse effects of constant high tumour burden, and the competitive capabilities of resistant phenotypes, to name a few.

Resistance to treatment can either be pre-existing, or can develop gradually as a response to environmental conditions. Pre-existing drug resistance exists before treatment has been initiated, and can be caused by a number of factors: cells can possess genetic mutations which result in decreased drug responsiveness; treatment initiation in heterogeneous tumours can select for pre-existing resistance subpopulations, leading to eventual treatment failure; and activation of intrinsic defence pathways which protect cells from environmental toxins. All large solid tumours are thought to have some cells with intrinsic resistance (22). Resistant mutations can also occur during treatment. Activation of new oncogenes (that is, genes with cancer-causing potential) after treatment can become the primary driver gene in cancer cell replication. Driver mutations or alterations in gene expression levels can arise as a result of treatment. Alternatively, resistance can arise as a result of microenvironmental changes (23).

381 Acquisition of resistance may not necessarily incur any cost, but often, alteration of cell
382 phenotype will require some sort of expenditure of resources that will impact the cell's pro-
383 liferative fitness. The mechanism of this acquired resistance, and the disparity between the
384 fitness of sensitive cancer cells and resistant ones, will have an impact on the strategy used for
385 treatment. In (24), the authors conduct a theoretical analysis of various mathematical models
386 for tumour containment to show that clinical gains of tumour containment depend strongly
387 on the intensity of competition between sensitive and resistant cells.

388 In (20), the authors explain that in order to resist treatment (assuming treatment resistance
389 is not intrinsic), cells must upregulate their natural molecular defence mechanisms, such as
390 increased expression of P-glycoproteins (25), which was shown in (22) to use up a large
391 portion of a cell's energy budget that could otherwise be used for invasion or proliferation. In
392 optimal conditions where resources are abundant, changes to resist treatment may not inhibit
393 proliferation, but in *in vivo* conditions with limited resources, this energy cost tradeoff must
394 be made, as cells must prioritise survival. In other words, this fitness cost of resistance is
395 essential to the cell's survival during treatment, but when no treatment is applied, this fitness
396 cost will be to the cell's detriment, and can be exploited.

397 1.1.6 Common ODE Models

398 It will help when discussing pre-existing modelling studies of adaptive therapy later to briefly
399 outline some common models which are often used to investigate population dynamics. Typ-
400 ically, when considering evolutionary processes, we must have model terms which describe
401 birth and death as a bare minimum. Common extensions when modelling cancer might in-
402 clude treatment terms which inhibit growth or incur additional deaths and terms which sim-
403 ulate competition for space and resources between phenotypes.

404 1.1.7 Modelling growth

405 We refer to Gerlee's review (26), which presents an overview of tumour growth laws.

406 One of the most straightforward growth models is the exponential growth model. Here,
407 we assume that growth occurs uninhibited by any other factors in the system. Microorgan-
408 isms, like cells, will grow exponentially until they exhaust their supply of nutrients, or until

there is no more space for them to divide into, so we might use an exponential growth model when we wish to study a system in which we assume abundant nutrients and space. For a population N , the ODE model which describes exponential growth is as follows:

$$\frac{dN}{dt} = aN,$$

where a is the rate of exponential growth. If $a < 0$, then the population will decline exponentially, which we call exponential decay.

If we wish to model a population whose size is limited by some external factor, there are a number of options. One option is the logistic model, which assumes that the population growth rate decreases as the size of the population increases, until it reaches a critical point where the size of the population can no longer increase, often due to some limiting resource or lack of space. This critical population size is called the carrying capacity. As the cell number approaches this carrying capacity, the growth rate will approach zero. We model the logistic growth of a population N as

$$\frac{dN}{dt} = aN\left(1 - \frac{N}{K}\right),$$

where a is the growth rate and K is the carrying capacity.

Typically, when modelling tumours, we do not assume that the population of the tumour will grow indefinitely, since cells will have to compete for nutrients more and more as the population increases. In (27), the author was able to fit a sigmoidal growth curve, called a Gompertz growth curve, to tumour growth data. A Gompertz curve can simulate this increased competition with an increasing population, and is modelled as follows:

$$\frac{dN}{dt} = aN \log\left(\frac{K}{N}\right),$$

where a and K are as previous. Although the Gompertz and logistic models are functionally similar for high populations (at least, in that their population is limited by a carrying capacity), taking a look at the limits of these growth functions for small populations reveals a key distinction. As $N \rightarrow \infty$ in the logistic case,

$$\lim_{N \rightarrow 0^+} a(1 - \frac{N}{K}) = a,$$

431 and in the Gompertz case,

$$\lim_{N \rightarrow 0^+} a \log(\frac{K}{N}) = +\infty.$$

432 So we see that the growth rate for small populations in the Gompertz function is unbounded.

433 A modification of the Gompertz model was proposed by Wheldon in (28), called the
434 Gomp-ex model, which assumes that growth is exponential until a critical threshold, denoted
435 N_C , is reached, at which point growth becomes Gompertzian. This is modelled as follows:

$$\frac{dN}{dt} = \min(aN, \alpha N \log(\frac{K}{N})).$$

436 Von Bertalanffy (29) posited that net growth rate should arise from a balance of birth and
437 death, and observed that metabolic rates scale with total population size. The model proposed
438 is as follows:

$$\frac{dN}{dt} = aN^\gamma - bN.$$

439 Here, a is the growth rate, b is a constant death rate, and $0 < \gamma < 1$ is a parameter which
440 governs the decrease in growth rate as population increases. The second term in this model
441 is a standard way to include a constant death rate which occurs independently of external
442 factors like predation or disease (in the case of animals) or treatment (in the case of cancer).

443 1.1.8 Modelling Competition

444 The Lotka-Volterra model (30; 31; 32) can be used to describe the dynamics of two interact-
445 ing species. With an appropriate model setup, we can use these equations to describe:

- 446 • predator-prey dynamics, where the predator population's density depends on the abun-
447 dance of prey, and the prey population suffers when predators are abundant;
- 448 • mutualism or symbiosis, where both populations benefit from one another's presence;

- commensalism, where one organism benefits but the other is neither helped nor hindered;
- parasitism, where the parasite organism harms the host organism to its own benefit whilst relying on it for survival;
- amensalism, where one organism harms the other without benefitting itself, and;
- competition, where both organisms suffer due to one another's presence.

Competition is of particular interest when investigating adaptive therapy, since both resistant and sensitive cell populations must compete in the tumour microenvironment for space and resources. A simple competitive Lotka-Volterra model, in which species X and Y compete with one another, might look like:

$$\begin{aligned}\frac{dX}{dt} &= \alpha X - \beta XY, \\ \frac{dY}{dt} &= \gamma Y - \delta XY,\end{aligned}$$

where α and γ are the birth rates of X and Y , respectively, and β and δ describe the antagonistic effect of competition on X and Y , respectively.

Models including a carrying capacity, like the logistic, Gompertz and Gomp-ex models above, can also be used in models with multiple populations to implicitly model density-dependent competition. If we have two organisms sharing one nutrient source in one shared space, then we might have it that they both share a carrying capacity. In the case of two organisms exhibiting logistic growth, X and Y , this might look like the following:

$$\begin{aligned}\frac{dX}{dt} &= \alpha X \left(1 - \frac{X + Y}{K}\right), \\ \frac{dY}{dt} &= \gamma Y \left(1 - \frac{X + Y}{K}\right).\end{aligned}$$

1.1.9 Modelling Treatment

There are several different ways we might model treatment in a tumour growth model. Typically, a treatment will act in two ways: either it will directly kill living cancer cells, or it

will inhibit cell proliferation to prevent further growth. We can incorporate these effects by having some term act as a multiplicative value between 0 and 1 that depends on the concentration of drug present, or by simply multiplying the level of treatment by the density and removing this from the population. In either case, our “treatment term” will have no effect on our cancer cells when the drug is absent, and as the drug concentration increases, the level of inhibition conferred by the treatment term will also increase. A general model which describes a treatment which kills cells might look like the following:

$$\frac{dN}{dt} = \alpha N - f(C)N,$$

where $f(C)$ is a function governing the effect of treatment, depending on the level of treatment, C . Likewise, if treatment inhibits growth, then we might have a growth model resembling the following:

$$\frac{dN}{dt} = \alpha g(C)N,$$

where $g(C)$ is a function which governs the level of inhibition on growth. The forms of the functions f and g will vary depending on the model, but there are a few popular options. The Hill equation (33) is used in pharmacology to describe drug-response relationships, and takes the general form

$$f(C) = \frac{\beta}{1 + \left(\frac{IC_{50}}{C}\right)^H},$$

where β is the maximum level of inhibition, IC_{50} is the half-maximal inhibitory concentration (which, given that the Hill function is sigmoidal, acts as the inflection point of the curve), and H is the Hill parameter which governs the shape of the curve. If we take $H = 1$ in the above equation, we have a special case called the Michaelis-Menten function (34), which takes the form

$$f(C) = \frac{\beta C}{IC_{50} + C}.$$

1.1.10 Adaptive Therapy Models

Some of the first models which advocated for a containment strategy in favour of MTD were published in 1992 by Martin *et al.* In (35), the authors consider density-dependent models with mutation and two types of cells – sensitive and resistant – to investigate the optimal tumour size needed to maximise survival time. Density-dependent Gompertz growth and Lotka-Volterra models with logistic and exponential growth are considered, with constant dose treatment or treatment delays used. Treatments which maintain a high tumour burden were optimal for Gompertz growth (at least doubling survival time) and were close to optimal with logistic and exponential growth (although the authors found that tumour burden during therapy in the latter cases had little impact on survival time). This study was extended to combination chemotherapy in (36) with similar findings.

The term “Adaptive therapy” was coined in Gatenby *et al.*’s 2009 paper (3). In this paper, tumour growth is modelled using a power law in the absence of treatment, and in the presence of treatment, a modified version of this power law with a constraint that describes treatment. Their analytical findings suggested that in tumours with resistant subpopulations, survival time is maximised when tumours are treated with the aim of stabilising, rather than curing. They show that the optimal treatment will be one which adapts to maintain a stable sensitive cell population, which suppresses the growth of resistant subclones.

In (37), the authors use both non-spatial models of coupled ODEs and a hybrid cellular automaton (which we cover in detail in a later section) to model adaptive therapy. They present an array of coupled ODE models, using both exponential and Gompertzian growth laws, and model competition between resistant and sensitive phenotypes using a frequency-dependent relative fitness function which inhibits the growth of resistant cells. Treatment is assumed to kill cells instantly when applied. They include a variation of the Gompertzian growth model which allows carrying capacity to vary over time, depending on tumour vascularisation. Their conclusion was that the benefit of AT protocols over MTD protocols will be limited unless the fitness differential between sensitive and resistant cells is large. Their hybrid cellular automaton model HCA predicted that spatial structure was a key factor in the efficacy of adaptive therapy; experimental results confirmed this hypothesis.

517 In (38), the authors utilise an off-lattice agent-based model to investigate spatial evolu-
518 tionary dynamics of different treatment schedules. Space limits growth of cells, and if no
519 space to proliferate is available, cells enter quiescence – if a cell can find space to divide
520 without overlap, it does so. Death due to turnover is excluded and resources are assumed
521 to be abundant when cell density is below carrying capacity. Cell death depended on drug
522 dose and cell sensitivity. Cell sensitivity was continuous, where the degree of sensitivity is
523 linearly dependent on the individual cell's cell cycle time. Drug dose for adaptive therapy
524 varies depending on population size in the tumour. Phenotypic drift and cell migration were
525 also included. Their model findings found that MTD treatment would cure sensitive tumours,
526 but tumours would recur if resistance was present. They propose a dose modulation protocol
527 for a majority of tumours, with a drug holiday strategy for more invasive tumours.

528 In (39), the authors again validate the AT hypothesis by testing it using a hybrid cellular
529 automaton. A phenotype's degree of resistance was assumed to be inversely correlated to
530 its potential to proliferate. Their model incorporates diffusion of oxygen, glucose and pH
531 buffers through the simulation grid. Cells in the grid could metabolise glucose and oxygen
532 into ATP, lactic acid and CO₂. Cell division depended on these diffusible factors and the by-
533 products of cells metabolising them. The model is split into two-phases; tumour free growth
534 and therapy. During the former, cells divide depending on the previously outlined factors,
535 their proliferative capacity, and available space; during the latter, the cells can be killed not
536 only by the drug itself, but can also fall victim to competition with the glucose transporters
537 used by the drug, meaning that treatment could also kill cells via a lack of available glucose.
538 In (40), a similar model is parameterised with *in vitro* data, and an adaptive treatment protocol
539 in which dose was increased (or decreased) if the tumour increased (or decreased) in size was
540 tested, and it was found that it could outperform MTD protocols.

541 In Zhang *et al*'s 2017 paper (41), the authors utilise a Lotka-Volterra model of three
542 different phenotypes to investigate patient specific data from a pilot clinical trial of adaptive
543 therapy with Abiraterone used to treat metastatic prostate cancer. We investigate this model
544 in a later section of this review.

545 Smalley *et al* (42) design a two-compartment ODE model with sensitive and resistant

cells to investigate the dynamics of BRAF inhibitor treatment of melanoma. This Lotka-Volterra model was calibrated with *in vivo* mouse xenograft data, and the authors used an implicit filtering algorithm to determine optimal parameters for their model. Using these parameters, tumour trajectories were predicted based on previous treatment decisions, and the results used to inform future treatment decisions for each mouse. The personalised adaptive protocols for each mouse were found to outperform continuous and fixed intermittent therapies.

Brady-Nicholls *et al* (43) developed an ODE model which incorporates stem cells, non-stem prostate cancer cells, and prostate-specific androgen (or PSA – a biomarker used to estimate the progression of prostate cancers, discussed in a later section) to examine treatment of metastatic castrate-resistant prostate cancer with abiraterone. They also consider a model extension which considers metastatic burden. The model was calibrated and validated against patient PSA responses to adaptive treatment with abiraterone, and then used to predict treatment responses. The model was able to accurately describe patient-specific clinical data and identify potential resistance mechanisms.

In a paper by Strobl *et al* (44), the authors investigate the conditions under which a cancer might respond to adaptive therapy. Using a Lotka-Volterra model of competing sensitive and resistant cells, they examine the importance of a fitness cost of resistance in an adaptive therapy protocol. They found that larger tumours close to their carrying capacity do not require a fitness cost for adaptive therapy to be successful, but smaller tumours growing further from their carrying capacity may require a fitness cost of resistance to see any gains from adaptive therapy. They also present evidence that cell turnover may be an important factor in modulating the impact of a fitness cost of resistance. They test these hypotheses using data from a cohort of prostate cancer patients undergoing intermittent androgen deprivation therapy.

A theoretical analysis of adaptive therapy is carried out in a paper by Viossat and Noble (24). Considering an array of different mathematical models, they outline the model conditions under which we will see a clinical benefit from aiming to contain a tumour, rather than using MTD protocols. A particularly important finding from this analysis was that a fitness cost of resistance is not necessary to see a benefit from adaptive therapy, but outcomes are better when initial resistance is rare. One finding is that if no competition between resistant

and sensitive cells takes place, then containment strategies will likely be outperformed by MTD strategies. If cells are not fully resistant, then containment alone may not be optimal, since an MTD strategy could be beneficial once partially resistant cells escape competition. They find that containment is most successful when tumours are stabilised at a size which is as close as possible to a theoretical maximum size threshold. Clinically, this is impractical, but they present results which show that containment strategies can still outperform MTD strategies when tumours are contained at sizes safely below this threshold.

1.2 Bayesian Inference and *Stan*

Here, we review the theory behind Bayesian inference and investigate the algorithms used in our chosen inference software. In particular, we shall look at the theory behind Bayesian inference and an extension of this method called hierarchical Bayesian inference, Markov chain Monte Carlo algorithms and the Metropolis-Hastings update algorithm, the Hamiltonian Monte Carlo algorithm, and some algorithms and methods used in the inference software of choice, *Stan* (45). We refer primarily to Gelman *et al*'s book *Bayesian Data Analysis* (46) for our explanations of Bayesian Inference and the algorithms which we intend to use – other citations are stated where relevant.

1.2.1 Bayesian Inference

In broad terms, the objective of Bayesian inference is to take a dataset, fit a probabilistic model to that dataset, and then describe this model using newly generated probability distributions for the model parameters, which, if desired, can permit us to generate predictions for future events based on the model. Our particular focus lies with parameter estimation. This process can be broken into three steps. First, we establish our probabilistic model, which contains a joint distribution for any observable quantities pertaining to our problem, and which is designed with both our data and the nature of the problem we are investigating in mind. Next, we use our data to condition this model and generate a posterior distribution estimate. Finally, we test this posterior distribution to determine whether the model fits the data in an appropriate way by comparing it to our model assumptions and deciding whether the conclusions we can take from this posterior distribution are realistic.

Specifically, we want to investigate an unknown parameter, θ , based on our observed data, y . To do this, we require a joint probability distribution for both θ and y . This joint probability distribution, $p(\theta, y)$, can be written as a product of two distributions: the prior distribution, $p(\theta)$, which should ideally encapsulate any information about what we expect the parameter estimates to be based on initial data analysis, and the sampling distribution $p(y|\theta)$, the likelihood of our data y given fixed model parameter θ . So we have that

$$p(\theta, y) = p(\theta)p(y|\theta).$$

We wish to find the posterior density that describes the parameters given the data, $p(\theta|y)$, in terms of the prior, sampling and joint probability distributions. We can do this using Bayes' theorem directly:

$$p(\theta|y) = \frac{p(\theta, y)}{p(y)} = \frac{p(\theta)p(y|\theta)}{p(y)},$$

$$\text{where } p(y) = \int p(\theta)p(y|\theta)d\theta.$$

We can rewrite this expression using an alternate form of Bayes' theorem, since we know that in our case, the data, here represented by the event $p(y)$, is not variable and so can be fixed. We can thus say that the posterior distribution is proportional to the product of the prior distribution and the sampling distribution or likelihood:

$$p(\theta|y) \propto p(\theta)p(y|\theta).$$

This is the mathematical backbone which underpins the inference we carry out in the project. Of course, in actuality, carrying this inference out is more involved since we are working with multiple parameters and datasets, so we shall look at the algorithms we can use to find the posterior distributions of interest.

1.2.2 Hierarchical Bayesian Inference

If we have several related datasets, a hierarchical Bayesian model is a good choice. Rather than inferring a set of parameters for each data set individually, it is often useful to try

and pool all of the data together to estimate parameters. This is precisely how hierarchical Bayesian inference works – we infer a set of parameters for each individual dataset, called the group-level parameters. These are all drawn from a shared parameter distribution which gives an overview of all of the datasets – this distribution allows us to infer what we call the hyperparameters. This is very useful for tackling variability between replicates in an experiment: for example, if we have several replicates of an experiment which shows the growth of a cell line over time, they will invariably exhibit different growth rates – with a hierarchical model, we obtain a replicate-specific growth rates as well as a catch-all distribution that reflects the variability between the replicates which only influences the data via these replicate-specific growth rates.

Since we are introducing a new parameter, we must accordingly adjust the mathematics we use to find our desired posterior distribution. We need to infer both the group-level parameter θ (typically there will be several of these, but to keep the explanation, we shall pretend for the moment we only have one) and the hyperparameter ψ . As such, our posterior distribution of interest is now comprised of a vector, (ψ, θ) , and the joint prior distribution is

$$p(\psi, \theta) = p(\psi)p(\theta|\psi),$$

with corresponding joint posterior distribution of

$$p(\psi, \theta|y) \propto p(\psi, \theta)p(y|\psi, \theta).$$

The distribution $p(y|\psi, \theta)$ only really depends on θ , since the hyperparameter ψ only influences the data via θ , and so we can simplify this joint posterior:

$$p(\psi, \theta|y) = p(\psi, \theta)p(y|\theta).$$

The challenge with this new model is that a prior distribution for ψ must also be chosen. This can be a fairly involved task and the method will vary from model to model, and so we omit the discussion pertaining to this selection, and discuss this selection in more detail in the methods.

1.2.3 Markov Chain Monte Carlo

Markov chain Monte Carlo (MCMC) simulations sample random values of θ from distributions we can approximate and then adjust these samples such that they better resemble the desired posterior distribution. The distribution of the sampled values depend only on the last values sampled, and so the drawn samples form a Markov chain.

As a reminder, a Markov chain is a sequence of random variables $\theta^1, \theta^2, \dots$, where, for any time t , the distribution of θ^t depends only on the value of the previous variable θ^{t-1} (note that these indices are not powers, but labels, unless otherwise stated). The fact that predictions can be made using only the current state, and are independent of the stochastic process's history, is called the Markov property.

We require some terminology and definitions to understand our objective when executing MCMC algorithms. A Markov chain is called Ergodic if, during the random walk, there is positive probability that we can reach any state from any other. In more precise terms (which we require for a later proof sketch), a Markov chain is ergodic if and only if it has at most one recurrent class and is aperiodic. We unpack these terms one at a time. Two states i and j in a Markov chain are recurrent if we can access state i from state j , and vice versa, with nonzero probability. A recurrent class is a subset of states in the Markov chain which are all recurrent with one another. If a Markov chain has only one recurrent class, then every state is recurrent with every other one. We can call a Markov chain in which we can reach any chosen state from any other in a finite number of steps irreducible. A state is considered periodic if we can only return to that state at times which are multiples of an integer larger than 1 – so an aperiodic state is one which is not periodic. An aperiodic Markov chain, then, is one in which every state is aperiodic. As stated, an Ergodic Markov chain satisfies both of these properties.

A key distribution when implementing an MCMC algorithm is the stationary distribution of the Markov chain. If a Markov chain starts with its stationary distribution, then the marginal distribution of any state at any time point will be the stationary distribution. Ergodic Markov chains have a unique stationary distribution. Proving that any given Markov chain is Ergodic is not a straightforward task, and so we aim to have the distribution we obtain from our MCMC process approximate the stationary distribution within some reasonable margin.

If the posterior distribution yielded from our algorithm is within this margin, we say that the Markov chains converge. How we determine whether or not the chains have converged will be discussed in detail later.

We refer to Geyer (47) for a more specific explanation, starting with the theory behind ordinary Monte Carlo simulation. Say we wish to calculate the expectation of some real-valued function g which cannot be approximated by some exact method. We denote this expectation $\theta = E\{g(X)\}$. Supposing we can somehow simulate the values in the state space $X_1, X_2, \dots$ independently and identically distributed (IID) whilst keeping the same distribution. Then we define

$$\hat{\theta}_n = \frac{1}{n} \sum_{i=1}^n g(X_i).$$

Allowing $\gamma_i = g(X_i)$, we have that all of γ_i are IID, with mean θ and the variance

$$\sigma^2 = \text{var}\{g(X)\},$$

where $\hat{\theta}_n$ is the sample mean of γ_i . The Markov chain central limit theorem (MCCLT, which varies from the standard central limit theorem, and is discussed in (48)) tells us that

$$\hat{\theta}_n \approx N\left(\theta, \frac{\sigma^2}{n}\right),$$

with variance estimated by

$$\hat{\sigma}_n^2 = \frac{1}{n} \sum_{i=1}^n (g(X_i) - \hat{\theta}_n)^2,$$

the empirical variance of γ_i .

MCMC plays out in a similar fashion, but our standard error term changes due to the stochastic dependence in the Markov chain. Supposing the sequence $X_1, X_2, \dots$ is a stationary Markov Chain whose prior distribution is the same as the distribution of the state space X , and assuming the MCCLT holds, we have

$$\sigma^2 = \text{var}\{g(X_i)\} + 2 \sum_{k=1}^{\infty} \text{cov}\{g(X_i), g(X_{i+k})\}.$$

This holds only for stationary Markov chains, which we never actually use for MCMC simulations (if we knew the stationary distribution already, we could just do an ordinary Markov chain simulation rather than MCMC). It is difficult to actually calculate this variance in the first place, and in practice we only ever estimate it via simulation.

For our purposes, we shall create multiple independent Markov chains starting at the point θ^0 . The next value in each chain is drawn from something called the transition distribution, $T_t(\theta^t|\theta^{t-1})$, $\forall t > 0$, which is conditional only on the previous draw in the chain. The key to successfully designing an MCMC algorithm is in the selection of this transition probability, since how this is done will dictate whether or not each chain in the process will converge to a unique stationary distribution which is the posterior distribution of interest. This is not a straightforward task, and so rather than sampling θ from the posterior directly, we use iterative methods which allow us to draw samples from a distribution which will resemble the posterior more and more as time increases.

1.2.4 The Metropolis Algorithm

In our inference for the project, we implement a Metropolis-Hastings Hamiltonian Monte Carlo algorithm in *Stan* (45). Metropolis-Hastings is a generalised variant of the Metropolis Algorithm, and so we shall cover this first for context. The Metropolis algorithm itself is just a modified random walk which checks each newly sampled θ^t value at each iteration to determine whether or not it better resembles the stationary distribution than the previous value. It was originally presented in (49), but we refer to (46) for the explanation of the algorithm and sketch of the proof that the algorithm works. It proceeds as follows:

1. Sample the first value, θ^0 , which must satisfy $p(\theta^0|y) > 0$, from the starting distribution $p_0(\theta)$ (our prior distribution, which will ideally be data-motivated, but can be a vague estimate if no data is available to estimate this distribution).
2. For $t = 1, 2, \dots$;
 - (a) Draw a proposed value θ^* from a proposal (or jumping) distribution at time t , which we call $J_t(\theta^*|\theta^{t-1})$. This jumping distribution must be **symmetric** – that is, $J_t(\theta^a|\theta^b) = J_t(\theta^b|\theta^a) \forall \theta^a, \theta^b, t$.

(b) To determine whether this proposed value is accepted, calculate the ratio of the densities,

$$r = \frac{p(\theta^*|y)}{p(\theta^{t-1}|y)}.$$

(c) Update the value for θ^t depending on the following:

$$\theta^t = \begin{cases} \theta^* & \text{with probability } \min(r, 1) \\ \theta^{t-1} & \text{otherwise.} \end{cases}$$

Hence, our transition distribution $T_t(\theta^t|\theta^{t-1})$ depends on a point mass at $\theta^t = \theta^{t-1}$ combined with the weighted proposal distribution $J_t(\theta^t|\theta^{t-1})$. To actually execute this algorithm, we need to be able to calculate r , which requires approximation of stochastic integrals. We also need to be able to draw a random value from the proposal distribution and another uniform random value to determine whether or not the new value is accepted in step (2c).

Gelman *et al* (46) sketch the proof that this algorithm yields a sequence of values which will converge to the Markov Chain's uniform stationary distribution as desired. Gelman *et al* begin by demonstrating that the sequence we simulate is a Markov chain with unique stationary distribution, and then showing that this stationary distribution is the posterior distribution we desire.

We first show that our Markov chain is irreducible, aperiodic, and recurrent. Aperiodicity and recurrence are satisfied for a random walk on most proper probability distributions (bar a few trivial exceptions). Irreducibility will hold as long as the chosen proposal distribution J_t allows us to jump to any and all states in finite time with a positive probability.

Now, we consider the Metropolis algorithm at time $t - 1$, having drawn a value θ^{t-1} from target distribution $p(\theta|y)$. Consider two points drawn from this target distribution, θ^a and θ^b satisfying $p(\theta^b|y) \geq p(\theta^a|y)$. The unconditional probability density of a transition from θ^a to θ^b is

$$p(\theta^{t-1} = \theta^a, \theta^t = \theta^b) = p(\theta^a|y) J_t(\theta^b|\theta^a) \frac{p(\theta^b|y)}{p(\theta^a|y)},$$

743 where the fraction on the right hand side is the ratio of densities, r , seen in step (2b) in the
 744 above algorithm. Our conditions on θ^a and θ^b ensure an acceptance probability of 1 here,
 745 since $r > 1$. Now, the unconditional probability density of a transition from θ^b to θ^a (the
 746 opposite direction of the previous transition), is

$$p(\theta^{t-1} = \theta^b, \theta^a = \theta^t) = p(\theta^b|y)J_t(\theta_a|\theta_b)\frac{p(\theta^a|y)}{p(\theta^b|y)}.$$

747 Cancelling the denominator of the ratio of densities, we obtain

$$p(\theta^a|y)J_t(\theta^a|\theta^b),$$

748 which is the same as the unconditional probability density of a transition from θ^a to θ^b since
 749 we require J_t to be symmetric for the Metropolis algorithm to work. Hence the probability
 750 of transitioning from θ^a to θ^b and from θ^b and θ^a are equal, and so θ^t and θ^{t-1} have the same
 751 marginal distributions. Hence $p(\theta|y)$ is the stationary distribution for the Markov chain of θ ,
 752 and since this Markov chain is aperiodic, recurrent and irreducible, this stationary distribution
 753 is unique.

754 1.2.5 The Metropolis-Hastings Algorithm

755 Hastings proposed a generalisation of the Metropolis algorithm in (50), which relaxes the
 756 requirement for the proposal distribution to be symmetric, compensating for this by adjusting
 757 the ratio r . The algorithm is as follows:

- 758 1. As before, sample the first value, θ^0 , which must satisfy $p(\theta^0|y) > 0$, from the starting
 759 distribution $p_0(\theta)$.
- 760 2. For $t = 1, 2, \dots$;
 - 761 (a) Draw a proposed value θ^* from a proposal distribution at time t , which we call
 762 $J_t(\theta^*|\theta^{t-1})$. This jumping distribution **need not be symmetric**.
 - 763 (b) Calculate the ratio of densities,

$$r = \frac{p(\theta^*|y)/J_t(\theta^*|\theta^{t-1})}{p(\theta^{t-1}|y)/J_t(\theta^{t-1}|\theta^*)},$$

which has been replaced by a ratio of ratios to account for the fact that J_t need not be symmetric (if it is symmetric, this simplifies to the original ratio of densities presented in the Metropolis algorithm). This ratio is always defined since jumping from θ^{t-1} to θ^* is only possible if both of the terms in the denominator are nonzero.

(c) Update the value for θ^t , as before, depending on the following:

$$\theta^t = \begin{cases} \theta^* & \text{with probability } \min(r, 1) \\ \theta^{t-1} & \text{otherwise.} \end{cases}$$

Proving that this algorithm yields a Markov chain whose unique stationary distribution is the desired posterior distribution follows in the same fashion as it did for the Metropolis algorithm. The only key difference is that we define θ^a and θ^b in such a way that they now satisfy $p(\theta^b|y)J_t(\theta^a|\theta^b) \geq p(\theta^a|y)J_t(\theta^b|\theta^a)$, accounting for the altered ratio of densities used in the new algorithm.

1.2.6 Assessing Convergence

In order to have any confidence that our inference was correct, we must determine ourselves whether or not the estimated posterior distribution has indeed converged to the target distribution we desire. We do this in a number of ways. The first is discarding early iterations in a simulation run. This is really best practice, since we do not want our final estimated posterior distribution to be too heavily influenced by our prior estimate for the target distribution. There is no fixed answer for how many iterations we discard, and often this decision has to be made once we have already tried to run the simulation at least once, but if we are being conservative, as many as half of the iterations can be discarded. We usually call these discarded iterations the “warm-up” or “burn-in”.

There are two values which we can calculate after the simulation has run which give us a simple and effective means of assessing convergence: the $\hat{R}$ value and n_{eff} , the effective sample size (ESS). The former is a value which tells us whether the chains have mixed properly whilst also telling us if these chains are tracing a common distribution – we discuss shortly what this means at more length. The latter is a summary statistic whose necessity

arises from the fact that MCMC processes will often cause each draw in the iteration to become autocorrelated. In an ideal world, all of our samples would be independent, but we will seldom achieve this with an iterative MCMC method. As such, we have the ESS as a sort exchange rate between these ideal independent samples and the correlated dependent samples in our Markov chain. If the Markov chain draws are highly correlated, then the ESS value will be lower than that of a different Markov chain with the same number of draws which were weakly correlated (51).

We first discuss the purpose and derivation of the first summary statistic, $\hat{R}$, called the potential scale reduction factor. This is a useful, dual-purpose statistic that shows us immediately whether or not the chains have mixed, as well as the stationarity of the chains. Mixing of chains refers to the mixing time of the Markov chain. The mixing time tells us how long the Markov process will take before it reaches a distribution which resembles the target distribution (within some chosen margin of error). Chains that are well-mixed will quickly converge to the target stationary distribution, and chains that mix poorly will take a long time to reach it (52). Stationarity is more straightforward – this just describes whether or not the draws taken resemble the target distribution. To check both, we split the chains after the warmup in two, and then analyse the sub-chains. If the chain is properly mixed, the sub-chains will also mix; similarly, if the sub-chains both traverse the same distribution, then the overall chain can be considered stationary. After splitting, we have m sub-chains of length n .

To check the mixing and stationarity of all of the chains together, we assess the variances between different chains, denoted B , and the variance of sub-chains within each individual chain, denoted W , simultaneously. For an estimated parameter ϕ , we label each draw in the simulation ϕ_{ij} , where $1 \leq i \leq n$ and $1 \leq j \leq m$, and compute both B and W as follows;

$$B = \frac{n}{m-1} \sum_{j=1}^m (\bar{\phi}_{\cdot j} - \bar{\phi}_{\cdot\cdot})^2, \text{ where } \bar{\phi}_{\cdot j} = \frac{1}{n} \sum_{i=1}^n \phi_{ij}, \bar{\phi}_{\cdot\cdot} = \frac{1}{m} \sum_{j=1}^m \bar{\phi}_{\cdot j},$$

$$W = \frac{1}{m} \sum_{j=1}^m s_j^2, \text{ where } s_j^2 = \frac{1}{n-1} \sum_{i=1}^n (\phi_{ij} - \bar{\phi}_{\cdot j})^2.$$

The between-sequence variance, B , contains the factor of n since it is based on the variance of the within-sequence variance means, $\bar{\phi}_{\cdot j}$, which are each an average of n values. We then

816 use W and B to estimate the marginal posterior variance of ϕ with the following weighted
 817 average:

$$\widehat{\text{var}}^+(\phi|y) = \frac{n-1}{n}W + \frac{1}{n}B.$$

818 This is an overestimate of variance if the prior distribution is sufficiently wide, but unbiased
 819 under a stationary distribution (or when $n \rightarrow \infty$). When n is finite, however, W should
 820 underestimate this variance, since the individual chains have not had enough time to stretch
 821 across all of the target distribution yet, meaning there is less variability in the sampled values
 822 – as such, as $n \rightarrow \infty$, $W \rightarrow \text{var}(\phi|y)$ in expectation. We finally have all of the components
 823 needed to derive the potential scale reduction statistic, $\hat{R}$, which will estimate the factor by
 824 which the scale of the current distribution for ϕ might need to be reduced if we continued the
 825 simulation in the limit $n \rightarrow \infty$. We have

$$\hat{R} = \sqrt{\frac{\widehat{\text{var}}^+(\phi|y)}{W}},$$

826 which has the property that $\hat{R} \rightarrow 1$ as $n \rightarrow \infty$. For our purposes, if, in our inference, we have
 827 that $\hat{R} > 1.01$, then we consider the inference unreliable. This is an indication that we must
 828 analyse the trace plots of each chain to see whether or not this is an issue of poor mixing or a
 829 lack of stationarity. To try and improve the $\hat{R}$ value, there are a few options. We can run the
 830 simulation for more iterations, but if the model is configured poorly, then this likely won't
 831 be a good solution unless we increase the number of iterations by a huge amount, which
 832 may take a very long time if our model is not as efficient as possible. Often, the solution
 833 is to reconfigure the model itself to try and improve the model geometry. The objective is
 834 to try and craft a Markov chain which is Ergodic, since these Markov chains behave very
 835 nicely and have unique stationary distributions. In practice, making a Markov chain with this
 836 property for complicated problems is very difficult, and so it is often easier to modify our
 837 prior distributions or the type of probability distributions we draw values from to improve
 838 model geometry instead. We can also use transformations to improve model geometry, but
 839 these transformations will be dependent on the current model.

840 To calculate the ESS, we first observe that in an ideal simulation in which all n draws

were completely independent, the value B would be an unbiased estimate of the posterior variance $\text{var}(\phi|y)$, giving us mn independent simulations in m sequences. Of course, we have already mentioned that MCMC simulations tend to lead to autocorrelation of ϕ values in the sequences, meaning $B > \text{var}(\phi|y)$ in general.

We need some idea of how much autocorrelation there is in the sequences – we call the sum of these autocorrelations ρ . To calculate this, we start by computing the variogram (the degree of spatial dependence of the process), V_t , at each lag t (where lag is the distance between two terms in the sequence). This gives us

$$V_t = \frac{1}{m(n-t)} \sum_{j=1}^m \sum_{i=t+1}^n (\phi_{i,j} - \phi_{i-t,j})^2.$$

The autocorrelation for some lag t is then estimated as

$$\rho_t = 1 - \frac{V_t}{2\widehat{\text{var}}^+},$$

where $\widehat{\text{var}}^+$ is as defined above. Ideally, we would simply take the sum of all of these correlations to get our ESS, but for large lag values, the correlation will be too noisy (noise will naturally be higher the more values in the sequence we include in our estimate). To offset this noise, we instead take a partial sum, starting from no lag, and ending when the sum of the ρ values for two successive lags t' and $t+1'$ is negative. This partial sum is used to finally estimate the ESS,

$$\hat{n}_{\text{eff}} = \frac{mn}{1 + 2 \sum_{t=1}^T \hat{\rho}_t},$$

where T is the first odd positive integer for which $\hat{\rho}_{T+1} + \hat{\rho}_{T+2} < 0$, and $\hat{\rho}_t$ is calculated as outlined above. What counts as an acceptable value for n_{eff} will depend on the number of iterations in the simulation. As few as 10 can be considered acceptable, but if we have a large number of iterations, this could be a negligible ESS. Naturally, we want this number to be as close as possible to the number of iterations after warmup.

It is worth stressing that even if our checks on $\hat{R}$ and the ESS pass, all this really allows us to do is declare that our chains *approximately* converge. It is possible that our choice of prior distribution makes it impossible for our sequence to explore large portions of the

target distribution, for example. We emphasise that the computation of these values are just diagnostic checks, not hypothesis tests, and that at best, we can only declare that the sequences *appear* to mix, and that they *appear* to be stationary. It is important to always be sceptical of a model output and to scrutinise it as much as possible – just because the output of our simulation looks good, does not necessarily mean that it is completely correct or that we have perfectly reconstructed the target posterior distribution.

1.2.7 Hamiltonian Monte Carlo

The Metropolis-Hastings algorithm can be inefficient for high-dimensional target distributions such as the ones we are interested in estimating, due to the way that the random walks navigate through the parameter space locally between each iteration. The Hamiltonian Monte Carlo (HMC) algorithm (53; 54) aims to correct some of this inefficiency by using ideas from Hamiltonian Mechanics to permit the random walk to move more quickly through the parameter space. Each parameter θ_j is given a momentum variable ϕ_j , which are updated simultaneously in the modified Metropolis algorithm, and making it so that the proposal distribution for θ_j is heavily dependent on the momentum parameter ϕ .

In a HMC simulation, we alter the posterior density $p(\theta|y)$ using a new momentum distribution, $p(\phi)$, giving us a new joint distribution $p(\theta, \phi|y)$. The simulation uses this joint distribution, but the vector ϕ is of little interest (at least, from the perspective of our final inference results) and is treated as auxiliary, existing only to improve algorithm efficiency. HMC simulations also require the gradient of the log-posterior density, which often needs to be computed analytically, since numerical differentiation of this vector for high-dimension problems is computationally expensive in practice.

The algorithm is as follows:

1. Update the momentum parameter ϕ with a newly sampled value from its posterior distribution, which we specify to be the same as its prior.
2. Simultaneously update θ and ϕ . This process involves L leapfrog steps, scaled by a factor ϵ , during which θ and ϕ are changed depending on one another. These leapfrog steps proceed as follows:

892 (a) Using the gradient of the log-posterior density of θ , make a half-update of ϕ :

$$\phi \leftarrow \phi + \frac{1}{2}\epsilon \frac{d\log p(\theta|y)}{d\theta}.$$

893 (b) Update the θ vector using the vector ϕ :

$$\theta \leftarrow \theta + \epsilon M^{-1}\phi,$$

894 where M is the covariance matrix of the momentum distribution $p(\phi)$. We ex-
895 plicitly state ϵ here rather than absorbing into M for reasons we outline shortly.

896 (c) Use the gradient of the log-posterior density of θ to perform the remaining half-
897 update of ϕ :

$$\phi \leftarrow \phi + \frac{1}{2}\epsilon \frac{d\log(p(\theta|y))}{d\theta}.$$

898 We note that during this leapfrog step, steps (a) and (c) need only be done in-
899 dividually on the first and last step of the leapfrog step— they can be done si-
900 multaneously on the steps in between. So this leapfrog step should start with a
901 half-update, after which we can have $L - 1$ full updates, and finish on a final half-
902 update. In the limit $\epsilon \rightarrow 0$, this leapfrog step preserves the joint density $p(\theta, \phi|y)$,
903 which we state without proof (Gelman *et al* give an intuitive explanation in (46)).

904 3. Set θ^{t-1}, ϕ^{t-1} as the parameter and momentum values prior to the leapfrog step, and
905 θ^*, ϕ^* as the newly proposed values after L steps. Compute

$$r = \frac{p(\theta^*|y)p(\phi^*)}{p(\theta^{t-1}|y)p(\phi^{t-1})}.$$

906 4. Update θ^t according to

$$\theta^t = \begin{cases} \theta^* & \text{with probability } \min(r, 1) \\ \theta^{t-1} & \text{otherwise.} \end{cases}$$

907 Note that since ϕ^t gets updated immediately at the start of the next iteration, and that
908 we do not really need this value after step 3 anyway, we do not bother with storing this
909 value.

910 We repeat this until we reach approximate convergence, using the same checks outlined
911 in the previous section. This algorithm is designed to work on target densities which are pos-
912 itive, which is useful for applications in cancer biology, since we seldom deal with negative
913 values when dealing with cancer cell populations, drug dosages and the like.

914 The algorithm can be calibrated as needed in three different ways. We can modify the
915 probability distribution for the momentum variables ϕ , change the number of leapfrog steps
916 L (which, for our purposes, is unnecessary, thanks to the no-U-turn sampler in *Stan*, which
917 we briefly discuss later on), or change the scale factor (often called the step size parameter)
918 ϵ . Typically, an initial choice for the values L and ϵ should be made so that $\epsilon L = 1$, which
919 vaguely informs the HMC algorithm that if we take L steps of length $\epsilon\phi$, we should travel
920 from one end of the target distribution to the other. If we change one value, we should change
921 the other accordingly, meaning that we either take a greater number of shorter steps, or fewer
922 longer ones. Choosing a mass matrix for the momentum distribution can be trickier. We
923 know that it should resemble the posterior distribution, and so our prior should resemble this
924 as much as possible. If we have a rough estimate of the variances of each parameter, we can
925 use this as a starting point – for example, if we have a data set in which each observed value
926 has a standard error close to 1, then a mass matrix which is the identity matrix would be a
927 good start: if the standard error was 10 for each observed value, then a diagonal matrix with
928 10s on the diagonal might be a better choice. The choice of mass matrix is part of the design
929 of the model itself, and calibrating our model as accurately as possible will improve both the
930 efficiency and approximated convergence of our inference.

931 1.2.8 No-U-Turn Sampler

932 The No-U-Turn sampler (NUTS) algorithm was originally presented in (55), and *Stan* uses
933 a modified version of this algorithm, which is outlined in (45). The need for the NUTS al-
934 gorithm arises from the fact that HMC algorithms have their use limited by the requirement
935 by the user to specify the number of leapfrog steps, L , and the step size parameter, ϵ . Setting

these parameters requires a specific expertise – one would need to perform several preliminary runs of a model and know how to interpret the results of said runs properly just to make a heuristic estimate of these parameters. The number of steps, L , is particularly difficult to estimate. The NUTS algorithm determines the value of L adaptively at every iteration of the HMC algorithm. The full algorithm is quite lengthy, so we just give an intuitive treatment here and refer readers to (55) for the full outline. Essentially, the NUTS algorithm forces the sampler to follow one trajectory for as long as it can. When the dot product of the momentum variable in the HMC algorithm, ϕ , and the distance between the current position of the sampler and its initial value θ is negative, the sampler stops following that trajectory (preventing U-turns in the algorithm trajectory, hence the name) and a new iteration is started.

Since the number of leapfrog steps is determined for us, it is easier to set the step size ourselves on the fly. *Stan* has a built-in feature that tells users at the end of the simulation how many divergent transitions occurred during the runs. A divergent transition occurs when the approximated Hamiltonian trajectory deviates too much from the true trajectory – this is measured by the distance between the Hamiltonian value and the initial value (56). These are undesirable and should be eliminated entirely where possible – one way to mitigate these divergent transitions is to alter the step size parameter ϵ . Other methods to mitigate these divergences are more involved, or rely on better preliminary data analysis: one can either reparameterise their model, or use more informative priors. The latter is best practice, since priors should be as informative as possible; the former must be done on a case-by-case basis, and will depend both on the data being modelled and the model itself.

Betancourt outlines in (57) how we might generalise this sampler to Riemannian manifolds, which relates to an alternative method of HMC optimisation called Riemannian adaptation, in which the mass matrix M is adapted at every step so that it conforms to the local curvature of the log posterior. This is beyond the scope of this review since it would require a crash course in differential geometry, but it is worth noting that the NUTS sampler is not the only way to optimise HMC algorithms.

1.3 Spatial Agent-Based Modelling of Tumour Evolution

This section of the literature review is published at (58). Many thanks to Robert Noble for their editorial contributions, Maciej Bak for their contributions to the section on workflow managers, and David Basanta for their contributions on the section on hybrid cellular automata on the bone ecosystem in cancer, as well as other pedagogical contributions.

1.3.1 Introduction

Cancer initiation, progression, and treatment responses are Darwinian evolutionary processes (59; 60) that can be investigated using a wide range of mathematical and computational methods. Examples include evolutionary game theory (61; 62), branching processes (63; 64), and Moran processes (65; 66). Yet while many tools have yielded important insights into cancer evolution, the study of spatial aspects – especially important in carcinomas, constituting the majority of humans cancers – often necessitates a spatially explicit approach, such as a spatial agent-based model.

An agent-based (or individual-based) model is a computational model of a system made up of autonomous, interacting “agents”. Spatial agent-based models (SABMs) have long been used to study the evolution of spatially structured communities because they can reveal how the processes of selection, drift, and gene flow depend on localised interactions among agents (typically individual organisms) or between agents and their spatially varying environment. As new technologies generate better spatial tumour data, SABMs are proving ever more useful in oncology. Typical applications include understanding tumour development, inferring the effects of driver mutations, and predicting treatment outcomes. For example in recent studies, Aif *et al.* (67) used an SABM to investigate the evolutionary rescue of drug-resistant tumour subclones; Saha *et al.* (68) used an SABM to investigate adaptive cancer therapy; and Bull and Byrne (69) used an SABM to simulate interactions between macrophages and tumour cells.

To support this burgeoning research field, here we present a seven-step guide to designing and implementing spatial agent-based models in which the agents are locally-interacting tumour cells or cell subpopulations. Starting from the simplest cellular automata, we discuss options for adding greater complexity and biological realism, such as multi-level spatial

structure and environmental heterogeneity. Based on our extensive experience of developing and using SABMs (70; 71; 72; 37), we cover practical issues such as event scheduling, visualisation, and how to use SABMs to infer parameter values from experimental or clinical data. Each topic is illustrated with examples from our own demon-warlock modelling framework (72; 70), other state of the art modelling platforms, and studies that have used SABMs in cancer research. Whereas our focus is on tumour evolution, much of our advice applies equally to similar modelling methods used to study bacterial colonies, invasive species, and organismal development. The guide is designed to be accessible for biologists and clinicians without specialist mathematical knowledge.

1.3.2 Spatial structure

Spatial structure determines the evolutionary balance between selection and drift, the nature of gene flow between subpopulations, and the strength of ecological interactions. When a model fails to accurately represent the spatial structure of a biological system, the model's predictions and inferences for that system may be highly unreliable (70; 73). It follows that the parameters of spatial structure – such as the size of locally interacting cell communities and the manner of cell dispersal – should be accorded the same importance as evolutionary parameters in model design. Notwithstanding the trade-off between model simplicity and realism, spatial structure parameters should, as far as possible, be derived or inferred from empirical data.

Many of the simplest spatial agent-based models are cellular automata. A cellular automaton is a model that plays out on a grid of sites in one or more dimensions. Each site is associated with one of a set of at least two possible states. Each site also belongs to a subset of sites called a neighbourhood, of which some examples are shown in Figure 1.1. For example, the von Neumann neighbourhood in two dimensions contains the nearest sites in the cardinal directions (up, down, left and right). A cellular automaton sequentially updates itself according to a set of rules. The update rules for a given site depend on its own current state and the states of the sites in its neighbourhood.

Whereas the update rules of many cellular automata are deterministic (74), probabilistic rules are more appropriate for modelling stochastic processes such as biological evolution.




FIGURE 1.1: Some common neighbourhoods that govern the update rules for cellular automata and other agent-based models in two dimensions (a) and three dimensions (b). A focal agent (cell) is shown in blue and its neighbourhood sites in grey.

A stochastic cellular automaton is equivalent to a collection of locally interacting Markov chains, which means that each event is chosen according to probabilities that depend only on the current model state, not any of its previous states.

In biological terms, each state corresponds to a type of cancer cell or some other entity (such an immune cell or part of the extracellular matrix). Generally we will assume that the focal agents in our models are cancer cells and we will use the terms “agent” and “cell” interchangeably where appropriate. A cellular automaton permits a cell’s event probabilities (for example, its division, death, and dispersal rates) to depend on the number of neighbouring cells. This allows us to account for crowding or Allee effects, such that birth, death or dispersal rates depend on the local or global population size. Event rates can also vary according to the types of the neighbouring cells, for example to simulate cell competition or immune predation.

Models of asynchronous processes, such as cell division in a tumour, typically use asynchronous updating, meaning that only one or a small number of sites are modified per update (75). In addition to being more realistic, asynchronous updating is often necessary to prevent conflicts. For instance, if two cells are attempting to divide but only one space is available for the two potential daughter cells then one must take priority.

Among the simplest stochastic cellular automata is the Eden growth model. This model is typically implemented on a two- or three-dimensional regular square grid with only two

possible states: unoccupied (S_0) and occupied (S_1). With each iteration, the update rule causes a site in the neighbourhood of an S_1 site to switch from S_0 to S_1 . In this way new S_1 sites (cells) are added to the surface of a cluster. The Eden growth model on an n -dimensional grid self-organises to resemble an n -dimensional ball with a non-trivial surface. The growth curve of the S_1 population approaches a polynomial of degree n (76).

The three most popular options for the Eden growth model update rule can be labelled alphabetically:

- **A**vailable site-focussed: Choose at random an S_0 site in the neighbourhood of an S_1 site, and switch it from S_0 to S_1 .
- **B**ond-focussed: Choose at random an S_1 site with a probability proportional to the number of S_0 sites in its neighbourhood, and then randomly choose an S_0 neighbour and switch it to S_1 .
- **C**ell-focussed: Choose at random an S_1 site with at least one S_0 site in its neighbourhood, and then randomly choose an S_0 neighbour and switch it to S_1 .

Although these update rules result in similar large-scale patterns, they generate cluster surfaces with different microscopic properties. Indentations in the model surface are more likely to be filled, and spikes are less likely to form, under option C than under option B, and under option B than under option A. Hence option C generates the smoothest surface and option A the roughest (77).

Variants of the Eden growth model have been used to investigate the evolution of paediatric glioma (78), colon cancer (79), hepatocellular carcinoma (80), and non-small cell lung cancer (81). Many studies use a variant that includes stochastic cell death. By opening up spaces for cell division, cell deaths increase clonal mixing within the tumour and facilitate selection (80).

Other stochastic cellular automata can be more appropriate than the Eden growth model for modelling systems in which state changes are not confined to the surface. Spatial branching processes are similar to Eden growth models except that if a dividing cell has no space to divide then it can create space by budging other cells. An intermediate model can be




FIGURE 1.2: **(a)** An illustration of the CPM. Here, two cells, shaded in green and yellow, each occupy multiple sites on a grid. In the leftmost panel, we see the model's initial configuration; in the middle panel, the state of site P is copied onto site Q , and the green cell grows; in the rightmost panel, the state of site P is copied onto site R , which is initially occupied by the yellow cell, thus deforming the yellow cell and budging it into site S . **(b)** An illustration of the LGCA. Grid sites shaded in green represent those which may influence the focal cell node, shaded in red. The leftmost panel represents the initial configuration; the middle panel shows a cell dividing into free space on its grid site, with the new cell shaded in orange; the rightmost panel demonstrates how cells might move within a grid site or between grid sites: the red cell has changed direction, and the orange cell has moved from one site to another.

created by stipulating that only nearby cells can be budged, so as to simulate physical constraints on cell division. Chkhaidze *et al.* (82) recently used such a model to investigate how spatially constrained tumour growth alters signatures of clonal selection and genetic drift in cancer genomic data. Good practice is to implement budging along an approximately straight line between the dividing cell and the nearest empty site. If budging is instead restricted to the cardinal directions or the cardinal and intercardinal directions then the simulated tumour will self-organise into an approximate square or octahedron, rather than a more biologically plausible disc or ball.

Another option is to allow dividing cells to replace, rather than displace, their neighbours. In the voter model, the update rule is such that, with a certain probability, a randomly selected site copies the state of a neighbouring site. Biasses can be introduced by setting unequal copying probabilities, corresponding to differences in cell fitness. Simple (linear) voter models satisfy a convenient property called coalescing duality, which means that their typical behaviour can be explained through mathematical analysis (83). In a pioneering 1972 study, Williams and Bjerknes (84) used a biassed voter model to simulate the spread of skin cancer through the basal epithelial layer.

The cellular Potts model (CPM), also known as the Glazier-Graner-Hogeweg model (85; 86), more explicitly simulates physical interactions among cells and between cells and their microenvironment. The model takes place on a lattice and each cell is represented by multiple lattice sites (as opposed to only one lattice site, as in previously discussed models), corresponding to the cell's volume. Cells are deformable and can adhere to one another or to surrounding empty sites (which might represent extracellular matrix or growth medium). Hamiltonian mechanics describe the overall energy of the system depending on adhesion forces and resistance to changes in cell volume. A random lattice site is chosen at each time step and its state is copied to a random neighbouring site. If the new configuration has lower energy than the previous configuration then the change is always accepted; otherwise, the probability of accepting the change depends on the Boltzmann temperature. The CPM has been used in numerous cancer studies, such as for simulating tumour growth, invasion and evolution (87), or for investigating how cell compressibility, motility and contact inhibition shape tumour cell clusters (88). The CompuCell3D modelling environment compucell

provides an efficient, flexible CPM implementation.

The biological lattice gas cellular automaton (89) excels instead at modelling cellular movement, and especially collective migration, in a simple, computationally efficient, and physically correct fashion. The model must play out on a square or hexagonal lattice in 2 dimensions, or a cubic, dodecahedral or icosahedral lattice in 3 dimensions. States incorporate cell velocities. For instance, consider a 2-dimensional square lattice in which each site contains 5 nodes: one for each directional velocity and a resting node at the centre. A cell occupying any one of these nodes can divide into other nodes on the same site. A cell can also reorient itself by moving between nodes on the same site, and can move between sites according to its velocity, provided there is space to do so. This model has been used, for example, to give insights into breast cancer invasion plasticity (90). Illustrations of the CPM and LGCA are shown in Figure 1.2.

An important limitation of all the aforementioned cellular automata is that their uniform spatial structures are inconsistent with the biology of many tumour types. Various common cancers have glandular structures and grow via individual cells or small cell clusters invading neighbouring tissue (91; 92). Colorectal adenomas are also glandular but grow through gland fission (93).

Inspired by classical population genetics models (94), a simple, conventional way to account for multi-level spatial structure in tumours is to assign cells to local subpopulations, called demes, located on a regular grid. Thus each grid site is allowed to contain not only one but dozens, hundreds, or thousands of cells. The subpopulation size per deme is prevented from exceeding a certain threshold – known as the deme’s carrying capacity – by decreasing cell division rates or increasing death rates as the subpopulation size grows.

Deme-based models allow for more complicated modes of cell dispersal. As in the voter model, cells can be assigned some probability of invading neighbouring demes, either individually or in clusters. The dispersal probability can also be made to depend on the population of the deme being invaded, so that cells disperse more easily in less densely populated regions near the tumour periphery. Alternatively, each occupied deme can be assigned a probability of undergoing fission, resulting in some of its cells being moved to an unoccupied neighbouring deme. Depending on the degree of budding allowed, the deme-level

dynamics of the fission model can resemble an Eden growth model (no budding of demes) or a spatial branching process (unlimited budding). Deme-based models additionally allow for the explicit simulation of tissue invasion, such that a tumour can grow only via its cells invading demes that are initially filled with normal cells (70).

If the within-deme subpopulations can be assumed to be well-mixed then cells that belong to the same deme and have the same phenotype and genotype can be modelled collectively, rather than as individual agents. This model design not only improves computational efficiency but can also facilitate mathematical analysis. For example, when cells disperse by invading neighbouring demes, the model can be designed so that the dynamics are approximately equivalent to the well understood spatial Moran process (70). Cells can be randomly selected within a deme by sampling from a hypergeometric distribution.

Even greater efficiency can be realised by not modelling inter-deme dynamics at all, and simply making the demes themselves the model agents (95; 96). Although such coarse-graining enables the simulation of much larger tumours, it comes at the cost of reduced precision. Care should be taken in translating between mutation rates per cell and effective mutation rates per deme.

Instead of confining agents to a regular grid, we might instead locate them in continuous space. This structure is potentially more realistic but also entails more parameters, more decisions to be made, and typically higher computational costs (97). To prevent multiple cells occupying the same space and to maintain tumour integrity, we now must model the movement of cells in response to physical forces such as cellular adhesion and repulsion (98). We may also choose to model directed movement under the influence of diffusible factors (haptaxis).

There are several practical ways to prevent cells overlapping in an off-lattice model, depending on how the agents are implemented. Suppose we have spherical cells, each with fixed radius r . We can then specify that when, as a result of cell division or movement, the distance between two cells' centres is less than $2r$, both cells will simply be pushed in opposite directions. Alternatively, to account for cell deformation, we might implement repulsion only when the distance between cell centres falls below some threshold value smaller than $2r$ (99). Some modelling platforms achieve greater realism and tractability by implementing

adhesion and repulsion forces using functions rooted in physics, which are beyond the scope of this guide.

1.3.3 Mutation

Having chosen an appropriate spatial structure, we next will decide which cell phenotypes and genotypes to include in our state space, and how to model mutations between these states. As ever, the goal is to balance model simplicity, realism, and computational demands.

A good part of the difficulty in designing a useful model stems from the fact that much of the experimental data gathered by cancer biologists focusses on genetic mutations while the rules that govern the behaviour of the agents in an SABM assume an understanding of the key cancer phenotypes. The most basic actions a tumour cell might perform at any given time step are apoptosis/death, proliferation, and motility. These are often considered as simple probabilistic events and often modelled in an exclusionary manner, so that if a cell is moving then it is neither proliferating nor dying. The required probabilities can either be taken directly from experimental data (which is often hard to measure *in vivo* and unrealistic *in vitro*) or calibrated with *in vivo* pre-clinical models. In hybrid SABMs, cells can also respond and secrete diffusible factors such as cytokines. Many of these have been experimentally characterized but many assumptions about key parameters such as the rate of diffusion of the cytokine in the microenvironment or the rate of secretion by the cell either lack experimental support or, if available, has been measured in a different context.

Using hard-coded rules to model the phenotype of a tumor cell, while relatively simple, does not capture the flexibility shown by biological cells in the mapping between genotype and phenotype. Gerlee and colleagues have instead proposed capturing some of the complexity of this mapping by embedding neural networks inside each agent, so that the phenotype emerges in a non-linear way as a result of the agent's state and the different microenvironmental inputs to which the agent is receptive (100).

Once phenotypes have been defined, the next step is to determine how these phenotypes will change as a result of mutations. One option is to model mutations as phenotypic switches. Many studies consider models with only two possible tumour cell states – mutated and unmutated – which differ in fitness (95), degree of drug resistance (38), or some




FIGURE 1.3: The result of running an Eden growth model with nearly neutral mutations (top) and beneficial mutations (bottom). Model produced in HAL using some in-built examples as a skeleton for the code. (105).

other trait. Grow-or-go models assume that cells can reversibly switch between predominantly migratory and predominantly proliferative phenotypes (101). Other models examine the evolution of continuous traits, such as levels of glycolysis and acid production (102).

If we are more interested in clonal dynamics then we can explicitly track changes to the (epi)genome. These mutations are conventionally assigned to three groups according to how they affect cell fitness: driver mutations (which increase cell fitness), passenger mutations (no effect), and deleterious mutations (negative effect). For simplicity, most studies assume an infinite sites model (103), such that no two mutations can occur at the same site. Finite sites models must be parameterised based on observed mutation frequencies (104).

We can, for example, convert an Eden model into an evolutionary model by implementing mutation. The grid and neighbourhood are defined as before but now we have multiple cell states $S_1, S_2, S_3, \dots$ and mutation rates between each pair of distinct cell states. A simple option, assuming infinite sites, is to set all mutation rates to be zero except in the case of S_i to $S_i + 1$ for all $i \geq 0$, so that every S_i cell has exactly i mutations. Let us assume that all these mutations are drivers and their effects combine multiplicatively, such that each mutation increases the division rate by a factor of $1 + s$, with $s \geq 0$. Assume also that mutations occur only at the time of cell division, and the number of new mutations per daughter cell is Poisson distributed. We then arrive at a reasonable toy model of spatial tumour evolution that can be implemented in not much more than 100 lines of code, as we illustrate with an R script (106). Figure 1.3 shows results of implementing a similar model in the HAL platform (105).

Modelling the evolution of a quantitative trait, such as cell division or death rate, leads

to further design decisions. As in our toy model, it can be wiser to draw mutation fitness effects from a probability distribution instead of setting them all equal. To see why, consider a model of an expanding tumour that, in the absence of mutation, has radius growth rate c_0 , and in which the spread of mutants is not confined to the periphery (for example, a biased voter model). When a new mutant arises within the wildtype population, its long-term fate, in the absence of further mutation, will be sensitive to its radius growth rate, c_1 . If $c_1 < c_0$ then the mutant will remain forever rare; if $c_1 > c_0$ then the mutant is bound to take over the entire tumour; if $c_1 = c_0$ then the mutant will become relatively more abundant over time without ever fully replacing the wildtype. Randomising the fitness effect randomises c_1 and so randomises mutant fates. Our demon-warlock framework draws each selection coefficient (relative increase in cell division rate) from an exponential distribution.

Strictly multiplicative fitness is best avoided in all but the smallest-scale models as it can lead to unrealistically high fitness values. This is especially problematic if mutation is implemented at the point of cell division, which creates a feedback loop in which lineage fitness grows at an ever increasing rate. A simple solution implemented in our demon-warlock framework is diminishing returns epistasis. When the selection coefficient of a driver mutation is s , instead of multiplying the division rate by $1 + s$, we instead multiply by $1 + s(1 - b/b_{max})$, where b is the previous division rate and b_{max} is an upper bound.

1.3.4 Event scheduling

The next step is to consider how to implement cell events algorithmically. Event scheduling can be the most important factor in determining computational efficiency, especially in simpler grid-based models. The optimal choice strikes a balance between efficiency, simplicity, and biological realism.

The Gillespie stochastic simulation algorithm (107) is an especially simple and popular solution to event scheduling. Event rates are assumed to depend only on the current state of the model and the time between events is exponentially distributed (as in a Poisson process), such that two events cannot occur simultaneously. The steps of the algorithm are as follows:

1. Initialise the system.
2. Set event rates (birth rates, death rates, dispersal rates, etc.).

- 1237 3. Randomly determine the next event such that $\mathbb{P}(\text{event} = E) = \text{rate}(E) / \Sigma(\text{rates})$
- 1238 4. Implement the chosen event.
- 1239 5. Advance the timer by $\delta t \sim \text{Exp}(1 / \Sigma(\text{rates}))$
- 1240 6. Repeat from step 2 until a stop condition is reached.

1241 This algorithm is more efficient than the event timer approach (see below) and is very
 1242 easy to implement. In statistical terms, the simulated sequence of events corresponds to
 1243 a trajectory of a set of stochastic differential equations, called the master equations. This
 1244 means we have a good mathematical understanding of how the algorithm behaves.

1245 Our toy Eden growth model (106) provides an example implementation of Gillespie's
 1246 algorithm. This model further improves computational efficiency by keeping track of the
 1247 cells that have space to divide, so that the next dividing cell can be chosen from among this
 1248 subset (which in n dimensions scales with the radius to the power of $n - 1$) rather than from
 1249 the entire cell population (which scales with the radius to the power of n). The drawback
 1250 is that cells without space to divide never undergo mutation, which may be an unjustifiable
 1251 assumption in a serious research model.

1252 Modifications of Gillespie's algorithm, such as tau leaping (108), are even faster but less
 1253 accurate. Tau leaping allows multiple events to occur simultaneously, which may be prob-
 1254 lematic in a spatial model if the events affect multiple sites in close proximity (for example,
 1255 if two cells are chosen to divide into the same empty site). Moreover, tau leaping improves
 1256 performance only when the system is dominated by a small number of large, homogeneous
 1257 subpopulations, which is typically not the case in SABMs.

1258 A shortcoming of the Gillespie algorithm is that some events, such as cell division, are not
 1259 true Poisson processes with exponentially distributed waiting times. In effect, the Gillespie
 1260 algorithm permits arbitrarily short cell cycles. Some cells may divide several times while, in
 1261 the same period, others with identical division rates fail to divide at all.

1262 One way to achieve more realistic cell cycle periods without sacrificing very much com-
 1263 putational efficiency is to use a phase-type probability distribution. Whereas an exponential




FIGURE 1.4: An example of using a binary tree to select an event (Event 3) from four options. Selected nodes are shown in blue.

1264 distribution models the time until the next event in a Poisson process, a phase-type distri-
 1265 bution models the time taken for an entire sequence of events, which may occur at different
 1266 rates.

1267 In practical terms, this entails executing the Gillespie algorithm as above, except that
 1268 when a cell is selected for division, it doesn't necessarily divide immediately, but instead
 1269 changes its position in the cell cycle. Given a target probability distribution for cell cycle
 1270 periods, we can use an algorithm to choose transition rates such that the resulting phase-type
 1271 distribution has the same mean, variance, and skew as the target (109). For example, suppose
 1272 that all cells begin in division state 0. When a cell is selected (according to a state-dependent
 1273 probability), its state is updated from 0 to 1, 1 to 2, or 2 to 3. When a state 3 cell is selected
 1274 it divides and both progeny are reset to state 0 (110). The method's greater realism comes
 1275 at the cost of additional memory demands and longer execution time, compared to the basic
 1276 Gillespie algorithm.

1277 When we have more than a handful of events to choose from it will be much more efficient
 1278 to implement event selection using a binary tree. Suppose, for example, that we have four
 1279 possible events with rates p_1, p_2, p_3 and p_4 , where $p_1 \leq p_2 \leq p_3 \leq p_4$. If we store the rate
 1280 sums $p_1 + p_2$, $p_3 + p_4$, and $p_1 + p_2 + p_3 + p_4$ then we can choose an event as follows. First we
 1281 generate a random number r from a uniform distribution between 0 and $p_1 + p_2 + p_3 + p_4$,
 1282 and we examine whether $r < p_1 + p_2$. Supposing r is greater than $p_1 + p_2$, we then test
 1283 whether it is less than p_3 . If so then we choose event 3; otherwise event 4. Effectively, we
 1284 have traversed a binary tree, beginning at the root node associated with the sum of all event
 1285 rates, and ending at a terminal node associated with a single event (Figure 1.4).

The binary tree method is efficient because both the number of steps needed to choose an event, and the number of nodes that need updating following a change in an event rate, grow only with the logarithm of the number of possible events. For example, we need only twenty steps to choose between one million possible events. As long as the cell population keeps growing, there is little benefit to pruning nodes and it is easy to ensure that the tree remains balanced. The rate sums together take up only as much computer memory as the individual rates. The main costs are in terms of code development time and code complexity. Binary trees require careful implementation and error checking to ensure that existing nodes are updated and, when required, new nodes are added after each model event. Our demon model implements binary trees and periodically recalculates event rate sums to prevent excessive accumulation of rounding errors.

A less efficient alternative to using phase-type distributions is to draw inter-division times directly from a chosen probability distribution. This approach enables more precise tracking and adjustment of individual cell cycles. An algorithm used in recent studies (102; 38) is as follows:

1. Initially assign every cell i a countdown timer set to time t_i drawn from some probability distribution (dependent on the cell's phenotype).
2. Subtract δt from every countdown timer, where $\delta t \ll t_i$ for all i .
3. For all cells i , in random order:
 - A. Implement cell death and dispersal events for i ;
 - B. If i is alive, has space to divide, and $t_i \leq 0$, then i divides;
 - C. Assign each new cell a countdown timer, set to some random time dependent on the new cell's phenotype.
4. Repeat from step two until a stop condition is reached.

How much this approach reduces computational efficiency will depend on other aspects of the model. It is likely to be much slower than a well implemented Gillespie algorithm when applied to a simple grid-based model, due to the additional burdens of updating every cell (Step 2) and shuffling all the cells (Step 3) at each small time step. In an off-lattice model,

where cells move much more frequently than they divide, and where a shuffling algorithm may already be required to randomise the order in which cell positions are updated, the cost of updating cell division state at the same time as position may be negligible.

There is not a correct answer when choosing how to schedule cell events – we must choose the most appropriate one for our model based on the system we wish to model and how we wish to model it.

1.3.5 Microenvironment

Whereas many SABM studies focus on the effects of spatial structure and cell-cell interactions, real tumours evolve in a complex microenvironment that varies over space and time. This tumour microenvironment, comprising both molecular elements, such as cytokines, and other (non-cancer) cells, constitutes the cancer ecosystem (111) – a key element of the selection process driving somatic evolution. Given a good rationale and sufficient parameterisation data, we may choose to extend our model by explicitly simulating microenvironmental factors in the form of agents (in the case of immune cells or stromal cells) or diffusible factors (such as oxygen and drugs). Permitting cancer cells to modify their selective environment creates potential for emergent complexity and niche construction (112; 113).

Hybrid cellular automata (or HCA) have been used to model interactions between tumour cells and diffusible factors for more than twenty years. As described in a pioneering 2001 paper by Patel and colleagues (114), these models consist of two interdependent components: stochastic cell events, and deterministic reaction-diffusion partial differential equations. The latter component dictates how chemicals or other factors work their way through the system as they are consumed and processed by cells. Local concentrations of diffusible factors contribute to the cell update rules.

Typically we assume that diffusible factor concentrations rapidly re-equilibrate following changes in the configuration of cells. We can then numerically solve the equations to find the equilibrium concentrations either after every cell event or, trading some accuracy for greater efficiency, after a relatively small number of cell events have occurred. Suitable procedures for solving partial differential equations as initial value problems can readily be found in textbooks and software libraries. These range from simple but inefficient algorithms

based on the classical Gauss-Seidel method, which require only a few dozen lines of code (115; 114; 37), to the highly sophisticated BioFVM solver (116), which is specifically optimised for hybrid SABMs. Several SABM platforms include their own methods for solving reaction-diffusion equations in two or three dimensions, which we review briefly later.

To add biological realism, we might make cell division and death rates in our model depend on the local oxygen and glucose concentrations as these factors diffuse through the tumour from the surrounding medium (in very small tumours and tumour spheroids) or from point sources representing blood vessels (in larger, vascularised tumours). We might also modify dispersal rates so that cells follow oxygen or glucose gradients. Potential adverse factors include acid produced through tumour cell metabolism, and drugs that diffuse from blood vessels. Hybrid cellular automata are especially suitable when the supply of an influential factor is highly variable over space or time, such as in the case of intermittent drug treatment (117).

1.3.6 Parameterisation and inference

Although theoretical models can be valuable for generating hypotheses and providing proof of concept, if we want to apply an SABM to studying a particular biological system then we must ensure that its influential parameter values are set appropriately. Parameterisation should ideally be based on clinical or experimental data specific to the biological system of interest; otherwise values can be estimated from studies of similar systems or theoretical considerations (for instance, diffusion coefficients approximately correlate with molecular weight). Influential parameters might pertain to the effects of mutations, drugs, oxygen and glucose; rates of chemical supply, diffusion, consumption and decay; cell dispersal modes and rates; baseline cell death rates, crowding effects and the size of interacting cell communities. Since calibrating SABMs is often computationally demanding, high-performance computation may be required to generate the necessary resources to calibrate them properly.

Bacevic and Noble *et al* (37) parameterised a HCA to mimic tumour spheroid evolution under drug treatment. In spheroids the limiting factor for cell survival and proliferation is oxygen. Other diffusible factors such as glucose were therefore omitted to simplify the model without compromising its usefulness. The oxygen concentration in the medium and oxygen

diffusion rates were drawn from previous studies (118; 119; 120), as were the mathematical relationships between oxygen consumption rate, cell proliferation rate and local oxygen concentration (121; 122). The different maximum proliferation rates of drug-sensitive and resistant cells, reflecting a fitness cost of resistance, were determined from new monolayer growth assays. Cells with insufficient oxygen supply were assumed to die.

Since oxygen effects alone fail to account for the extent of quiescence observed in tumour spheroids, Bacevic and Noble *et al* implemented crowding effects by permitting cell budding only within a specified radius. New monolayer growth assays revealed that the relationships between cell proliferation rate, death rate and drug dose could be well approximated with piecewise linear functions. The drug's impact on proliferation was further assumed to multiply the oxygen effect, consistent with prior observations (121). Drug consumption was also modelled using Michaelis-Menten kinetics, with a diffusion rate chosen according to the drug's molecular weight and an appropriately low consumption rate. Thus parameterized, the SABM accurately predicted the outcomes of new tumour spheroid experiments (37).

Araujo and colleagues (123) developed a hybrid cellular automaton for which the goal was to capture the ecosystem of the bone. A crude approximation of this ecosystem includes the bone itself, the myeloid-derived cells such as osteoclasts that resorb bone, and the cells derived from mesenchymal stem cells, such as osteoblasts, that deposit new bone. Each of these cell types can be modelled as discrete agents regulated by diffusible factors – such as TGF- β , RANK ligand, and other factors embedded in the bone matrix – described by partial differential equations. Parameterisation of the model is facilitated by the fact that non-cancerous cells have more predictable phenotypes, and the model's overall behaviour can be calibrated to ensure it recapitulates bone homeostasis. Araujo and colleagues thus studied how bone metastatic prostate cancer cells could infiltrate the bone ecosystem, take advantage of it, and grow (124). They also investigated what prostate cancer cells in the primary tumour should be of concern to physicians, and why conventional treatments that fail to disrupt tumour-ecosystem interactions also fail to provide long-term cancer cures in bone metastatic prostate cancer (125).

Unknown parameter values can be inferred by combining an SABM with a statistical method. This is, in fact, often the main objective of an SABM study. Approximate Bayesian

computation is a popular approach that, in its simplest form, infers the value of a parameter θ as follows;

1. From our data, calculate some summary statistic μ_{data} ;
2. Set $i = 1$;
3. Run the model using a candidate parameter value θ_i drawn from some prior distribution;
4. Calculate the summary statistic μ_i for the model output;
5. If the difference between μ_i and μ_{data} is less than a predefined tolerance then add θ_i to the posterior distribution;
6. Increment i ;
7. If i is less than some threshold then repeat from step 3.

Although simple in principle, approximate Bayesian computation requires careful implementation. The accuracy and precision of inferences depend on the choices of prior distributions, summary statistics, and tolerances, as well as the number of iterations. Typically multiple parameter values cannot be precisely derived from prior data or models, in which case each should be assigned a vague (high variance) prior distribution. Tolerance values should be tuned such that neither too many nor too few candidate parameter values are accepted to the posterior distribution. Summary statistics should capture features of the system that provide useful information about the parameters of interest. A useful template is a 2010 study (126) in which Sottoriva and Tavaré inferred aspects of stem cell dynamics in the colonic crypt by combining a cellular Potts model with approximate Bayesian computation, using a summary statistic based on methylation patterns.

An alternative to this approach was recently outlined in (127), in which the authors describe a novel method utilising neural networks to reduce both tumour images and SABM simulations to low-dimensional points. The distance between these points acts as a quantitative measure of how the two differ. This enables direct comparison, and by using parameter fitting algorithms to minimise the distance between the two sets of points, parameters can be estimated directly from the images and the simulations.

Whatever the objective, an essential step in any modelling study is to examine, as far as is practical, how the results and conclusions depend on uncertain aspects of the model. A common approach is to run a large number of model variants with different combinations of plausible parameter values. Varying one parameter at a time can provide useful insight into which parameters have the greatest impact on model output, with the shortcoming that non-linear interactions between parameters are often neglected. A more sophisticated approach is to infer a multivariable “metamodel” function – a model of the model – that approximately describes how the model’s parameters determine its outputs.

Since varying many parameters systematically on a continuous scale is infeasible, sampling methods such as Sobol sequencing (128) or Latin hypercube sampling (129) can be used to generate a set of near-randomly sampled combinations of parameter values. Both methods were used in a recent SABM study of cancer cell response to ATR-inhibitors (79). A recent introductory review explains specifically how to apply these methods to cancer ABMs (130). It is important to note that thorough sensitivity analysis involves varying not only parameter values but also mathematical relationships, aspects of spatial structure, and any other influential model components.

1.3.7 Visualisation

Having built and parameterised a model, we next require useful ways to visualise its output. Typical methods represent spatial information, multidimensional phenotypic information, or evolutionary dynamics. Representing all of these aspects in a single image is generally impossible.

A spatial plot represents the state of an SABM at a moment in time. Producing a two-dimensional grid plot of a two-dimensional on-lattice model is straightforward – we simply output the state of each site as a matrix of numbers and input this matrix into a bitmap (or raster) plotting function in R, Python, MATLAB, or similar software, using different colours to represent the different states (Figure 1.3). Our toy Eden growth model (106) provides an example implementation. Diffusible factor concentrations can be shown outside the tumour using a colour gradient and inside the tumour by adjusting brightness (37). We can apply the same method to off-lattice models by specifying a grid and assigning each grid square




FIGURE 1.5: **a)** Plots of a 3D off-lattice ABM, produced in PhysiCell (131), showing a cross-section of model states of a hanging-drop spheroid growth simulation at different time points, using either a deterministic or a stochastic SABM. Cells are coloured according to cell cycle position. Cells in the K_1 cell cycle state are green, post-mitotic K_2 cells are magenta, quiescent cells are pale blue, apoptotic cells are red, and necrotic cells are brown. Cell nuclei are shown in dark blue. **b)** Muller plot showing phylogenies and phenotype frequencies over time. **c)** Fish plot showing phylogenies and phenotype population sizes over time. **d)** 2D grid plot corresponding to same simulation as the Muller and fish plots in previous panels, with the same colour scheme, at the final time point. Plots b and c were produced using the R package *ggmuller* (132). Image a is reproduced from (131) under the terms of a Creative Commons Attribution License and with the approval of Paul Macklin. Plots b-d are reproduced from (71) under the terms of a Creative Commons Attribution License.

a value that summarises the states of all points within the square. Given multi-level spatial structure, we can represent the most abundant state in each deme (70).

Illustrating three-dimensional information is more technically demanding as we need to project the object onto a two-dimensional plane, determine the visible surface, and add shading (as in Figure 1.5a). Suitable computational methods include rasterisation and ray tracing, which can be performed in R and Python or using dedicated software, such as POV-Ray. Further details can be found in the PhysiCell documentation (131). A much simpler solution is to plot only two-dimensional slices.

Muller plots represent subpopulation dynamics and phylogeny, while disregarding spatial information. The horizontal axis represents time and the vertical axis corresponds to subpopulation frequency. Each subpopulation is depicted as a shaded area emerging from its immediate ancestor (Figure 1.5b). Fish plots are similar but show population size rather than frequency (Figure 1.5c). Software packages for producing these plots include ggmuller (132) and EvoFreq (133).

In a phenotype space plot, the axes correspond to continuous traits such as cell fitness, metabolic type, and degree of drug resistance, and each point represents a cell. We can visualize phenotypic evolution by animating phenotype space plots from a series of time points. Robertson-Tessi and colleagues pioneered the use of these plots in cancer research in their 2015 study of the effects of metabolic heterogeneity on tumour growth (102).

1.3.8 Reproducibility

Reproducibility is a cornerstone of the scientific method. A reproducible modelling study not only allows others to easily regenerate its results but also permits further data processing, downstream analysis of generated data, generation of summary statistics, ease of production for visual representations or plots, and even adaptation of the existing model for novel purposes.

Gundersen (134) describes three categories of reproducibility:

- Outcome reproducibility: The reproduction experiment's result matches the original. If the same analysis of the result is performed, the same conclusions can be drawn, and the original hypothesis is supported by both experiments.

- Analysis reproducibility: The reproduction experiment's result differs from the original, but if the same analysis method is used, the interpretation of the results still matches the original.
- Interpretation reproducibility: The reproduction experiment's outcomes and the analysis of said outcomes both differ, but the interpretation matches the original interpretation.

Computational modelling studies should typically aim for the highest standard of outcome reproducibility. If care is taken to construct a well-packaged computational study in a controlled digital environment, then in principle, given a suitable machine, the study should be easily reproduced exactly. This entails not only comprehensively explaining methods, results, analyses, and interpretation, but also sharing the model code and scripts used at every step of pre-processing and analysis, providing a detailed description of how to execute the code, and sharing any associated data and parameterisation and configuration files.

In their outline of best practices to observe throughout a computational research project, Sandve and colleagues (135) advocate tracking how every result is produced and reporting intermediate results as well as final outcomes. To make code easier to reproduce, one should catalogue the versions of software used at every point, record the seeds used in any random number generation, and implement version control (136). Manual data manipulation should be avoided in favour of using automated methods to reformat and process raw data files. The raw data used to produce summary data plots should be easily accessible to facilitate easy plot reproduction and to allow readers to check individual data points. Textual descriptions of methods and results should link to the associated raw data and code so that a reader can easily follow all the steps leading to interpretations. Lastly, modellers are highly encouraged to share each full study, ideally with a dedicated public server. One such research-oriented database is zenodo (137), where scientists may freely upload their research output permanently as a citeable piece of software.

A complex computational model will often require multiple steps to be carried out in sequence. If a high-performance computing (HPC) cluster is required to run the model efficiently – as is typical for complex models – it is essential to utilise a workflow manager to properly orchestrate the steps (138). Open-source workflow managers allow researchers to

package a model into a reproducible, cross-platform workflow. Nextflow (139) and Snake-make (140) are among the most popular workflow managers with several published pipelines (141; 142; 143; 144), strong community support, and extensive documentation, giving users flexibility when designing their own custom pipelines. Snakemake is based on Python, a popular language among computational biologists and bioinformaticians. Nextflow uses the Java-based language Groovy, which has a Python-style structure and is relatively easy to for Python users to learn. Both also enable automatic parallelisation for HPC clusters, which can be essential for complex SABMs or for running multiple instances of smaller models simultaneously.

Another option is to utilise container technologies, considered by many to be the gold standard in computational research. These are less computationally demanding than running an application on a computer directly or using a virtual machine and so permit faster deployment, patching, and scaling. Containers also allow users to deploy the application on multiple operating systems or machines without reformatting, and will run the application the same way no matter where they are deployed (145). Docker (146) is a popular container design platform which permits packaging applications into distribution-independent containers. Another option, Bioconda (147), enables easy dependency management, and can be deployed inside a container.

For many research projects, the easiest option can be to build on an existing open-source agent-based modelling platform. Some of these platforms – such as Chaste (148), CompuCell3D (149), HAL (105) and PhysiCell (131) – excel in simulating off-lattice cell populations in complex microenvironments. Others, such as demon (150) (which has an automated computational workflow, Warlock (72)), J-SPACE (151) and SMITH (152), focus on efficient modelling of evolutionary dynamics. Several are modular platforms, which facilitate reproducibility by making it easy to create and share extensions of the generic software. Nevertheless, even the most flexible platform is necessarily based on certain fundamental assumptions, structures, and algorithms. If we want to create an especially innovative model, requiring several novel components that pre-existing modelling platforms lack, then we might find it best to start from scratch. In principle, specialist rather than generalist models permit greater optimisation in terms of memory demands and execution time.

As the volume of publicly available research data has been growing exponentially in recent decades (153), proper digital data management and annotation is recognized as an essential step in computational research – crucial for research reproducibility. Most notably, the FAIR principles have become a cornerstone in modern data management, particularly in the realms of scientific and research data (154). FAIR is an acronym that encapsulates a set of guiding principles: Findable, Accessible, Interoperable, and Reusable. To be FAIR, data must first be Findable, meaning that it is easy for both humans and machines to discover, thanks to comprehensive metadata and proper indexing. Data should be Accessible, ensuring that access rights and permissions are clear and well-defined, thus minimizing barriers to entry. Interoperable data is structured in a way that allows integration with other datasets by adhering to common standards, formats, and vocabularies. Lastly, data should be Reusable, with thorough documentation, contextual information, and availability in a format that facilitates easy replication and reuse. Altogether, the FAIR principles serve as a framework for enhancing data sharing, management, and collaboration, ultimately driving scientific progress and fostering open science initiatives. Major organisations that have embraced FAIR guidelines include the European Open Science Cloud (155), the European Life-Science Infrastructure for Biological Information (156), the US National Institutes of Health (157), and the Global Alliance for Genomics and Health (158). There are a number of open-source frameworks one might use to develop an agent-based model of cancer evolutions. The cell-based version of Chaste (148) is a highly sophisticated, multiscale computational framework for modelling cell populations. Chaste permits both on- and off-lattice models and has built-in code for simulation of specific biological systems, such as cancer development within colonic crypts. Chaste has its own ODE and PDE solver, called SUNDIALS (159).

CompuCell3D (149) is a general-purpose platform for implementing tissue development models, including the Glazier-Graner-Hogeweg (or cellular Potts) model that its developers pioneered. Its bespoke CC3D-Bionetsolver package solves ODEs and PDEs using a finite element method. CompuCell3D has been used in dozens of studies of cancer and morphogenesis.

Demon (150) specializes in simulating intratumour population genetics. Its multi-scale

spatial structure makes it especially well suited to studying the evolution of glandular tumours. Demon can be configured to implement mathematically tractable models such as the Eden growth model, biased voter model, spatial Moran process, and spatial branching processes. An automated computational workflow called warlock (72) facilitates running demon simulations in parallel on a high-performance computing cluster.

HAL (105) is a generic and highly customisable platform comprised of modular components which allow for multiple grids to operate simultaneously, each performing different tasks. For example, one grid might handle cell-cell interactions while another implements oxygen diffusion. HAL has multiple ODE and PDE solvers to suit different modelling needs. It also includes several pre-built model templates.

J-SPACE (151) is a modelling platform designed specifically for phylogenetic modelling. It simulates cancer evolution on a grid (or some other graph) and generates synthetic reads from next-generation sequencing platforms. A primary goal of J-SPACE is to help researchers assess the impact of incomplete data or experimental error on downstream bioinformatics pipelines.

PhysiCell (131) is a flexible framework that can implement physics-based off-lattice models of large numbers of cells in dynamic tissue microenvironments, with dynamic cell-cycle state tracking. PhysiCell uses a custom-built, open-source package for ODE and PDE solving, called BioFVM (116). Potential functions are used to describe cell-cell interactions including adhesion, repulsion, and cell-matrix interactions (99).

SMITH (152) implements a branching process with quasi-spatial constraints that separate the tumour into a proliferating shell and a static core. By simulating the dynamics of clones rather than individual cells, SMITH is able to simulate the evolution of a tumour up to a billion cells in only a few minutes on a standard desktop PC. This computational speed comes at the cost of the model's strong simplifying assumptions.

Morpheus (160) is a highly accessible open-source platform in which users can develop multi-scale, multicellular systems which couple ODEs, PDEs and cellular Potts models, with automatic scheduling. Rather than coding models manually, users can describe the model in biological and mathematical terms in Morpheus' GUI, and utilise provided tools for visualisation and parameter estimation.

1.4 Bipolar Androgen Therapy

Prostate cancer is among the most prevalent cancers on the planet (161), and is a leading cause of death in males (162). These cancers produce prostate specific androgen (PSA), which circulates in the blood, and although these androgens are present in the absence of prostate cancers, their elevated levels when prostate cancer develops allow easier detection and progression monitoring with basic blood tests. Even though these tests allow many prostate cancers to be detected early before they can progress, inevitably, some patients are diagnosed with advanced or metastatic disease. A study conducted in 1972 by Huggins and Hodges (163) highlighted the essential role that androgens play in the progression of prostate cancers, and so androgen deprivation therapy (ADT) has been and continues to be the standard of care for advanced and metastatic prostate cancers. In ADT, the aim is to reduce the amount of testosterone in the body as much as possible; doing so starves the prostate cancer cells of this androgen and either kills them outright or prevents them from proliferating. Regular PSA tests during treatment are a useful tool for monitoring disease progression (164), and although ADT patients with metastatic disease typically see positive PSA response after initial treatment (165), eventually, cancer cells adapt to their new environment and develop a resistance to ADT (166), which we call castrate resistant prostate cancer (CRPC).

CRPC is a more recent classification given to what was once called hormone-refractory prostate cancer (167). Previously, it was thought that the cancer simply adapted to grow and progress without any androgen activity. Recent studies have shown, however, that androgen receptor activity in these cancers is often still essential to their survival and progression (168; 169). The mechanisms of resistance vary and which mechanisms develop may need to be considered as we look forward to modelling any CRPC treatments.

Castrate levels of serum testosterone can be achieved either via surgical removal of the testes (orchiectomy), which account for a vast majority of the body's testosterone production (the rest being produced mainly by the adrenal glands in the brain, and once a prostate cancer develops, often by the prostate cancer cells themselves (170)), or through application of various drugs (171) that either interrupt signalling events that produce testosterone or halt testosterone production altogether. When these methods fail, adjunctive therapies such

as chemotherapy (172) or anti-androgen agents (173) can be used to improve overall survival. There is no fixed optimal strategy for ADT – treatment is continuous and typically is maintained even once CRPC develops. Even if we see an improvement in progression free survival, ADT brings with it a litany of side effects, including, but not limited to, sexual dysfunction, fatigue, hot flushes, osteoporosis, fluctuating mood, anaemia, cardiovascular morbidity and depression (174; 175).

As such, an alternative to standard continuous ADT (which, from this point on, will be referred to as CADT where relevant, to avoid any ambiguity) was proposed by Klotz *et al* (176), called intermittent androgen deprivation therapy (IADT). In IADT, given acceptably low serum testosterone levels and a similarly low PSA level will be given an intermission from CADT during which testosterone levels (T levels) are allowed to recuperate free of inhibition. There are multiple motivations for doing this. With no inhibition on testosterone production, serum T levels are allowed to creep back up to their regular levels, offsetting a number of the previously mentioned side effects of CADT and giving the patient some semblance of normality (at least from a quality of life perspective) before returning to treatment again.

In IADT, since treatment is halted briefly, testosterone levels are allowed to rise, if only briefly. In a noninferiority randomised trial of IADT versus CADT conducted by Crook *et al* (177), they found that in a trial of 690 IADT patients and 696 CADT patients, IADT was no worse than CADT in terms of overall survival, but the rise in testosterone levels seen in 79% of IADT patients (with 35% seeing a return to pre-treatment levels) meant that several of the adverse side effects of CADT were offset. A similar study by Hussain *et al* (178) found that although the test for noninferiority of IADT in terms of overall survival were statistically inconclusive, similar findings were made with regard to quality of life (QoL) improvements in the IADT cohort. A review of several randomised clinical trials of the efficacy of IADT versus CADT by Magnan *et al* (179) had similar findings to Crook *et al*'s study, finding that although IADT was no better than CADT in terms of overall survival, QoL improvements, and in particular an improvement in sexual and physical function, were seen in a majority of the IADT cohorts in the studies reviewed.

A takeaway from these studies is that IADT is, for the most part, no better or worse

than CADT from a survival standpoint, since it does not delay time until CRPC progression any more than its continuous counterpart. One of the rationales behind IADT is that during periods of treatment, hormone-sensitive cell populations are reduced before resistance can develop (ideally), and periods of intermission where testosterone is present in above-castrate levels repopulate the tumour with hormone sensitive cells (which would not be possible with CADT) (180). The increase in testosterone levels is critical to the success of IADT, both in line with this rationale, and to trigger any QoL improvement. Although the above castrate T-levels failed to delay time to progression in IADT trials, a new therapy currently under investigation, called bipolar androgen therapy (BAT) attempts a similar approach, but with a very different execution.

The idea behind BAT, proposed by Denmeade and Isaacs in 2010 (181), arose from the somewhat surprising (although by no means recent) finding that hormone-sensitive human prostate cancer cells which express the androgen receptor (AR) could grow optimally in castrate levels of testosterone (182) despite their supposed reliance on the presence of androgens. Stranger still was the finding that the growth of the cells in these conditions was actually hampered when additional amounts of androgen were introduced to the growth medium (183; 184). These findings are corroborated by *in vivo* studies (185; 186) which found that athymic mice with AR expressing human prostate cancers saw cancer cell growth inhibition when given supraphysiologic systemic testosterone levels, and conversely, saw an increase in cancer cell growth under antiandrogen treatment.

Utilising these findings, the proposed method for BAT would be to give suitable patients CADT as normal, but periodically give patients an injection of testosterone which would raise the amount of testosterone in their system to be well above its pre-treatment levels. After this injection, CADT is resumed to reduce testosterone levels again as quickly as possible. The key idea here is that this “shock” of high testosterone can have a number of beneficial effects. Besides the potential hampering of cancer growth observed in the clinical trials mentioned above, a key biological rationale is that these high testosterone levels, although only maintained briefly, should fool the cancer into potentially abandoning its adaptive response to ADT, and potentially even restore sensitivity to CADT even after this sensitivity has been lost (187; 188). BAT patients should, in theory, also benefit from the same QoL improvements as

IADT patients.

Although there has been some interest in BAT as a new strategy for maintaining ADT sensitivity in prostate cancers, clinical trials have not progressed beyond small phase II trials, and no randomised phase III trials are currently in progress. BAT has not yet been subject to any proper mathematical modelling, either, and so we shall review some of the key details that we might need with modelling this new therapy as our aim. We will take into account the biological mechanisms at play, look over currently completed clinical trials to see how or why testosterone injections are scheduled as they are, and look at some models used for investigating CADT and IADT to see if they may be of any use when modelling BAT.

1.4.1 Testosterone production and the Androgen Receptor

The androgen receptor (AR) is key to the progression and survival of both castration-sensitive and castration resistant prostate cancers. To understand the role this receptor plays, and how we might intervene in this role, it is important to understand how the body produces testosterone. First, the hypothalamus (a part of the brain) releases luteinising hormone-releasing hormones (LHRH) which signals to the pituitary gland (another part of the brain which controls hormone levels in the body) to release luteinising hormones (LH). These then bind to the Leydig cells in the testes, whose function it is to produce testosterone. Once testosterone is produced, it flows through the bloodstream to perform various functions, but particular to our focus, it binds to the effector tissues in the prostate. In the prostate, testosterone is converted into a more potent ligand, dihydrotestosterone (DHT), which, along with testosterone itself, can bind to the AR once it reaches the prostate cells (189), at which point the AR will bind to the cell nucleus.

The AR itself acts as a transcription factor, which is a protein which activates (or prevents) genetic transcription by binding to DNA. This protein has three domains: the ligand-binding domain (LBD), which allows DHT and testosterone to bind to the AR; the DNA-binding domain, which is responsible for the interaction between DNA and the AR in the cell nucleus; and the N-terminal domain, which is responsible for directing proteins to the correct location in the cell, and whose inhibition has been shown to hamper AR transcription (190).

1.4.2 Androgen Deprivation Therapies and Castration Resistance

Huggins and Hodges (163) were the first to demonstrate the testosterone dependence of prostate cancers. Following this discovery, surgical or medical castration has been a common treatment for advanced and metastatic prostate cancers. Surgical castration prevents production of a majority of the body's testosterone via surgical removal of the testes (called orchiectomy). There are multiple different avenues to pursue for medical castration – which one is chosen will depend on how advanced the disease is and the response to current treatment. LHRH agonists and antagonists are a typical choice for preventing testicular testosterone production (191). LHRH agonists can prevent the pituitary gland from releasing LH by binding to the LHRH receptor in the pituitary gland, and LHRH antagonists prevent LHRH from binding to the receptors in the pituitary gland (192). Part of the reason that these agents are used is that the effects are reversible if side effects resulting from treatment are too severe (172). Adjuvant therapies with anti-androgen drugs like Abiraterone or Enzalutamide are often used alongside LHRH (ant)agonists, and have been shown to confer overall survival benefit, as have chemotherapeutic agents such as docetaxel (173; 172).

Even with these treatments, castration resistance inevitably emerges. Castration resistance is defined by a progression of the disease despite treatment with ADT. It can manifest either as a continuous rise in serum PSA, progression of pre-existing disease, and/or the appearance of new metastases (193). This does not indicate that the tumour has become androgen independent, and the varying mechanisms of resistance reinforce this fact. Moreover, it has been found that intratumoural levels of testosterone in CRPC are higher than that of hormone-naïve prostate cancers despite castrate levels of serum testosterone, indicating that these cancers will adapt to ensure that there is still testosterone available for their development in some cases (194).

1.4.3 Mechanisms of Resistance

As we have mentioned, testosterone is produced primarily by the testes, but a small amount is produced by the adrenal gland in the brain. Since LHRH (ant)agonists act on the hypothalamus-pituitary-Leydig cell axis, the production of testosterone by the adrenal glands is not prevented, meaning that low levels of testosterone still circulate throughout the system. In response to this low testosterone microenvironment, cells adapt to compensate via upregulation of the AR, allowing AR activation despite low levels of testosterone(195; 196; 197; 198). Hypersensitivity of the AR has been found in anywhere between 25% to 80% of CRPCs (199; 200; 198).

Another way CRPCs adapt to the absence of testosterone is to mutate in such a way that the AR can be activated by other ligands. Mutations on the AR lead to increased AR activity in the presence of non-androgenic steroids like cortisone (201) and progesterone (202).

Testosterone can also be supplied to tumour cells via alternative steroidogenesis pathways, either due to the pathway of steroidogenesis via the adrenal gland or testosterone production within the tumour itself. The adrenal gland produces the androgen precursors dehydroepiandrosterone (DHEA) and dehydroepiandrosterone sulfate (DHEA-S), which can be converted to DHT, and are unaffected by ADT. Without any ADT treatment taking place, these precursors typically are converted androstenedione (AD), and then to testosterone, and finally to DHT by the enzyme steroid-5 α -reductase. With ADT treatment removing a majority of available testosterone to facilitate this conversion from AD to testosterone, the enzyme steroid-5 α -reductase has been shown to act directly on AD molecules, producing 5 α -androstenedione, which can be converted directly into DHT using other available enzymes in the system (203). This 5 α -dione pathway has been found to be dominant in some CRPCs (204). Increased levels of a variety of enzymes in CRPCs also indicated a capability for the tumour to produce androgens on their own (194; 205; 206; 207).

Being a transcription factor, the AR binds to DNA sequences to promote or prevent transcription of genes. To do this, the AR binds to co-regulatory proteins which can either activate or suppress gene expression, respectively named co-activators and co-suppressors, of which there are over 170 documented (208). Fluctuations in the levels of these co-regulators

can lead to variations in regular transcriptional activity. Studies have shown that certain co-repressors in CRPC are downregulated in the presence of some mutations (209), leading to increased transcriptional activity, while others have found elevated levels of some co-activators in CRPC (210) in the presence of other mutations, leading to increased AR activity.

Aberrant activation of the AR has also been seen in CRPC. Outlaw pathways involving growth factors, cytokines and kinase pathways have been shown to increase AR signalling in CRPC, allowing the disease to progress independent of ligand AR activation by promoting expression of mitotic cell phase genes (211).

Although these highlight the essential role the AR plays in the progression of CRPC, it is possible for the disease to progress independent of AR activation. Signalling cascades initiated by the activation of other ligand receptors activate the expression of genes which are capable of promoting cell proliferation or inhibiting apoptosis (212). Some of these bypass pathways can act both with and without the AR.

1.4.4 Bipolar Androgen Therapy

Denmeade and Isaacs (181) proposed bipolar androgen therapy in response to the findings that hormone-sensitive prostate cancer cells could grow optimally in an androgen-free environment (182) and that these same cells could see growth inhibition after the introduction of large quantities of androgens (183; 184). They followed these observations up with an *in vitro* study which confirmed that all but one of the CRPC cell lines which demonstrated AR activity demonstrated an inhibition in growth when exposed to large amounts of exogenous testosterone.

Although no standardised treatment regimen exists for BAT, the core idea is that a patient presenting with CRPC, on whom ADT therapies alone no longer yield positive response, be given an injection of a large amount of testosterone periodically (typically every month) to yield a blood testosterone level which is higher than the baseline testosterone levels one would typically see in the absence of any androgen ablation therapy. A typical blood testosterone level is somewhere between 300-600 ng/dL (213), but in BAT, one might typically aim for a level of 1000-3000 ng/dL. Once testosterone has been injected, ADT is immediately resumed until the next testosterone injection, meaning that patients will fluctuate between

near castrate levels of testosterone (100-200ng/dL, where castrate levels are considered to be around 50ng/dL) and supraphysiologic levels after testosterone injection. Testosterone injections, as opposed to transdermal or transbuccal gels or intranasal sprays, are chosen not only because they are inexpensive, but because these alternative testosterone delivery methods cannot achieve the rapid supraphysiological testosterone levels needed to see any benefit from BAT.

The rationale behind BAT is straightforward – by abruptly introducing large doses of testosterone, we can potentially interrupt the adaptive response of the prostate cancer cells to castration, such as upregulation of the AR, and maintain for as long as possible a hormone-sensitive tumour on which androgen ablation is effective (214) by gearing for selection of hormone-sensitive clones, with the added benefit of induction of CRPC cell senescence or death seen in previous studies. Clinical trials mentioned above (183; 184) demonstrated that high testosterone levels can induce growth inhibition in CRPC where AR activity was observed – as such, we would expect BAT to be an effective treatment whenever AR activity is found in CRPCs. It may still be effective when CRPC progression is independent of AR activation by creating an environment which reverses the cells' adaptive response to a low-testosterone environment, but we cannot say with any certainty without data to support these claims.

PSA is often used as a biomarker for disease progression in prostate cancer. Prostate cancer cells produce PSA when exposed to testosterone, regardless of whether or not these cells' growth is inhibited by that testosterone, meaning that unlike treatments for prostate cancer like CADT and IADT, PSA level is likely an unreliable measure for disease progression during BAT – something corroborated by a number of clinical trials.

An added benefit of BAT is that patients experiencing the adverse effects of ADT will likely see marked improvement in QoL, since testosterone levels during BAT should always remain above castrate levels, especially after testosterone injection (215).

1.4.5 Clinical trials

Since BAT is a relatively new treatment, there are only a handful of clinical trials to date which have been completed, and with small patient cohorts. Schweizer *et al* (187) enrolled

a cohort of 16 patients with asymptomatic CRPC with a low to moderate metastatic burden, who had been undergoing ADT for at least one year, into a pilot clinical trial on a 28 day treatment cycle. They were given a 400mg injection of testosterone cypionate on day 1 alongside a daily oral administration of 100mg of etoposide (a chemotherapeutic agent) from days 1 to 14. After three cycles, if a PSA decline was seen, the treatment was continued without etoposide, and throughout the trial, ADT was continued to prevent further testosterone production as much as possible. Two patients were taken off the study due to toxicity, and 7 of the remaining 14 patients showed PSA response after three cycles and were allowed to continue into the next phase of the study. 8 out of 14 patients demonstrated a radiographic response after BAT, but only 6 showed a positive PSA response. 10 patients were given hormonal therapy after the study concluded, and all of them showed a PSA decline after initiation of these treatments. A noteworthy point is that 3 of these patients were re-challenged with ADT alongside antiandrogens they had previously stopped responding to, and saw a PSA response, indicating the potential of BAT to restore sensitivity to ADT therapies which had previously failed.

Schweizer *et al* (187) conducted another study investigating the efficacy of BAT on hormone-sensitive, asymptomatic, metastatic prostate cancer. A cohort of 33 patients were enrolled following 6 months of ADT treatment. Those with PSA < 4 ng/dL after the initial ADT were allowed to continue the study, receiving 3 month cycles of BAT and ADT followed by ADT alone. 400mg of testosterone cypionate or enanthate were given on days 1, 29 and 57 during BAT cycles. 29 patients were allowed to start the study following the initial 6-month ADT phase, 17 of whom showed a PSA < 4 ng/dL after 18 months. 10 of 29 patients had a radiographically measurable disease, of whom 8 showed objective response – 4 partial, 4 complete. QoL improvement was observed in all patients following the first BAT cycle after ADT.

Teply *et al* (188) carried out a phase 2 trial on 33 patients with metastatic CRPC and castrate testosterone levels who had received at most two second line hormonal therapies – in particular, patients had to have shown disease progression after treatment with enzalutamide and continued PSA increase after treatment with enzalutamide had ceased. Patients were given 400mg of testosterone cypionate every 28 days until disease progression was seen,

alongside an LHRH agonist. Once disease progression was seen, patients were rechallenged with enzalutamide to determine whether the drug had any effect following BAT. 30 patients were eligible for the trial: 29 completed the BAT trial and a further 21 went on to the enzalutamide re-challenge phase. 9 of 30 patients who underwent BAT showed a PSA reduction of 50%. 15 of the 21 patients re-challenged with enzalutamide after cessation of BAT showed a PSA reduction of 50%, reinforcing the findings of Schweizer *et al*'s 2015 study (187) and demonstrating the potential of BAT for restoring sensitivity to anti-androgen agents.

In one of the biggest trials to date, Denmeade *et al* (216) set out to compare progression-free survival (PFS) rates when comparing BAT to treatment with enzalutamide after abiraterone acetate treatment failed. Patients were allowed to switch treatments once progression was observed. Patients in the BAT cohort were given 400mg intramuscular testosterone cypionate injections every 28 days, and enzalutamide cohort patients given daily oral doses of 160mg, both until clinical or radiographic progression was observed. Both cohorts were maintained on androgen ablative therapies, either via orchiectomy or LHRH (ant)agonists. Overall survival, PSA and objective response rates, PFS after switching treatment (PFS2), treatment safety and QoL were also considered. 222 patients with asymptomatic metastatic CRPC were selected for the study, of whom 195 underwent the trial, with 94 initially receiving BAT and 101 receiving enzalutamide treatment. PFS time for both cohorts was 5.7 months, with a 50% decline in PSA seen in 28.2% of BAT patients versus 25.3% for enzalutamide patients. 77.8% of patients switching from BAT to enzalutamide saw a 50% PSA reduction, who also saw a PFS2 of 28.2 months and OS benefit of 37.1 months. On the other hand, only 23.4% of patients switching from enzalutamide to BAT saw a 50% PSA reduction, with a PFS2 of 19.6 months and OS benefit of 30.2 months. PSA-PFS for enzalutamide was shown to increase from 3.8 months after abiraterone to 10.9 months after BAT. Patient-reported QoL surveys also consistently favoured BAT.

Markowski *et al* (217) carried out a trial in which 45 metastatic CRPC patients were given 400mg testosterone cypionate injections every 28 days alongside an IV of 480mg nivolumab (an immunotherapeutic agent) on the same dosing schedule. Patients were given a 12-week lead-in of BAT alone. Patients were also allowed one cycle of taxane-based chemotherapy to treat the mCRPC prior to the start of the trial. 18 of the 45 patients saw a PSA decrease

of 50% after the conclusion of the trial, and 10 out of 42 patients with measurable disease showed partial or complete response to treatment. Median radiographic PFS was 5.7 months, and a further 5 patients saw no radiographic progression for 11 months, with one patient showing ongoing complete radiographic response (at the time the results were published).

Markowski *et al* (218) carried out a further phase 2 trial in which they compared the effects of BAT in mCRPR patients after treatment with abiraterone and enzalutamide. 29 patients were enrolled after abiraterone treatment and 30 patients were enrolled after enzalutamide treatment. Patients were given a 400mg intramuscular testosterone cypionate injection every 28 days alongside ongoing LHRH (ant)agonist treatment until disease progressed, at which point the abiraterone and enzalutamide patients were re-challenged with their respective prior treatments, to determine BAT's capability for resensitising patients to these drugs. No statistically significant differences in PSA₅₀ responses were seen in either cohort after BAT alone, but the PSA₅₀ responses in the enzalutamide cohort were much higher after re-challenge (68%) than in the abiraterone cohort (16%), suggesting that BAT may be more effective at resensitising patients who undergo treatment with enzalutamide than those who undergo abiraterone treatment.

None of these clinical trials suggest any particular reason for the scheduling of these testosterone injections, but we note that the dose and frequency of these injections is consistent across every study. A potential question we might answer with our modelling is whether or not the dose scheduling and dosage itself is optimal.

1.4.6 Mathematical Models of Prostate Cancer

Although BAT itself has never been formally modelled, a variety of mathematical models, both theoretical and data-driven, have been used to investigate treatment of prostate cancers, many with a particular focus on optimising ADT treatment regimens and estimating time taken until treatment resistance emerges. Whilst a large bulk of current models focus mainly on population dynamics, some focus on cellular kinetics, and others have a primary focus of determining the efficacy of PSA as an indicator of tumour size and progression. Phan *et al*(219) provide a comprehensive review of such models – we refer the reader here for such a review, but here, we shall focus on models which have particular relevance to us when

1925 modelling BAT. This review also contains a table of data-validated parameter value ranges
 1926 for cell growth, death and mutation, as well as values corresponding to PSA production and
 1927 clearance, which will be of use when designing our own model.

1928 **1.4.7 Population Models**

1929 Jackson (220) developed a model which showed that emergence of castration resistance after
 1930 initial ADT response can be explained by the cancer being comprised of multiple cell sub-
 1931 populations, all with varying androgen dependence. Jackson uses the following PDE model,

$$\begin{aligned}\frac{dR}{dt} &= u(R(t), t), \\ D_p \frac{\partial p}{\partial r}(R, t) - u(R, t)p(R, t) &= 0, \\ D_q \frac{\partial q}{\partial r}(R, t) - u(R, t)q(R, t) &= 0,\end{aligned}$$

1932 where $u(R, t)$ is collective cellular motion, $R(t)$ is tumour radius, $p(R, t)$ is the volume
 1933 fraction of androgen dependent (AD) cells, and $q(R, t) = 1 - p(R, t)$ is the same for andro-
 1934 gen independent (AI) cells. The model assumes that an androgen-rich environment will aid
 1935 proliferation of androgen-dependent cells only, but castrate environments will increase apop-
 1936 tosis in AD cells whilst decreasing apoptosis of AI cells. The model also assumes spherical
 1937 symmetry of a tumour. It led to two conclusions: first, that castration resistance emerges in
 1938 the presence of treatment due to this disparity in apoptosis rate as opposed to an increased
 1939 proliferation rate; second, that ADT only works in the model for a small subset of model
 1940 parameters. These findings are interesting, but the assumption that castrate resistant cells are
 1941 androgen independent disagrees with the modern understanding that the AR often remains a
 1942 key component of CRPC cells even in castrate environments.

1943 Ideta *et al* (221) also modelled prostate cancer in terms of two respective populations of
 1944 AD and AI cells, but explicitly model androgen level as well, aiming to compare CADT and
 1945 IADT – the first model to do so. The ODE model is as follows;

$$\frac{da}{dt} = -\gamma a + \gamma a_0(1 - u),$$

$$\frac{dx_i}{dt} = \alpha_i p_i(a) x_i - \beta_i q_i(a) x_i + (-1)^i m(a) x_1.$$

Here, a is androgen level, x_i is the population of AD ($i = 1$) and AI ($i = 2$) cells, $u = \{0, 1\}$ determines whether ADT treatment is ongoing ($u = 1$) or halted ($u = 0$) and $p_i(a)$, $q_i(a)$ and $m(a)$ describe the effects androgen levels have on proliferation, cell and mutation from AD to AI, respectively.

The authors also explicitly model PSA level with the following equation:

$$y(t) = c_1 x_1(t) + c_2 x_2(t).$$

They consider three different assumptions about net growth rate of AI cells, assuming a constant death rate of these cells when androgen level is normal ($a = a_0$), which are as follows:

- (i) Proliferation rate of AI cells remains constant ($\alpha_2 p_2(a_0) = \alpha_2$);
- (ii) When $a = a_0$, AI cell population does not change ($\alpha_2 p_2(a_0) = \beta_2 < \alpha_2$);
- (iii) When $a = a_0$, AI cell population decreases ($\alpha_2 p_2(a_0) = 0$).

Under these assumptions, the authors found that only in the case of assumption (iii) can relapse be averted, and further that if AD cells have a competitive advantage over AI cells without any treatment, IADT is better if the aim is to prolong time to CRPC emerging (although the term “competitive advantage” here should be used loosely, since competition between cells here is not explicitly modelled). Like Jackson (220), however, the authors assume that AI cells have at best no competitive advantage over AD cells – evidence exists to suggest that this is the case when androgens are abundant (41), but since AI cells are more prevalent once CRPC develops, this may not be the case in general. A key feature of this model was explicit modelling of PSA and incorporation of androgen dependence – we shall see that several modellers have taken inspiration from these decisions.

Aiming to address the absence of competition in this model, Shimada and Aihara (222) modelled competition via a constant competition term, and demonstrated numerically that not only can IADT prolong time until CRPC emergence, but that the success of IADT increases the more of an advantage AD cells have over AI cells. Yang *et al* (61) had similar

findings with their own competitive model, but they chose instead to have their competition terms depend on androgen levels. The competition in this model was also limited by shared resources, addressing an issue with Ideta *et al*'s model taking place in an environment where cancers can grow without any limitation.

1.4.8 Data-validated Models

Hirata *et al* (223) developed a piecewise-linear dynamical model with the aim of reproducing the nonlinear trends in PSA levels during IADT. They observed that during treatment, PSA levels initially decrease sharply to a point, after which the rate of decrease itself decreases until the PSA level is undetectable. The model was inspired loosely by Ideta *et al*'s model, which explicitly model PSA. However, the authors of this model observed that describing the two-phase PSA decline with a linear model would require two linear variables with different decreasing factors, and a third to account for disease relapse. As such, their model considers three separate cell states, defined by whether or not the mutations occurring between these states are reversible or not. Cells in state one correspond to AD cells; those in state two correspond to AI cells which can revert back to androgen dependent cells during treatment intermission – biologically speaking, cells which have acquired reversible mutations during IADT; cells in state three represent AI cells which cannot regain androgen dependence, due to somatic mutations of the AR (to name an example). The resulting model, which had two distinct forms corresponding to on-treatment and off-treatment periods, respectively:

$$\frac{d}{dt} \begin{pmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \end{pmatrix} = \begin{pmatrix} w_{1,1}^1 & 0 & 0 \\ w_{1,2}^1 & w_{2,2}^1 & 0 \\ w_{1,3}^1 & w_{2,3}^1 & w_{3,3}^1 \end{pmatrix} \begin{pmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \end{pmatrix},$$

$$\frac{d}{dt} \begin{pmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \end{pmatrix} = \begin{pmatrix} w_{1,1}^0 & w_{2,1}^0 & 0 \\ 0 & w_{2,2}^0 & 0 \\ 0 & 0 & w_{3,3}^0 \end{pmatrix} \begin{pmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \end{pmatrix},$$

where $w_{i,j}$ are the transition rates from x_i to x_j , and the PSA level is represented simply as the sum of x_1 , x_2 and x_3 . This phenomenological model was able to accurately reproduce the nonlinear PSA kinetics observed in several sets of individual patient data, and was also

used to determine whether CADT or IADT would be best for patients based on PSA kinetics by taking patient PSA data and numerically estimating model parameters.

Aiming to develop a model that is more biologically tractable, Portz, Kuang and Nagy (PKN) (224) developed a model inspired by Droop's marine ecological model (225) which treats androgen supply as a limiting nutrient for the proliferation and death of cells, acting as a cell quota in the model. We have two cell subpopulations – AD and AI – which both have distinct minimum androgen requirements in order for growth to take place. The model separates intracellular androgen levels and circulating serum androgens, with the latter being interpolated using an exponential model. PSA is also explicitly modelled, giving the following:

$$\begin{aligned} \frac{dx_i}{dt} &= \underbrace{\mu_m \max\{0, 1 - \frac{q_i}{Q_i}\} x_i}_{\text{Cell growth}} - \underbrace{d_i x_i}_{\text{Cell Death}} - \underbrace{\lambda_i(Q_i) x_i}_{\text{Mutation from } x_i \text{ to } x_j} + \underbrace{\lambda_j(Q_j) x_j}_{\text{Mutation from } x_j \text{ to } x_i}, \\ \frac{dQ_i}{dt} &= \underbrace{\nu_m \frac{q_m - Q}{q_m - q_i} \frac{A}{A + \nu_h}}_{\text{Androgen Influx to Cells}} - \underbrace{\mu_m(Q_i - q_i)}_{\text{Androgen Uptake}} - \underbrace{bQ_i}_{\text{Androgen Clearance}}, \\ \frac{dP}{dt} &= \underbrace{\sum_{i=1,2} (\sigma_0 + \sigma_i \frac{Q_i^m}{Q_i^m + \rho_i^m}) x_i}_{\text{PSA Production}} - \underbrace{\delta P}_{\text{PSA Clearance}}, \end{aligned}$$

with respective AD and AI cells x_1 and x_2 , intracellular androgen concentration Q_i , PSA level P , maximum proliferation rate μ_m , serum androgen concentration A , androgen level minimum for cell proliferation q_i , death rate d_i , intracellular androgen-dependent mutation rates λ_i, λ_j , maximal androgen uptake rate ν_m , half-saturation of uptake rate ν_h , maximum androgen uptake quota q_m , androgen clearance rate b , baseline PSA production rate σ_0 , cell-specific PSA production σ_i , androgen-dependent, cell-specific PSA production parameter Q_i^m , PSA production half-saturation level ρ_i^m and PSA clearance rate δ . This model was validated with clinical data, and like the Hirata *et al's* model, it successfully reproduced trends seen in IADT clinical data and showed the capability for fitting and forecasting PSA trends.

A key limitation of PKN's model is its dependence on interpolation of androgen concentration data, meaning that predictions are difficult to make, owing to the need for some guess as to future androgen levels. Furthermore, since this model focussed solely on IADT, it failed to reproduce continuous ADT treatment dynamics. Baez and Kuang (226) offered an altered version of the PKN model, with androgen instead being included as a variable, much like Ideta *et al*'s model, which allows direct comparison between model forecasts and clinical data. Furthermore, the authors only permit mutations from AD to AI rather than allowing transitions back and forth, which allowed the model to capture disease relapse for both IADT and CADT. A key takeaway from this model was about model design itself – with cells being allowed to mutate between AD and AI and back again, tumour control is simpler, but does not guarantee the eventual emergence of resistance which is seen in reality.

1.4.9 Cellular Kinetics Models

Zhang *et al* (41), like Hirata *et al*, designed a phenomenological model of three different cell subpopulations. Their aim, however, was to investigate competition between these three subpopulations, rather than mutational processes. They use a standard Lotka-Volterra Model,

$$\frac{dx_i}{dt} = r_i x_i \left(1 - \frac{\sum_{j=1}^3 a_{ij} x_j}{K_i} \right),$$

$$\frac{dP}{dt} = \sum_{i=1}^3 x_i - 0.5P,$$

where $x_i, i \in \{1, 2, 3\}$ are AD cells, treatment-sensitive AD cells and treatment-resistant AI cells, respectively, r_i is the maximal growth rate of each cell type, a_{ij} are elements of the competition matrix, and P is the PSA level. Their aim was to demonstrate with their model, alongside a clinical study they were carrying out, whether or not IADT was better at prolonging resistance onset than CADT, since they felt that previous clinical studies were not well supported. Their simulation results suggest that treatment outcomes of IADT and CADT are similar in terms of prolonging resistance onset. They demonstrated that an adaptive intermittent treatment protocol, in which a drug dose is supplied in response to PSA progression,

2036 could delay resistance onset even more. A pilot clinical trial presented in the paper supported
2037 these findings in a small cohort of patients.

2038 West *et al* (227) expanded on this work and worked into the model the potential for using
2039 multiple drugs for treatment at once. Addressing a limitation in model parameterisation, they
2040 assume that some parameters across patients will fall within similar ranges, and so they can
2041 be fixed across a patient cohort, using instead the initial size of each cell subpopulation as
2042 a distinguishing feature for each patient. They found that using multiple drugs instead of a
2043 single drug can further delay time until treatment resistance emerges.

Chapter 2

Bayesian and Computational Modelling of Adaptive Therapy

2.1 Foreword

The results in this chapter are based entirely on analysis of experiments performed by our collaborators. However, whilst conducting these experiments, an error was made with the preparation and storage of the drugs used, rendering the results herein are unpublishable. When this error was detected, the mathematical modelling and analysis of these experimental results was only partially finished. The analysis carried out and results obtained herein were done in good faith, prior to us finding out that the experiments were flawed, at which point investigation into these results ceased.

2.2 Introduction

Mathematical modelling is a key component of the wider interdisciplinary field of mathematical oncology. Close collaboration between mathematicians, experimental biologists and clinicians is thus instrumental in the success of modelling ecological phenomena in the tumour microenvironment. The results, however, can uncover otherwise overlooked mechanisms and interactions in the tumour microenvironment (37), (228), optimise and improve upon pre-existing treatments (3), and even give us the tools to carry out simulated clinical trials without the need of live test subjects (229). The former two are of particular interest when investigating cancer adaptive therapy (AT).

A mathematical model of adaptive therapy was first presented by Robert Gatenby and collaborators in (3). The aim was to address the disparity between fixed, non-adaptive cancer treatment protocols and the dynamic nature of the tumour microenvironment. Gatenby presents an adaptive treatment protocol which changes in response to temporal and spatial variations, both inherent and treatment-induced, in the tumour microenvironment, as well as the phenotypic makeup of the tumour. An important factor in the success of such treatment regimens is that cells which resist treatment incur some "fitness cost" in comparison with the treatment-sensitive wild type. Resistant cells might require more energy or substrate to resist treatment which can no longer be dedicated to proliferation, or might be more susceptible to growth inhibiting factors, like ischemia or hypoxia, in the microenvironment. It follows that in situations where these cells are pre-existing in the microenvironment, they appear in small numbers. Typically, when a tumour is detected, a maximal tolerated dose (MTD) of a chemotherapeutic (or equivalent treatment) will be applied in order to reduce the size of this tumour as much as possible. These low-frequency resistant phenotypes will survive this MTD treatment and grow without any competitive pressure from the fitter sensitive cells, leading to eventual re-emergence of the now fully resistant, and thus untreatable, tumour. Adaptive therapy aims instead to try and preserve as much of these sensitive cells as possible to suppress proliferation of the weaker resistant phenotype, while trying to maintain a stable tumour burden. The results presented predicted prolonged survival in comparison to MTD therapies and were validated by *in vivo* experimental results. Several studies on the theoretical and clinical advantages and limitations of adaptive therapy have been conducted in (37), (38), (24). We focus in particular on the notion of containment, where we aim to keep the cell number or tumour volume approximately constant for as long as possible. Initially, it was thought that AT requires a fitness cost of resistance. We see in (24), however, that this is not essential – rather, the key assumption is that large numbers of sensitive cells reduce the growth rates of resistant phenotypes when these phenotypes are rare. When aiming for containment, then, we want to maintain a large enough sensitive cell population to prevent resistant cells from growing at the same rate that they would in isolated culture. This will allow us to prevent emergence of resistant phenotypes which we cannot treat, and in turn will allow us to keep our tumour the same rough size for as long as possible. We call the point at which

the resistant cells grow at the growth rate they display in isolation after previous inhibition by sensitive cells competitive release – it is very unlikely that we will never reach this point, but it has been shown that it is possible to prevent this from happening for longer using an adaptive protocol than a MTD protocol in the studies cited above.

Building on previous works, in collaboration with experimental biologists Medhi Morchikh, Dan Fisher and Liliana Krasinska (IGMM, CNRS, Montpellier), we conduct a combined mathematical and experimental analysis of cancer adaptive therapy with epidermal growth factor receptor tyrosine kinase inhibitors (EGFR TKIs). We first analyse the results of monolayer homogeneous and co-culture growth assays of the PC-9 non-small cell lung cancer (NSCLC) cell line under treatment with erlotinib and osimertinib, the results of which already allude to some interesting behaviours that might inform how we choose to model this data. We present a hierarchical Bayesian model which we can use to model different data sets with several replicates which share properties across different experiments, and consider a few models which we can use to analyse the results of these monolayer assays and inform future experiments.

We also analyse the results of a tumour spheroid assay conducted using the same cell lines, the objective of which was to determine what qualitative effects (if any) we see when adding the additional competitive pressure of spatial competition to the system. Spatial interactions can be difficult to model using ODEs, and so we instead turn to agent-based computational models to try and replicate behaviours seen in the spheroid experiments. Finally, we trial some preliminary models to determine whether they could feasibly fit some experimental data obtained from a set of in-vivo mouse experiments.

2.3 Methods

2.3.1 Monolayer Co-culture Assay Data Analysis with Wild Type, Erlotinib-Resistant and Osimertinib-Resistant Cells

Our collaborators provided us with several sets of experimental results from monolayer growth assays of both sensitive and resistant cell phenotypes. Cells were grown in both

	WT-OR WT-ER Growth	FACS Frequency	Exchange Media
Initial Ratios in Co-Culture	99:1, 1:1	99:1, 19:1, 9:1, 1:1	n/a
OSM Dose (nM)	0,1,10	0,1,5,10	0,5
Replicates	2	2	2
Initial Cell Number	50,000	50,000	50,000
Duration (Days)	10	18	4
Data Types	Cell Numbers, Frequencies	Frequencies	Cell Numbers

FIGURE 2.1: Tables summarising the experiments and data types they yielded for this chapter.

homogeneous and heterogeneous culture with varying drug doses, and at varying initial ratios of sensitive to resistant cells. We have several experiments to look over. The first has only a few doses and initial seeding ratios but we have access to both cell frequencies and densities, and so we shall use these results for inferring our model parameters later; the second has a wider variety of initial seeding ratios and drug doses but we only have frequencies, and no cell densities; the third set of results was to determine how cells might interact with growth medium and influence growth of other phenotypes. We present these results and the experimental procedures used to obtain them now. A summary of the 3 assays is shown in figure 2.1.

For the first set of experiments, 21 total assays were carried out using three different cell phenotypes – PC9-WT, the sensitive wild-type, OR, the osimertinib- (OSM) resistant phenotype, and ER, an erlotinib-resistant phenotype which is still sensitive to OSM treatment. 9 homogeneous assays, each with two replicates, were conducted to determine how varying treatments would affect growth of each cell line and 12 heterogeneous assays, also each with two replicates, were conducted to determine how competition affected the cell growth when combined with these varying treatments. 5000 cells were seeded at day 0 in every assay. The WT and ER cells were expressed with enhanced green fluorescent protein (EGFP), and the OR cells with mCherry (a red fluorescent RFP derivative), so that the frequencies could be determined via Fluorescence activated cell sorting (FACS) analysis. The cell number was also recorded. Each cell phenotype was grown in isolation in dimethyl sulphoxide (DMSO, the solvent in which the drug is dissolved), and then in the presence of 1nM and 10nM of OSM, the latter of which is our maximum tolerated dose (MTD) in these experiments. Growth media and drug supply were replenished every 24 hours. In order to develop a resistance to erlotinib, WT cells were grown with 0.5nM of the drug for one month, with

regular replenishing of both the medium and the drug. To develop the osimertinib-resistant cells, erlotinib-resistant cells were cultured with 0.5nM of osimertinib for one month, again with regular replenishing of the medium and drug.

2.3.2 Calculation of Selection Coefficients

Using the cell frequencies, we can calculate a selection coefficient to determine whether there was any unexpected interactions between the two cell lines. The method for calculating these selection coefficients outlined is based on the method outlined on (37), which is an extension of the method outlined in (230). We first normalise the frequencies recorded from the FACS competition assay data so that they sum to unity. For two competing cell lines, P and Q , with corresponding frequencies p and $q = 1 - p$, we define the selection coefficient of cell line P relative to cell line Q , denoted s , as the rate of change of the ratio of their frequencies: explicitly, we define s as

$$s = \frac{d}{dt} \log \frac{p}{q} = \frac{d}{dt} \log \frac{p}{1-p}.$$

If $s > 0$ then P will increase relative to Q – the opposite holds if $s < 0$.

We can also estimate the selection coefficient we might expect from a heterogeneous growth assay between two different cell lines if we use their homogeneous exponential growth rates. Given that P and Q have respective initial population sizes P_0 and Q_0 and exponential growth rates r_P and r_Q , then we have that

$$\log \frac{p}{1-p} = \log \frac{P}{Q} = \log \frac{P_0 \exp(r_P t)}{Q_0 \exp(r_Q t)} = \log \frac{P_0}{Q_0} + (r_P - r_Q)t,$$

so it follows that in this case, $s = r_P - r_Q$.

As such, for each assay, the selection coefficient was predicted to be the difference between the exponential growth rates of the cell lines in co-culture. The growth rate was estimated as the mean gradient of the log-transformed growth curves of each cell type in homogeneous culture over the first 72 hours of growth. These predicted selection coefficients were then compared to the selection coefficients we calculated from the co-culture assay frequency




FIGURE 2.2: A flowchart describing the various methods used to calculate the selection coefficient of a cell type P relative to cell type Q .

dynamics. The various methods used to calculate the selection coefficients are summarised in Figure 2.2.

2.3.3 Adjusting for Loss of EGFP and mCherry

It was observed in the FACS assay frequency data that even in monolayer culture, frequency of cell lines was not 100%. This is likely the result of natural loss of EGFP and mCherry proteins during cell growth.

Following the method outlined in (37), we estimated loss of EGFP by fitting a regression curve to the log-transformed FACS assay frequency data for WT-EGFP in monolayer culture. The adjustment factor for the WT cells, a_{WT} , was then calculated as

$$a_{WT}(t) = \exp(m_{WT}t + c_{WT}),$$

where m_{WT} is the gradient of the regression line and c_{WT} is the intercept (expressions for the ER and OR cell lines are defined analogously). We then multiplied the WT cell frequencies by $1/a_{WT}$, the ER cell frequencies by $1/a_{ER}$ and the OR frequencies by $1/a_{OR}$. Finally, we normalised these new data so that the relevant frequencies summed to unity. No changes in the qualitative outcomes were observed after adjusting for the loss of the markers.

2.3.4 Monolayer Co-culture Frequency analysis

Experimental setup here is similar to the design of the previous one, with a few key differences. Given the poor performance of the ER cells in the previous assays, we opt to study only on the WT and OR cells, again expressed, respectively, with EGFP and mCherry. In order to get a better understanding of frequency behaviour over a longer period, these assays are run for a total of eighteen days. Cells were once again grown in DMSO for the treatment-free assays and with doses of 1nM, 5nM and 10nM of OSM for the treatment assays. Every two days, 80% of the cells grown were removed and a FACS analysis carried out to determine cell frequencies. The remaining 20% of the cells were allowed to continue growing with replenished medium and drug dose. For co-culture growth assays, initial cell seeding ratios used were 99:1, 19:1, 9:1 and 1:1. Selection coefficients were calculated and normalised frequencies plotted (since we found that EGFP and mCherry loss had a negligible effect on results).

2.3.5 Monolayer Exchange Media Assays

In order to determine whether cells were affecting their growth medium during experiments, assays were conducted in which WT and OR cells were allowed to grow in homogeneous culture for 24 hours. At the end of the 24 hours, the WT and OR cells were removed and placed in the medium that the opposing cell line was grown in, and allowed to grow for a further 48 hours, after which the cell count was obtained. A control experiment was conducted in which cells were allowed to continue growing in their own medium for the remainder of the 72 hour experiment. Both experiments were carried in DMSO and with 5nM OSM. Cell count was recorded once every 24 hours.

2.3.6 General Hierarchical Model: Inference Software and Code Design

We estimate parameters using the *CmdStanR* interface to the *Stan* Package in the R programming language (231; 56). This package is widely available and allows us to perform easily reproducible Bayesian Inference for estimating parameter values without having to write custom-made code from scratch. We use the *CmdStanR* interface rather than the *Rstan* package itself because it is more regularly updated and has a more intuitive interface.

The rough structure of the code in every inference is the same, but how the code looks and how it is utilised will depend on the model we are trying to infer parameters for. For example, in our first proper inference, we only infer growth rates specific to each assay, so we run the same short model multiple times with different data to obtain these growth rates. A later model accounts for competition and treatment, and so we input all of the data at once and run the more complicated model only once. We make use of an ODE solver built into *Stan*, which itself utilises a 4th and 5th order Runge-Kutta method to solve the growth equations at every step.

We use a hierarchical model to infer parameters which gives both a general overview of the system and a tighter fit for each individual replicate. This allows our model to reduce posterior variance by pooling all of our available data. The downside of this approach is that it complicates the model geometry and reduces geometric ergodicity, which can bias inference and lead to divergent transitions, in which the sampler fails to explore the sample space effectively. To offset this, we utilise a non-centred implementation of the model, in which we fit a latent Gaussian variable which we use to recover group-level parameters rather than fitting them directly. Though we must fit multiple parameters for each replicate, we just give a brief treatment of how we infer one parameter for one experiment's replicates since the method is essentially the same for each.

In a hierarchical model, we use hyperparameters, denoted μ and τ , as the mean and variance, respectively, for sampling the possible values of the actual parameters $\theta_1 \dots \theta_n$ we are interested in, called the group-level parameters, across n replicates. When utilising a non-centred implementation, we do not use these hyperparameters directly for sampling group-level parameters, but instead recover them with a scaling and translation using latent Gaussian variables $\tilde{\theta}_1 \dots \tilde{\theta}_n$. Altogether, for $a, b, c, d \in \mathbb{R}$ and $1 \leq i \leq n$, we get the following:

$$\mu \sim \mathcal{N}(a, b),$$

$$\tau \sim \mathcal{N}(c, d),$$

$$\tilde{\theta}_i \sim \mathcal{N}(0, 1),$$

$$\theta_i = \mu + \tau \cdot \tilde{\theta}_i.$$

How we choose values for a, b, c and d is outlined separately for each different experiment. To ensure that these parameters fit our experimental data, for each replicate in each experiment, we require an additional value σ as a variance in a normal distribution, whose mean, $\hat{y}_i$, is the approximate solution to the ODE model at each step in the iteration during inference. We thus require an additional pair of hyperparameters, ν and ξ , as well as a latent Gaussian variable, $\tilde{\sigma}$, which is used in the same way as above. We thus have the following:

$$\nu \sim \mathcal{N}(0, 0.01),$$

$$\xi \sim \mathcal{N}(0, 0.01),$$

$$\tilde{\sigma} \sim \mathcal{N}(0, 1),$$

$$\sigma = \nu + \xi \cdot \tilde{\sigma},$$

$$y_i \sim \mathcal{N}(\hat{y}_i, \sigma),$$

where y_i is our estimated solution to the ODE system, which allows us to see how closely our Stan model output fits the data.

2.3.7 Spheroid Growth Experiments

In the first spheroid assay, an initial population of 1000 cells per well was seeded, with various initial ratios of WT to OR cells over a time period of 26 days. The objective here was to see how varying initial cell ratios and dosages of OSM would influence cell growth dynamics – in particular, we wanted to observe if any spatial competition seemed to be taking place. Three initial cell ratios were used: 100% WT cells, 100% OR cells and a ratio of 99:1 WT cells to OR cells. In addition to the treatment free assays (with the exception of the treatment free 100% WT, in which the spheroids were lost at day 16 of the assay), assays were carried out in the presence of 1nM and 10nM OSM dissolved in DMSO for each initial ratio. Spheroid Area was recorded and growth media exchanged every 2 days. These assays were replicated at least 5 times each. The second spheroid assay was set up in the same

fashion and carried out in the same way, but with extra sets of assays to determine if a 5nM OSM dosage or a 19:1 WT to OR initial cell ratio would have any qualitative effect. Again, these were replicated at least 5 times, with most assays being carried out at least 6 times. Some data points were missed during the course of these assays, but our collaborators did not provide a reason for this. FACS quantification was conducted at the end of these assays to determine differences in the initial and final cell ratios. Some replicates were imaged using IncuCyte, a real-time quantitative live-cell imaging and analysis system – these can be seen in figure 2.20.

2.3.8 Modelling spheroid growth in HAL

Our stochastic hybrid cellular automaton was built in Hybrid Automata Library (117) in the Java programming Language, adapted from some model examples included in the package to incorporate oxygen diffusion and budding. A short custom script was written into the HAL *Util* package which allows for more straightforward cell budding, and the cell step structured to allow for more experimentation in the maximal size of the division neighbourhoods. Although HAL has a built-in method for cell disposal, this removes the cell entirely from the grid, and so we add a few lines of code which allow us to keep dead cells on the grid – we can easily toggle this on and off if need be. The random number generator “rng” uses the java function “SplittableRandom” to generate a pseudorandom number between 0 and 1 on a uniform distribution. A summary of the boundary values and parameter values used is included in figure 2.3. We used a simplified competitive release model included in HAL as a basis for our own model (117).

HAL is highly customisable. We can schedule events however we like as a result. For simplicity’s sake, however, we use the “default” method built into HAL which can be seen in several of the examples included therein. At every model step, each living agent has a probability of division or death which we calculate at every time step (we outline how below). At the next model step, we shuffle the list of agents in some random order so that we do not have a spheroid which only grows in one direction at a time.

In our model, cells can “budge” other cells around them across a short distance to account for spatial competition. If we extend this distance, it relaxes the effects of competition, since

cells can budge cells further and further away to make space to divide: if we decrease it, then spatial competitive pressure increases since cells have fewer opportunities to make space to divide. We use an ADI approximation method to evaluate PDE diffusible fields at every model step, one of the three methods built into HAL. We calculate chances of proliferation and death, as well as drug and oxygen consumption rates using the methods outlined in (37), which we will briefly outline below. We run our model on a 2D grid – the result effectively corresponds to a 2D slice of a 3D spheroid.

Oxygen consumption and cell death to hypoxia (an insufficient oxygen supply) both depend on the local oxygen concentration relative to medium concentration of 2.2×10^{-7} – we call this g . Before every cell attempts proliferation, it consumes an amount Q of oxygen, where, given that Q_{max} is the maximum possible consumption rate,

$$Q = Q_{max} \frac{48g}{48g + 1}.$$

If $g < 0.001$ then cells enter a state of quiescence – in this state, given ample oxygen supply, cells can turn back into actively dividing cells, but still need ample space to divide. If $g < 0.0001$, then cells will die due to lack of oxygen, resulting in a necrotic core of cells. We define drug consumption rate, R , as

$$R = R_{max} \frac{48g}{48g + 1},$$

where R_{max} is maximal drug consumption.

To calculate drug and oxygen effects on cell proliferation, P , and death, M , we first estimate impact of treatment. Assuming a plentiful supply of oxygen, we want $P = 1$ in absence of treatment. For now, we shall assume that proliferation is best in absence of treatment for both cell lines – more evidence is needed to decide how low dose treatments may aid in OR cell proliferation, since we have not got enough cell density data with varying initial WT-OR cell ratios. Effects of treatment will vary depending on cell line and drug dose. Initially, we used our first inferred model inference to model treatment in our ABM, but this led to an excessive treatment effect at a 1nM dose of OSM. As such, we use an alternative approach, in which we model treatment using a Hill function. This sees a much

better agreement with our 3D spheroid images, but we note now that due to a lack of data, we cannot infer parameters for this hill function as accurately as we would like – for now, the parameters used in this function are essentially ballpark estimates. We also note that we have not carried out any formal qualitative comparisons between the model outputs and the spheroid data. Regardless, we use this Hill function to calculate a proliferation chance of

$$P = P_{max}(\frac{30g}{30g + 1})(\frac{1}{1 + \frac{\delta}{IC50}^\beta}),$$

where $P_{max} = 31/30$, $IC50 = 5.75$ for WT cells and $0.9*31/30$, $IC50 = 12$ for OR cells. We include this impact on P_{max} on OR cells since our inference from **Fig. 16** shows us that in 1:1 competition, the hyperparameter mean growth rate of OR cells is roughly 90% of that of the WT cells in absence of treatment.

We have from the 3D modelling images in **Fig. 25** that even though our drug is cytostatic, at high doses, there are clearly a higher number of cells which do not express our fluorescent marker – more so than the 1 – 2% we have seen previously. Hence, we can assume that some number of cells are dying as a result of high dose treatment. Since we have insufficient data to estimate the number of cells which are killed by high-dose treatment, save for the fact that nearly every cell dies under a dose of 100nM of OSM regardless of phenotype, the following estimation is not particularly well driven by the data. We want chance of cell death to be slightly greater than proliferation chance at a 10nM dose for WT cells, and we reasonably assume that cell death at much higher dose – say, 25nM – is at least 50% – a conservative estimate. We also assume that since resistant cells are able to withstand treatment at the doses we are looking at despite their partial sensitivity to OSM, we do not add a chance of cell death for these cells. We also assume that cells do not die as a result of drug dose at low doses or in absence of treatment. We have

$$M = \max(0.0236\delta - 0.09, 0).$$

There are two main parts to the model – the Cell Step, which governs how agents on the grid behave, and the grid setup, which defines the framework for whole model and governs how agents interact with the grid, as well as how diffusible factors in the model behave (in

Parameter	Value
Timestep	1 hour per tick
Maximal Sensitive Prolif. P_{max}	31/30
Maximal Resistant Prolif. P_{max}	$0.95 * 31/30$
$IC50_{WT}$	5.75×10^{11}
$IC50_{OR}$	12×10^{11}
Hill Coefficient β	5
Drug Dose, δ (nM)	$\delta \times 10^{11}/, \delta = 0, 1, 5, 10$
Sensitive Cell kill	$Max(-0.0236 \times 10^{11} * \delta - 0.09 \times 10^{11})$
Maximal Drug Uptake, R_{max} ($mol\ cm^{-3}s^{-1}$)	$1 \times 10^{-11}/1 \times 10^{-8}$
Drug Diffusion ($cm^{-2}s^{-1}$)	5×10^{-6}
Relative Max. Oxygen Concentration ($\mu mol\ cm^{-3}$)	$2.2 \times 10^{-7}/2.2 \times 10^{-7}$ (1)
Oxygen Diffusion ($cm^{-2}s^{-1}$)	2×10^{-5}
Maximal Oxygen Uptake ($\mu mol\ cm^{-3}s^{-1}$)	$5 \times 10^{-9}/1 \times 10^{-8}$ (0.5)
Proportion of Resistant cells	0.05
Initial Tumour Radius	6 cells
Grid Area	500^2 cells

FIGURE 2.3: Parameter values used for the HAL model simulations.

particular, how these factors behave when agents are not consuming them). Algorithm 1 below outlines how the grid setup works, and Algorithm 2 explains the Cell Step. We show a few images of the results of running this model with a 19:1 initial cell ratio and with 0nM, 1nM, 5nM and 10nM doses of OSM in figure 2.23.

2.3.9 Pilot In Vivo Mouse Experiments

Two experiments were carried out, but data was only provided for the second set of experiments. These mirror the first in most ways, but a higher drug dose was used in both the AT and MTD treatment protocols. Additionally, these treatment protocols were also given to a set of mice whose tumours were comprised entirely of drug resistant cells, to determine what returns, if any, the AT protocol could yield. Flow cytometry analysis was carried out at the beginning (which required the sacrifice of 1 mouse per experiment) and the very end of these experiments to determine the proportions of resistant, sensitive, and dead (or marker negative) cells. We note that drug dose in the AT protocol was reduced from $2mgkg^{-1}day^{-1}$ to $1mgkg^{-1}day^{-1}$ at around day 40, before being increased again to the original dose.

Algorithm 1: Grid Setup

```

1 Define Parameters as outlined above
2 Initialise Drug and Oxygen Diffusion Fields
3 Initialise Random Number Generator rng
4 Function CompGrid:
5   Initialise new rng
6   Set oxygen and drug fields to initial conditions
7 Function Main:
8   Define Grid Dimensions, initial tumour radius
9   Set up Visual Options for model visualisation
10  Initialise Model according to Init
11  Set different drug boundary values for each model (according to dosage)
12  for  $0 \leq tick \leq 720$  do
13    run the model once per tick (one tick == one hour)
14    Update model visualisation
15 Function Init:
16   Create circular tumour with Initial Radius as defined above in the center of the
   grid and with the above defined initial proportion of resistant cells
17 Function ModelStep:
18   Shuffle agent list
19   Execute Cell Step on each cell in the newly shuffled order
20   update drug diffusion fields
21   update oxygen diffusion fields
22   Plot model for visualisation depending on most recent run of model step

```

Parameter	Description	Value
W_0	Initial value of W	0.95, 0
R_0	Initial value of R	0.05, 1
a	Growth rate of WT cells	0.18
b	Growth rate of OR cells	0.18
θ	Period between treatments (days)	$\frac{7}{5}$
δ	OSM Dose ($mgkg^{-1}day^{-1}$)	0, 2, 20
$IC50_W$	Half-maximal inhibitory dose of W ($mgkg^{-1}day^{-1}$)	5
$IC50_R$	Half-maximal inhibitory dose of R ($mgkg^{-1}day^{-1}$)	1000

FIGURE 2.4: Parameter values used for hypothesis testing in the null model.

Algorithm 2: Cell Step

```

1 Calculate  $g$ , relative oxygen concentration
2 Calculate Drug, Oxygen Uptakes
3 Update Drug, Oxygen supply on grid square containing live agent before
  proliferation
4 if Agent is Resistant then
5   Calculate and assign probability of division, quiescence and death depending on
    $g$ , and current drug supply, as well as resistant cell parameters
6 else if Agent is Sensitive then
7   Calculate and assign probability of division, quiescence and death depending on
    $g$ , and current drug supply, as well as sensitive cell parameters
8 else
9    $deathprob \leftarrow 1, divprob \leftarrow 0$  // Dead Cells cannot divide
10 if  $g < 0.0001$  OR  $rng < deathprob$  then
11   Agent Dies
12 if  $rng < divprob$  AND Grid Square non-empty AND Agent is Alive then
13    $int\ ring = 0, options = 0$ 
14    $double\ r = 0.5, R = 1.5$ 
15   while  $R \leq MaxRadius$  do
16     if  $options > 0$  then
17       BREAK;
18     else
19        $DivHood \leftarrow CircleHood(IncludeOrigin = FALSE, Radius = R)$ 
20        $options \leftarrow$  Number of Empty Spaces in  $DivHood$ 
21        $ring \leftarrow ring + 1\ R = ring + 1 + r$ 
22   Define Array  $RandomiseOptions[0, 1, \dots, options]$ 
23   Shuffle Order of  $RandomiseOptions$ 
24   if  $options > 0$  then
25     for  $0 \leq DivOption \leq options$  do
26        $opt \leftarrow RandomiseOptions[DivOption]$ 
27       // index of potential division / budge site
28        $x_{opt} \leftarrow$  X coord. of  $opt + 0.5$ 
29        $y_{opt} \leftarrow$  Y. Coord of  $opt + 0.5$ 
30        $d \leftarrow$  distance from dividing cell to  $(x_{opt}, y_{opt})$ 
31       if  $d > r + R$  then
32          $CheckProb \leftarrow 0$ 
33       else if  $R \geq r + R$  then
34          $CheckProb \leftarrow 1$ 
35       else
36          $CheckProb \leftarrow$  area of square within division radius
37       if  $rng < CheckProb$  then
38         if  $d \leq \sqrt{2}$  then
39           Agent divides into  $DivHood[opt]$ , no budging required
40           BREAK;
41         else
42           Define Array  $BudgeLine \leftarrow (x, y)$  -Coordinates of squares
           between dividing agent and empty square
43           Budge Each Agent lying on  $BudgeLine$  into the nearest empty
           space
44           Agent divides into newly created space
45           BREAK;

```

Parameter	Description	Value
W_0	Initial value of W	0.95, 0
R_0	Initial value of R	0.04, 0.8
F_0	Initial value of F	0.01, 0.2
a	Growth rate of WT cells	0.18
b	Growth rate of partially sensitive OR cells	0.18
c	Growth rate of fully sensitive OR cells	0.18
θ	Period between treatments (days)	$\frac{7}{5}$
δ	OSM Dose ($mgkg^{-1}day^{-1}$)	0, 2, 20
$IC50_W$	Half-maximal inhibitory dose of W ($mgkg^{-1}day^{-1}$)	5
$IC50_R$	Half-maximal inhibitory dose of R ($mgkg^{-1}day^{-1}$)	20

FIGURE 2.5: Parameter values used for hypothesis testing in the Heterogeneous OR Resistance Model.

Parameter	Description	Value
W_0	Initial value of W	0.95, 0
R_0	Initial value of R	0.05, 1
I_0	Initial value of I	1
a	Growth rate of WT cells	0.18
b	Growth rate of partially sensitive OR cells	0.18
c	Growth rate of Immune Cells	0.18
d	Immune cell impact via interaction with W, R	1
g	Cancer cell impact via Immune response	0.5
θ	Period between treatments (days)	$\frac{7}{5}$
δ	OSM Dose ($mgkg^{-1}day^{-1}$)	0, 2, 20
$IC50_W$	Half-maximal inhibitory dose of W ($mgkg^{-1}day^{-1}$)	5
$IC50_R$	Half-maximal inhibitory dose of R ($mgkg^{-1}day^{-1}$)	20
$IC50_I$	Half-maximal inhibitory dose of I ($mgkg^{-1}day^{-1}$)	5

FIGURE 2.6: Parameter values used for hypothesis testing in the lotka-volterra immune model.

2.4 Hypotheses and Results

2.4.1 OR Resistant Cells Do Not Pay a Fitness Cost Of Resistance


We first analyse the homogeneous growth assays to determine whether or not we actually have any fitness cost of resistance. We pay particular attention to the growth in the first few days, since there should be nothing inhibiting growth at this stage and confluence should not be reached and so growth should be exponential. We can see in the growth plots in Figure 2.7 that when grown in Dimethyl Sulphoxide (DMSO) (the solvent in which the drug is dissolved for treatment, which we use in our controls), cells do indeed grow exponentially for the first 4 days of the assay. We estimated from these log-transformed growth curves that the growth rates of the sensitive wild-type (WT), osimertinib resistant (OR) and erlotinib resistant (ER) cell lines were, respectively, around 0.66, 0.87 and 0.81 – so in fact, in the early growth phase, resistant cell lines actually have a stronger rate of proliferation.

We see that at our lower dose, ER cells suffer the greatest impact on growth. This is likely due to the combined reduced proliferative fitness incurred by development of resistance and their sensitivity to treatment with osimertinib. WT cells are impacted, but the dose is not high enough to prevent cell growth completely. The OR cell line also sees a substantially higher final cell number at a low dose when compared to the homogeneous cell cultures and the other cell lines at this dose.

At our high dose, we see that both sensitive cells and ER cells fail to grow much beyond their initial value, which is unsurprising. The OR cells are still able to proliferate, but at a reduced rate when compared to the other two assays. This suggests that these cells display partial sensitivity to OSM, which we note when modelling later, since we can no longer make the assumption that resistant cells see no impact from treatment.

2.4.2 Initial Frequency Assays Do Not Indicate Abnormal Frequency-Dependent Behaviour

In figure 2.8, we see the frequency data and selection coefficients for the 99:1 co-culture assays at every dose. The results of this assay are unsurprising – we saw in figure 2.7 that the exponential growth rates of each cell line at 0nM and 1nM are all similar, and so changes



Growth Rates	WT	OR	ER
0nM	0.6682	0.8716	0.8073
1nM	0.7636	0.7870	0.7127
10nM	0.0681	0.6663	-0.0070

FIGURE 2.7: Plot showing homogeneous growth curves for the WT (green), ER (blue) and OR (red) cells at varying doses of OSM (both replicates). Both columns display the same data – the left column has a linear scale on the y -axis, the right has a log scale. Table shows estimated growth curves from the first 72 hours of growth.




FIGURE 2.8: Frequency Data and Selection Coefficients for the 99:1 assays for both WT-OR replicates. Top row shows WT frequencies (green) and OR frequencies (red) on a linear scale. The middle row shows The selection coefficient of the WT cells relative to the OR cells. Dashed line represents estimated selection coefficient, black points show calculated selection coefficients from timepoint $t - 1$ to timepoint t , and bold line is the mean average of these points. The bottom row shows the same frequencies as the top row in logit scale.




FIGURE 2.9: Frequency Data and Selection Coefficients for the 1:1 assays for both WT-OR replicates. First and fourth rows show WT frequencies (green) and OR frequencies (red) on a linear scale. The second and fifth rows show the selection coefficients of the WT cells relative to the OR cells. Dashed lines represent estimated selection coefficient, black points show calculated selection coefficients from timepoint $t - 1$ to timepoint t , and bold lines are the mean average of these points. The bottom row shows the same frequencies as the top row in logit scale. Frequency data was only available for one replicate in this assay.




FIGURE 2.10: Frequency Data and Selection Coefficients for the 99:1 assays for both ER-OR replicates. Top row shows WT frequencies (green) and OR frequencies (red) on a linear scale. The middle row shows The selection coefficient of the WT cells relative to the OR cells. Dashed line represents estimated selection coefficient, black points show calculated selection coefficients from timepoint $t-1$ to timepoint t , and bold line is the mean average of these points. The bottom row shows the same frequencies as the top row in logit scale.




FIGURE 2.11: Frequency Data and Selection Coefficients for the 1:1 assays for both ER-OR replicates. Top row shows WT frequencies (green) and OR frequencies (red) on a linear scale. The middle row shows The selection coefficient of the WT cells relative to the OR cells. Dashed line represents estimated selection coefficient, black points show calculated selection coefficients from timepoint $t-1$ to timepoint t , and bold line is the mean average of these points. The bottom row shows the same frequencies as the top row in logit scale. Frequency data was only available for one replicate in this assay.

in frequency here likely will not occur. Calculated selection coefficients at all doses do not differ substantially from the predicted selection coefficient, indicating that there are no frequency-dependent effects occurring that are not immediately obvious from observing the data.

The 1:1 assay represents a case where we expect AT to fail, given that resistance is no longer rare here. During AT, frequency of resistant cells will gradually increase as the treatment proceeds over time. At this point in the treatment, using osimertinib as we did previously used would be inappropriate since half of the cell population resists the drug. In figure 2.9, we see the frequency of the OR cells overtake that of the WT. Given that the OR cell growth rate is much greater than the WT cell's growth rate at this dose, this is unsurprising.

In figure 2.10, we see the results of the 99:1 co-culture assays. The results line up with our findings from the homogeneous growth assays – the reduced fitness of ER cells compared to the OR cells combined with their sensitivity leads to a steady increase of OR cell frequency at a 10nM dose. We can even see that the OR frequency begins to steadily increase at our low dose (which is much clearer on our logit scale plots). Neither of these results are surprising given that the OR cells out-grew the ER cells in homogeneous culture – in fact, OR cells benefitted in the long term from a low dose of OSM. The results in figure 2.11 line up with what we would anticipate – ER frequency is overtaken by OR frequency at 1nM and 10nM, no doubt due to the difference between the growth rates at these doses.

We now analyse the second round of frequency assays, in which we look at a longer time period and a wider range of drug doses, but with only the WT and OR cells. We see in figure 2.12 that in the 99:1 competitive assays, without treatment or with low dose treatment, we see no new behaviour that was not already seen in the previous experiment. Now that the experiment goes on for a greater length of time, however, we see that at 10nM, OR cells frequency does eventually exceed that of the WT cells. In fact, this we see this clearly immediately after our first experiment ended. We see the same thing with a 5nM dose, but at a later point in the assay. What is interesting here is that the switch in frequencies is almost instantaneous rather than gradual. For example, in our 99:1 assays, WT cells occupy almost the entire well with frequencies at or around 99%. Then, suddenly, OR cells dominate the




FIGURE 2.12: Frequency Data and Selection Coefficients for the 99:1 assays for both WT-OR replicates. The top three rows show the results of the first replicate, and the bottom three the second. For each group of plots for each replicate, the top row shows WT frequencies (green) and OR frequencies (red) on a linear scale, the middle row shows the selection coefficient of the WT cells relative to the OR cells (Dashed line represents estimated selection coefficient, black points show calculated selection coefficients from timepoint $t - 1$ to timepoint t , and bold line is the mean average of these points) and the bottom row shows the same frequencies as the top row in logit scale.




FIGURE 2.13: Frequency Data and Selection Coefficients for the 1:1 assays for both WT-OR replicates. The top three rows show the results of the first replicate, and the bottom three the second. For each group of plots for each replicate, the top row shows WT frequencies (green) and OR frequencies (red) on a linear scale, the middle row shows the selection coefficient of the WT cells relative to the OR cells (Dashed line represents estimated selection coefficient, black points show calculated selection coefficients from timepoint $t - 1$ to timepoint t , and bold line is the mean average of these points) and the bottom row shows the same frequencies as the top row in logit scale.




FIGURE 2.14: Results of the monolayer exchanged media assays with and without OSM. Top left shows WT cell growth in WT and OR media in absence of treatment; top right shows WT growth in WT and OR media with 5nM OSM; bottom left and right show the same, respectively, for OR cell growth.

well instead. Were this just due to a disparity between growth rates, we would expect this switch to be more gradual. For brevity, we attach the results of the 19:1 assays and the 9:1 assays in appendix A, since behaviour here is the same as in the 99:1 assay. We see the same unusual instantaneous frequency “switching” in these cases. The 1:1 assay results in figure 2.13 display similar behaviour, with the frequency switch occurring the earliest in these assays at each respective dose.

2.4.3 Cells affect Growth Medium to Alter Growth Behaviour

In figure 2.14, we can see the results of the exchanged media assays. We see that in absence of drug, the WT cells grow in exactly the same way regardless of growth medium. Results of the remaining assays were less straightforward, however. When treated with a 5nM dose of OSM, the WT cells actually grew better in the OR medium than in their own, which suggests that something in the OR medium is either protecting the WT cells from the drug or encouraging faster proliferation (relative to the growth rate in their own medium). WT cells still grew slower in the presence of the drug in both medium than they did without any drug. The OR cells grew even faster than the WT cells when introduced to the WT medium, and we see that OR cells, like the WT cells, grow more successfully in the OR medium when supplied with a drug dose. These results suggest that the presence of different cells in the growth medium may affect the behaviour of these cells when medium is not replaced regularly, which would be the case *in vivo*.

2.4.4 Bayesian Hierarchical Modelling Yields Posterior Distributions of Model Parameters

We now infer parameters from this data using our Bayesian hierarchical modelling framework. The statistical model underpinning this framework is outlined in the methods section.

For WT cell population W and OR cell population R , we have

$$\begin{aligned}\frac{dW}{dt} &= \lambda_W W(t) \frac{\log(K/N(t))}{\log(K/N_0)}, \\ \frac{dR}{dt} &= \lambda_R R(t) \frac{\log(K/N(t))}{\log(K/N_0)},\end{aligned}$$

where $K = 1.2 \times 10^8$ is a fixed quantity chosen based off of the results shown in figure 2.7, $N(t) = W(t) + R(t)$ is the total cell number at time t , $N_0 = N(0)$ and λ_W, λ_R are the respective growth rates of W and R which we aim to infer. When inferring parameters with this model, we intentionally choose as uninformative a prior as possible, in a bid to minimise the possibility of the prior biasing inference. For all of the inferences in this section, we place a uniform prior between 0 and 3 on every growth parameter. We infer all of the replicates

of one assay at a time rather than inferring over all of the data at once, so as to zero in on exactly how the growth rates in each assay behave.

We see in Figures 2.15, 2.16, 2.17 and 2.18 that our inferences converge to a clear posterior distribution. There were only negligible divergent transitions (less than 10 out of 3000, if any). $\hat{R}$, a measure of how well chains have mixed, was within acceptable bounds. Altogether, we can be confident that these values have been reliably inferred. In appendix A, we present histograms for each group-level posterior as well as a comparison between the data and inferred mean cell populations with error bands for each fit inferred. In general, the fit agrees well with the data.

The growth rates in homogeneous culture steadily declines as the drug dose increases, with only negligible differences between WT and OR cell lines. OR cells showed the improved growth rate at a low doses that we saw previously. The growth rate of OR cells is notably impacted in both the 99:1 and 1:1 assays when compared to the homogeneous growth rate. We see improved growth rate of both WT and OR cells in the 1:1 assay under low dose treatment, which, along with the results of the exchange media assays, suggest that something is being expressed by OR cells into the medium that is either hindering drug effectiveness or aiding overall cell growth.

Our exploratory data analysis predicted that the relationship between growth rate and drug dose is roughly linear. We make a better estimation of this behaviour by taking 1000 random draws from each of the homogeneous growth posterior distributions, plotting them against the drug dose applied, and fitting a line to this data using the *polyfit* function in Numpy (232), a python package for a variety of generic mathematical applications. We do this 100 times and then take an average of the gradients and y-intercepts to give a good behavioural estimate. The y-intercept here corresponds to the homogeneous growth rate, essentially treating the homogeneous, drug-free growth rate as a constant once inferred. We see the results in figure 2.19. It is still reasonable to assume based on this that drug response in our model is linear. We shall, for the moment, assume that even though we see increased growth rates at low doses for the 1nM homogeneous OR assays, that the response here, too, is linear, until we have a sufficient amount of data to say otherwise. We note, however, that this trend in the data will no doubt present itself in the inferences, most likely as a higher




FIGURE 2.15: (Top) Histogram plots of prior and posterior distributions of the hyperparameters for the isolated monolayer WT culture data. (Bottom) Hyperparameter and group level (GL) posterior distributions for both replicates in the Assay.




FIGURE 2.16: (Top) Histogram plots of prior and posterior distributions of the hyperparameters for the isolated monolayer OR culture data. (Bottom) Hyperparameter and group level (GL) posterior distributions for both replicates in the Assay.




FIGURE 2.17: **(Top)** Histogram plots of prior and posterior distributions of the Hyperparameters for the 99:1 WT:OR monolayer co-culture experimental data. **(Bottom)** Hyperparameter and group level (GL) posterior distributions for both replicates in each assay.




FIGURE 2.18: (Top) Histogram plots of prior and posterior distributions of the hyperparameters for the 1:1 WT:OR monolayer co-culture experimental data. (Bottom) Hyperparameter and group level (GL) posterior distributions for both replicates in the Assay.




FIGURE 2.19: **(Top)** Plotting random draws from the WT homogeneous growth posteriors against dose, and the resulting line of best fit. **(Middle)** Plotting random draws from the OR homogeneous growth posteriors against dose, and the resulting line of best fit. **(Bottom)** Parameters for best fit lines.

group-level mean growth rate at 1nM dosage in OR cells, which will also result in higher hyperposterior means and variances to accommodate this.

2.4.5 In Vitro Tumour Spheroid Data Indicates Spatial Competition Between PC-9 Phenotypes

Fitting models to monolayer data is useful for predicting future experimental results, and other analysis of these results is useful for zeroing in on cell-cell interactions and cell mechanisms without confounding results, but these assays fail to capture any spatial interaction between the different cell phenotypes. Naturally, when considering the tumour microenvironment *in vivo*, we must take these spatial interactions into account. As such, our collaborators provided us with a large amount of tumour spheroid data.

We have several experiments, each with a several of replicates to analyse. In the interest of brevity, and due to a lack of major qualitative between replicates, we shall present average spheroid area growth curves – these are shown for the first set of experiments in figure 2.21. We present the homogeneous WT growth, 99:1 growth curve and 19:1 growth curve averages from the second experiment, since there were fewer missing data points as well as a larger number of replicates. The OR growth curve averages are taken from the first experiment. Full data is available in appendix A.

Looking first at the 100% WT assay plots, we see that the cells have a roughly exponential initial growth, but at around the day 20 mark, the growth seems to arrest and the spheroid remains roughly the same size. We see similar tapering off of the spheroid area in all of the DMSO assays, albeit at different time points. When grown in the presence of 1nM of OSM, we see that in all cases where the WT is present, the spheroid still grows to this tapering off point, but takes longer to do so. In contrast, we see that at dosages of 5nM and 10nM of OSM, the WT cells see less successful growth than at lower concentrations. This suggests that treatment with OSM impacts the carrying capacity as well as the growth rate and that the impact of both varies as the dose increases (in particular, the carrying capacity is not particularly sensitive to lower OSM dosages). The OR cells see very similar growth in DMSO and in 1nM. At 10nM, we see that though growth rate is notably impacted, which is to be expected, since OR cells are not fully resistant to OSM.




FIGURE 2.20: IncuCyte Imaging of tumour spheroids at various time points. WT cells are green, OR cells are Red, and the remaining dark-coloured cells are marker negative. The top image shows results of a few of the spheroids from the first assay, and the bottom image shows some results from the second.




FIGURE 2.21: Average growth curves of the 3D Spheroid Assays, both in linear and log scales on the y -axis.

% WT Cells	% OR Cells	OSM Dose (nM)	Growth Rate
100	0	0	0.284
100	0	1	0.224
100	0	5	0.098
100	0	10	0.123
99	1	0	0.272
99	1	1	0.212
99	1	5	0.120
99	1	10	0.133
95	5	0	0.291
95	5	1	0.257
95	5	5	0.149
95	5	10	0.134
0	100	0	0.311
0	100	1	0.342
0	100	10	0.156

Initial WT:OR Ratio	OSM Dose (nM)	Final WT Cell %	Final OR Cell %	Negative Cell %
100:0	0	98.10	0	1.90
100:0	1	96.85	0	3.15
100:0	5	81.80	0	18.20
100:0	10	82.90	0	17.10
99:1	0	98.65	0.10	1.20
99:1	1	97.50	0.35	2.10
99:1	5	77.55	1.95	20.55
99:1	10	6.50	89.10	4.40
19:1	0	95.90	1.40	2.65
19:1	1	95.35	1.55	3.05
19:1	5	61.45	21.95	16.30
19:1	10	4.35	83.20	12.40
0:100	0	0	98.6	1.40
0:100	1	0	99.10	0.90
0:100	10	0	82.40	17.30

FIGURE 2.22: (Top) Log-Transformed growth rates of the tumour spheroid curves. (Bottom) FACS data at the end of the Assay.

In the isolated WT Assay, we see little difference in the data obtained for the DMSO assay and the 1nM OSM assay – the spheroids grow unimpeded and reach roughly the same size. The FACS data is as expected – percentages are less than 10% due to the aforementioned natural loss of fluorescent marker. At 5nM OSM, we see a significant reduction in tumour size and growth rate across the assay, and a higher proportion of EGFP negative cells at the end. Though there is some loss of EGFP as a result of cells processing the fluorescent marker over time, over 20% of the cells are marker-negative, which is far more than we expect to see from cells processing the marker – this is instead due to cell death caused by the higher dose of OSM. At 10nM of OSM, we see that large numbers of sensitive cells are killed by the high drug dose and a huge reduction in growth rate is observed when compared to the control. In the homogeneous OSM assay, we see similar behaviour to the WT assay – the growth of the DMSO assay and the 1nM OSM assay are very similar and the sizes of the spheroids at the end of the experiment are roughly the same in both. At 10nM OSM, we again see a higher marker-negative cell proportion and a smaller tumour (but not nearly as small as in the WT assay).

In the co-culture assays, we see that the size of the spheroids are again roughly the same at the end of the DMSO and 1nM OSM assays. We note, however, that in both the 99:1 and 19:1 assays, the final OR cell proportion is lower than the initial one. In the 5nM assay, we see at the day 26 mark that the spheroids are notably smaller. However, the OR proportions in both the 99:1 and 19:1 assays are higher at day 26 than at start of the experiment (and in the 19:1 assay, the OR cells make up almost a quarter of the spheroid). In the 10nM assays, WT cells fail to grow at any substantial rate and OR cells quickly overtake them in terms of cell density even though they had a lower initial frequency.

The results of the co-culture assays are generally promising. The resistant cells make up a lower proportion of the spheroid at the end of the DMSO and 1nM assays than they do initially. This is likely due to spatial competition between cell lines. Both cell lines had similar homogeneous growth rates (in fact, OR cells consistently had better growth rates than WT cells, especially under treatment), and so *a priori* we would expect roughly the same proportion of WT cells to OR cells at the end as there are at the beginning of the experiments (if not a slightly higher OR proportion). Since this is not the case, however, it is likely that

spatial competition is a taking place.

The issue here is that the size of the tumour under 1nM of treatment is roughly the same as it would be if left untreated. Ideally, we could just leave the tumour spheroid to keep growing until the sensitive cells completely eliminate the OR cells, but this could take a very long time. In a proper clinical setting, waiting for this to happen could involve allowing the tumour to grow to a dangerous size. Treatment must be applied to keep the tumour at an acceptable size. What we deign acceptable, however, is difficult to quantify exactly – it could be keeping the tumour as small as possible whilst keeping the sensitive proportion as high as we can; it could be that we simply wish to maximise survival time whatever the cost; it could be that we wish to contain the resistant population for as long as is feasible. If we take it to be the former, the results we have for the 5nM competitive assays will likely be closer to an acceptable treatment regimen. The tumour is much smaller than in the DMSO and 1nM assays, and even though the resistant proportion is higher at the end of the assay than it is in the beginning, the resistant population is still contained (so, indeed, this would also satisfy the latter definition of acceptable treatment). We would hypothesise that the proportion the resistant cell population makes up of the total spheroid would continue to increase as the assay goes on, and eventually there would be some competitive release, with the tumour eventually reaching a similar size to the spheroids in the homogeneous OR assays. The crucial point here being that this happens “eventually” – we reasonably assume that it would happen later than it would without competitive pressure or any form of treatment. Too high a dose of OSM will cause this “competitive release” to occur sooner.

In the 10nM competitive assays, we see that the final FACS reading, shown in in figure 2.22, shows a higher final proportion of OR cells than there are WT cells. The IncuCyte imaging results in figure 2.20 give us a better look at when this competitive release occurs – at around day 18 in both cases, we see a larger area of red resistant cells than green resistant ones.

2.4.6 Using a hybrid cellular automaton model to capture qualitative behaviour seen in Spheroid Image Data

As previously mentioned, though a hierarchical model may work given the number of replicates, the complex spatial interactions would either render such a model difficult to work with. We would frequently need to model diffusible PDE fields and account for the interaction between the cells and diffusible factors. It would be better, then, to design a spatial, agent-based model.

We need a model which can account for cell growth which depends on cell-cell interactions as well as cell-medium interactions. Chances of cell proliferation should depend on local oxygen and drug concentration as well as phenotype. Cells should also only be allowed to proliferate when there is space to do so. As such, we present a stochastic hybrid cellular automaton model, built using Hybrid Automata Library (HAL)(117). The particulars of this model are discussed in the methods section.

We see in figure 2.23 the results of running the HAL model where we model the drug concentrations to be consistent with the spheroid assays. From a purely visual perspective, our model results are qualitatively consistent with the spheroid image data. Spheroids at 0nM and 1nM OSM exhibit similar behaviour, with WT cells effectively surrounding OR cells and preventing them from proliferating much beyond their initial density. At 5nM, spheroid density stabilises, but the lower overall density brings with it a higher OR cell frequency. We also note that this is only one visualisation of a stochastic process – the results of the 5nM simulations are highly variable. Sometimes, OR cells will be surrounded and eventually eradicated, other times they will make up as much as 40-50% of the tumour spheroid.

2.4.7 In Vivo Data Shows that a Low-Dose Treatment Protocol Allows for A Constant Tumour Burden

Results of the mice experiments using tumours with a 19:1 ratio of sensitive wild type (WT) cells to OR cells can be seen in figure 2.24. The first result of note here is that with the AT protocol, tumour volume roughly stabilises, unlike the MTD and vehicle cases, meaning that tumour control here has been achieved. Here, WT cells are abundant enough to contain the OR cells via what we presume to be selective competition, but treatment is not so intense




FIGURE 2.23: HAL modelling results at the same timepoints as the images taken in figure 2.20, with Day 2 at the top and day 22 at the bottom. From left to right, constant drug doses of 0nM, 1nM, 5nM and 10nM of OSM are given. Initial cell ratio is 19:1. Green squares represent WT cells; red squares represent OR cells; light blue squares represent WT cells undergoing quiescence; magenta squares represent OR cells undergoing quiescence; grey squares represent dead cells which are unable to divide due to lack of oxygen. Backgrounds of each models are heat maps representing drug concentration.




FIGURE 2.24: Experimental results for the mouse pilot experiments using mixed WT-OR tumours. The top row shows the results of the AT protocol (left), MTD protocol (middle) and vehicle control (right) for all 5 replicates. The bottom row shows the same results but on a log scale to better see small changes in growth behaviour.




FIGURE 2.25: Results of Flow Cytometry analysis results for the mixed WT-OR tumours.




FIGURE 2.26: Experimental results for the mouse pilot experiments using tumours comprised entirely of OR cells. The top row shows the results of the AT protocol (left), MTD protocol (middle) and vehicle control (right) for all 5 replicates. The bottom row shows the same results but on a log scale to better see small changes in growth behaviour.

that it eradicates this WT cell population. Typically, for an AT treatment protocol, we would expect a gradual increase in OR cell frequency as treatment progresses, but in figure 2.25, we see that the OR cell frequency is slightly lower at the end of the experiment than at the beginning.

Results of the MTD protocol are as we might have expected. Over the first 10 days of the experiment, tumour sizes decrease, during which time we reasonably assume that the WT cells are being eradicated due to the high-dose treatment. After this point, free of any selective pressure imposed by their WT competitors, OR cells proliferate quickly, leading to tumour growth similar to that seen in the vehicle. Flow cytometry results in figure 2.25 show that though WT cells have not been eradicated completely, they are vastly outnumbered by OR cells.

The results of the experiments using resistant tumours are shown in figure 2.26. Despite what we might anticipate, the AT protocol actually outperforms the MTD protocol here, despite the fact that the resistant tumour should not respond to AT treatment. Although the response to the low-dose treatment is not as good as in the case of the mixed tumour, we do see some stabilisation of tumour size in most replicates, and a slower growth rate in the latter half of the experiment when compared to the growth curves in the vehicle.

2.4.8 ODE Models that Exclude Competition Fail to Capture the Qualitative Behaviour Seen in The Data

We have two working hypotheses to explain this unusual growth behaviour in the resistant tumours. The first is that there is some heterogeneity in drug sensitivity in the OR cell line, meaning that some cells within this cell line resist the drug better than others. If the number of OR cells which are partially sensitive to treatment is abundant enough, we could see some containment of the more resistant subclones in the AT protocol. This would also explain the initial decline in tumour size in the MTD case. The cells which respond to treatment would be killed due to the high drug pressure, at which point the remaining cells which better resist treatment can grow freely without any spatial pressure imposed by the partially sensitive cells surrounding them.

Another possibility is that an immune response in the mice (which, though T-cell deficient, still possess the capacity for some immune response) could explain this behaviour. For the AT protocol, the lower drug dose will have less of a burden on the mouse's physical health, and by extension, its immune system. This would allow the mouse's immune system to combat the tumour itself, with the added help of the drug, potentially leading to the results we see in the AT case. In the MTD case, the high drug burden might suppress the immune system. This may mean that initially, while immune cells are still present and capable, they could decrease the size of the tumour. As the physical burden of prolonged drug treatment mounts, immune response may weaken, and the tumour will grow unimpeded.

To test these hypotheses, we will design some simple ODE models. We aim to keep these models as straightforward as possible. At present, we are not concerned with choosing biologically accurate parameter values – this is something that we will do later using our hierarchical bayesian framework. Instead, we wish to see if a model can be used to explain or replicate the qualitative behaviour seen in the data. Initially, we try a null model with which we make minimal assumptions. This has a dual purpose – simultaneously, it shows us that there must be more going on in these experiments than we can explain with these minimal assumptions, as well as showing us what sort of modifications we might make to improve the model.

We start with a null model, in which we make only two assumptions. We first assume that tumour cells will grow exponentially in absence of treatment – we see this behaviour clearly in the vehicle data above. We also assume that cells will be killed according to their phenotype's level of resistance and the drug dose being applied at point of treatment (37).

With our two cell lines and these minimal assumptions, our null model is as follows:

$$\begin{aligned}\frac{dW}{dt} &= aW - \frac{1}{\theta} \cdot \left(\frac{\delta}{IC50_W} \right) W, \\ \frac{dR}{dt} &= bR - \frac{1}{\theta} \cdot \left(\frac{\delta}{IC50_R} \right) R,\end{aligned}$$

where a, b are the growth rates of WT cells W and resistant cells R , respectively, θ is the interval between treatment applications, δ is the given drug dose and $IC50_W$ and $IC50_R$ are the half-maximal inhibitory concentrations of OSM for W and R . The parameter values used




FIGURE 2.27: **(Top 2 rows)** Model Results for each data set with the parameters outlined in figure 2.4. **(Bottom 2 rows)** The same data, zoomed in on the first 20 days, where most, if not all, non-trivial occurrences happened in the data, to see clearly how the model and data compare at these time points.

for hypothesis testing (which, we reiterate, are not strongly data-informed), can be found in figure 2.4 in the methods.

When comparing the growth curves this model produces, as seen in figure 2.27, we see several qualitative behavioural disparities in comparison to the data. The model performs satisfactorily for the mixed vehicle, mixed MTD and resistant vehicle cases. It does not capture the behaviour seen in the remaining cases, however. In the mixed AT case, we see no inhibition on the OR cells – the data shows that these are contained by the sensitive wild type cells present, but here, without any term accounting for containment, we see that the OR

cells grow uninhibited and escape competition. In the resistant AT data, we see an initial dip in population around day 10, and a much slower growth rate, which we cannot achieve with this model, since the OR cells are virtually unaffected by treatment. We see the same dip in population size in the resistant MTD data, which this model also fails to capture.

In short, this null model is insufficient. We conclude based on the Mixed AT modelling that we require some inhibition on OR Cell growth by the WT cells; by the current resistant AT model we conclude that the OR cells do not resist the drug as much as initially assumed, given that this population is temporarily inhibited by treatment more than we would anticipate. We did see some evidence of this happening in the monolayer culture data, but for this null model we wished only to include the bare minimum assumptions, and so left this out for the moment. The resistant MTD model shows that both of these additional factors would be insufficient. With no WT cells to inhibit the OR cells, we would not see any dip in the cell density in the first ten days. If we assume a less resistant OR cell population, we would only see a slower or halted growth curve, but no initial drop in the population. We hypothesise, then, that we may need (at least) a third, more resistant OR cell subclone.

We could also have some immune response, but this would need to involve some sort of drug-dependence since we don't see anything unusual in the vehicle which might suggest any immune interactions. Containment at lower drug doses (such as in the mixed AT data), however, could be induced by some symbiosis between the drug and the immune system. It could also be the case that immune response in these cases is simply not strong enough to induce any significant behaviour, either due to a lack of pressure on the tumour cells or a sufficiently high cell density to override any effect the immune response would have. If we have immune response and lower AT doses working together to contain cells, we would require any significant immune response to occur after application of drug, but also to taper off as drug dose increases due to drug toxicity. We shall test these hypotheses separately, but we note that they are not mutually exclusive and a combination of several could be present.

2.4.9 Three-Phenotype Model of Varied Drug Resistance Better Aligns with Qualitative Behaviour Seen in In Vivo Data

We make two modifications in this next model. We add to the resistant cell lines a linear relative fitness function, which suppresses growth of cells linearly based on their prevalence in the tumour. In (37), the authors point out that this modelling choice will likely lead to predictions which are overly optimistic. At this stage, however, we are only testing hypotheses, and are not yet at the point where we are fitting these models to data. Instead, we only wish to see whether the behaviour of these models matches that of the data from a qualitative perspective. As such, to keep the model as simple as possible, we opt for a relative fitness function which suppresses cell growth linearly based on their prevalence in the tumour. We adjust the IC_{50} of the resistant cell line to be much lower than in the null model, to reflect the surprising growth behaviour seen in the resistant tumours with both the AT and the MTD doses – as a starting point, we match the IC_{50} to the MTD dose. We also add a second resistant population, which does not respond at all to treatment, and makes up some proportion of the OR cell density in the previous model which is largely chosen arbitrarily, since we have no concrete way of establishing what proportion of the OR cells fully resist the drug from the data. With these modifications, our new model is as follows:

$$\begin{aligned}\frac{dW}{dt} &= aW \cdot \frac{W}{W + R + F} - \frac{1}{\theta} \cdot \left(\frac{\delta}{IC_{50_W}}\right)W, \\ \frac{dR}{dt} &= bR \cdot \frac{R}{W + R + F} - \frac{1}{\theta} \cdot \left(\frac{\delta}{IC_{50_R}}\right)R, \\ \frac{dF}{dt} &= cF \cdot \frac{F}{W + R + F},\end{aligned}$$

where $a, b, \delta, \theta, IC_{50_W}$ and IC_{50_R} are as above and c is the growth rate of our fully resistant cells, F . Parameter values used in this model can be found in figure 2.5.

Though parameter values need tweaking to see cell number increases more closely resemble the data, this model behaves much more like the data does, which we see in figure 2.28. Unlike the null model, we see better containment of resistant phenotypes in the mixed AT model, with the WT cell number decreasing very gradually, which is to be expected in reality – indefinite containment of resistant cells is an unrealistic objective.




FIGURE 2.28: **(Top)** New Model Results for each data set with the parameters outlined in figure 2.5. **(Bottom)** The same models, zoomed in on the first 20 days, where most, if not all, non-trivial occurrences happened in the data, to see clearly how the model and data compare at these time points.

In the mixed MTD protocol, though cells do not grow as much as they do in the data, we still see a competitive release of resistant cells. This disparity is important, however – we do not want our model to be overly optimistic, and this is, at present, a very optimistic estimate. There are three things in the model which could be to blame. It could be that the competition term's influence is stronger in the model than it would be in reality – we can adjust this with some additional parameter to decrease this influence, or a different competition term altogether (we have already discussed that this competition term was likely overly optimistic above, and a sigmoidal rather than linear competition term would likely be a better choice). The initial ratio of R to F could be incorrect, but this is difficult to address without more data. The parameters used for hypothesis testing could deviate too far from the as yet unknown true model parameters which would allow us to recapture the shapes of the curves in the data. Inferring growth rates and IC_{50} values will not be an issue. However, even if we assume the growth rate of F is the same as the other two phenotypes, the initial ratio of R to F is more difficult to pin down. If we know all the other parameters, however, this may be doable by using the points at which the OR cell populations start rising after the dip at around day 10 in the data (but even then, we would need to assume extinction of the partially sensitive OR cells, which may be unrealistic).

The resistant AT model is better, since cell growth is much slower and a stable cell number is achieved before cell growth increases in the late stages of growth, but we still do not see a dip in cell number at day 10 in the experiment. Provided the hypothesis that some OR cells resist treatment more than others is correct, this behaviour could likely be achieved by adjusting the initial R to F ratios and IC_{50R} , the latter of the two being the more likely solution. Our data would suggest that the partially sensitive cells could potentially be as sensitive as the WT cells, since the growth behaviour in the data for both the mixed and resistant tumours in the early stages of the experiment are so similar (and, again, if this hypothesis holds).

The extra cell line added in this model gives us something which much better resembles the data for the resistant MTD case. Cell number initially drops as the partially sensitive cells are killed off, and then the fully resistant subclone is allowed to grow unfettered by competition with the partially resistant subclone and the drug, growing exponentially as it

did in the data.

2.4.10 Immune Response Model Fails to Capture Qualitative Behaviour Seen in In Vivo Data

Our first attempt at modelling immune response tries to make minimal assumptions about how immune cells and tumour cells interact. For now, we do not assume that the resistant phenotype has heterogeneous drug resistance levels and assume instead that the drug affects all resistant cells the same way. We assume that cancer cell growth occurs the same way as we did previously – cells grow exponentially but their growth rate is also proportional to their relative frequency, meaning that resistant cells can be contained if they are vastly outnumbered by sensitive ones, and vice versa.

For immune cells, we assume that immune cells grow exponentially and that immune response is weakened as drug dose increases. We further assume that cancer cells are killed at a rate g via interaction with immune cells, and that immune cells are used up after interaction with cancer cells at a rate of d . This yields the following model:

$$\begin{aligned}\frac{dW}{dt} &= aW \cdot \frac{W}{W+R} - \frac{1}{\theta} \cdot \left(\frac{\delta}{IC50_W}\right)W - gWI, \\ \frac{dR}{dt} &= bR \cdot \frac{R}{W+R} - \frac{1}{\theta} \cdot \left(\frac{\delta}{IC50_R}\right)R - gRI, \\ \frac{dI}{dt} &= cI - \frac{1}{\theta} \cdot \left(\frac{\delta}{IC50_I}\right)I - d(W+R)I,\end{aligned}$$

where I is immune cell population, c is the growth rate of these cells, d and g are interaction parameters mentioned previously, and $IC50_I$ is the half-maximal inhibitory dose of I . Parameter values used with this model can be found below in figure 2.6. We note now that parameterising levels of competition between immune and cancer cells from the data will be particularly difficult since we have no way of estimating how many immune cells will be available to combat cancer cells locally, nor can we reliably estimate the impact these immune cells are having beyond looking at the points at which cancer cell populations initially drop and then rise, as we might have done in the case of the heterogeneous drug resistance in OR cells. In this case however, we needed only to estimate a cell number, whereas here we would be estimating three different parameters. Additionally, the values of these three

parameters (and the differences between them) will likely have a huge impact on model outcomes. As such, the parameter values here are guesses, and save for finding other relevant parameters published somewhere else, without more data pertaining to immune response, it will be difficult to improve upon them. At this stage though, we only wish to observe model dynamics and compare those dynamics to the overall behaviour of our data, so for our current purposes, these values will suffice.

The long term behaviour in the vehicle cases for this model, presented in figure 2.29, matches the data qualitatively, but we see an initial dip in cell population for the vehicles which is not present in the data. For the AT cases, largely, this model is adequate – we see containment of OR cells in the mixed case without any substantial cell number increase. For the resistant AT model, we do indeed see an initial dip in cell number due to the immune response. We find that the same can be said of mixed MTD model – WT cells see heavy impact from the drug and this allows OR cells to eventually escape containment and grow unimpeded. The resistant MTD model, too, is sufficient, since the cells grow unimpeded after an initial dip in cell number.

There are two key issues with this model. The first is a matter of parameter choice, which we have touched on already – immune cell population in each case takes quickly approaches zero. This is a result of poor parameter value choice, but it is difficult to choose better ones at this stage, as we have previously discussed. For example, if we were to increase the growth rate of our immune cells, with the current model set up, we may end up driving all of the cancer cells, to extinction, since we see that the cell number declines in every case so long as there is a sufficient number of immune cells. The other issue is that we see an immune response in the model for the vehicle cases that is not present in the data.




FIGURE 2.29: **(Top 2 rows)** First set of immune model results for each data set with the parameters outlined in figure 2.6. **(Bottom 2 rows)** The same models, zoomed in on the first 20 days, where most, if not all, unusual occurrences happened in the data, to see clearly how the model and data compare at these time points.

2.5 Discussion

We saw in figure 2.7 that we could maintain a roughly constant OR cell population at 10nM of OSM, which means that even if resistant cell population is no longer rare, containment may still be possible *in vivo* due to the partial OSM sensitivity that our OR cells exhibit. This is not especially useful if we want to consider adaptive therapy, of course, since we require resistant phenotypes to be rare for it to work. We do, however, now have some understanding as to the mechanism of resistance of the OR cells. We saw that the OR cells grow most successfully in the presence of a low drug dose. This is itself an interesting find, and one that may be important during a low-dose adaptive protocol. A potential risk arises – treating with a low dose may hamper WT cell growth, but may also promote OR cell growth, which we may need to keep in mind when considering an adaptive protocol. We also note that only at 10nM do we see a roughly constant cell population, which is what we seek if we aim to contain the population.

Our analysis of the data in figure 2.7 showed that OR cells have a higher growth rate than WT cells *in vitro*. This is not necessarily an issue – it was shown in (24) that we do not need a fitness cost of resistance in order to see better clinical returns from containment than we might see from MTD treatment. This would suggest, however, that an adaptive treatment protocol would only be effective if resistant cells were rare.

The results of our FACS frequency assays may suggest that we see an Allee effect in the OR cells, where they initially have a low growth rate but at a point some mechanism allows them to grow much more quickly – perhaps a certain number or proportion of cells need to grow before rapid proliferation can take place. Without density data, however, it is difficult to make a concrete hypothesis on this, since we do not know how quickly (if at all) WT cells are growing in these situations.

We also observed in the 5nM and 10nM frequency assays that even though OR cells had a higher frequency by the end, there was a consistent increase in WT cell frequency towards the end of the experiment. We currently have no explanation, nor any working hypothesis, for why this might be. Our exchanged media assays demonstrated that both cell phenotypes grow more successfully in a medium inhabited by OR cells when a drug dose is applied. It is possible that when there is a sufficient number of OR cells, they can produce some factor

that aids in cell proliferation. It is possible that this could be an influencing factor, since the growth media are replaced every 2-3 days, but the exchange media results do not suggest that the extent to which this might influence growth is substantially greater in WT cells than it is in OR cells.

The findings from the exchange media assays warrant further investigation. There could be a number of reasons that OR cells grow faster than WT cells in WT cell medium. It could be that OR cells have to consume more substrate in order to proliferate than the WT. The OR cells may acidify their medium more rapidly than the WT cells, which in general will hinder growth. There is even a possibility that WT cells produce some factor which aids in proliferation, but this seems less likely given that WT growth in absence of treatment was functionally identical in both WT and OR media.

The fact that OR and WT cells grow faster in the OR medium when drug dose is supplied supports the hypothesis that the OR cells excrete a protective factor or some other paracrine factor which reduces drug efficacy. This is further hinted at when we note that OR cells grow very poorly in WT medium when given treatment.

If the OR cells are capable of protecting all of the cells around them from treatment, then in co-culture, they may produce enough of this factor to shield WT cells from treatment as well as themselves. This could mean that we see greater growth rates for low-dose treatment in co-culture assays in which the OR cell population is high. Contradictory to the results seen in section 2.4.1, we have that the growth rate of OR cells in monolayer culture during the exchange media assays is less than that of the WT cells. Our Bayesian hierarchical modelling framework can capture heterogeneity between replicates, so this does not pose an issue, but it would be useful to have a larger number of replicates for the initial growth assays so that we could say with confidence whether or not we observe a fitness cost of resistance in the OR cell line. The results of the exchange media assays suggest that with a large enough number of OR cells, we may see a paracrine signalling mechanism which protects cells or promotes cell growth during low-dose treatment. If so, this will have to be worked into the model in some way, but whether this effect is density- or frequency-dependent is difficult to say without further experimental data.

Our agent-based model did qualitatively capture the behaviour we saw in the spheroid

image data, but detail on the cell cycle behaviour of specific cells was not modelled. Although it was not our objective, it would be worthwhile to model other aspects of the tumour microenvironment. Including every single aspect of would make the model difficult to interpret and draw conclusions from, but creating a model which captures more of the biological mechanisms involved in the tumour microenvironment might make it easier to compare agent-based modelling results with *in-vivo* experimental data. For example, we could add additional agents which represent tumour-associated macrophages, which might influence cell migration. We could also include components which model angiogenesis, which could help us to more accurately describe how local hypoxic conditions influence tumour development.

Given the time and resources, we would be able to model cell cycle phases in real time in a way which was adequately supported by experimental data – such a study was carried out in (233). This would also be a nice motivation for designing a similar model to the one above using PhysiCell (131), which permits off-lattice 3D modelling of cells in different cell cycle phases. Since off-lattice models necessarily require the modeller to properly implement physically realistic cell movement to prevent cell overlap, a model such as this which also incorporates cell cycle phases would be a useful extension of our work. We might include additional cell types beside the key cancer phenotypes, too – for example, we investigated the possibility that some immune response might be seen in the *in vivo* data, and it might be possible to model this using an ABM.

Initially, we designed our agent-based model such that any cell with a relative oxygen concentration less than 0.001 died instantly, but a discussion with our collaborators at IGMM revealed that it was highly unlikely that this was the case. As such, we added the potential for cells in area of low relative oxygen concentration to quiesce rather than dying outright. Functionally, this adds very little in the way of new behaviour, but gives a more realistic look at what is happening in the spheroid at a given time point.

Whether OR cells are surrounded by WT cells and eradicated in the 5nM model simulations is highly dependent on the initial position of the OR cells in the simulation. If they occupy the centre of the spheroid, they cannot proliferate. The closer to the spheroid periphery they are initially, the higher the chance of a higher OR cell frequency at the end of the simulation. As such, it may be prudent to add into our model a component which logs the

position of the OR cells at the beginning of the simulation to see how strong a correlation there is between initial OR cell distance to spheroid periphery and final OR cell frequency. We intend to investigate this in future work – intuitively, if fewer WT cells surround the OR cells, then they will have less trouble proliferating since there will be ample space for them to do so. We note again that we are yet to carry out any qualitative comparison between the results of the model and the 3D *in vitro* experimental data. To do this, we would add to our model a component which records the cell number and spheroid area after each run of the model, and run the model multiple times to get a large amount of data for comparison, which we consider for future works.

We observed unusual behaviour in the resistant AT and MTD *in vivo* experiments. This, combined with a lack of any visible oddities in the vehicle data could hint at some interaction between the immune system and the drug. Given that the mice used in the trial are immunodeficient, it could be that without the drug, the immune response is not strong enough to have any significant effect on the tumour, but when aided by the drug targeting the tumour, we have some synergistic behaviour that allows the immune cells to attack the cancer more effectively once it has been weakened by the drug. Making any firm conclusions on this would require an analysis into the immune response in this specific strain of mice, both with and without the presence of osimertinib.

Though we have *in vitro* experimental data which suggested partial sensitivity to the drug in the OR cell line, this was only observed at a high drug dose – in fact, OR cell growth was actually better in the presence of a low drug concentration in monolayer culture, likely due to the way the OR cells were developed from the WT cells. As such, to see control of the OR cell population during low-dose therapy in the *in vivo* data, despite the fact that they should see little to no hindrance from treatment at such doses, is surprising.

The MTD experiments may give us some insight as to why this is happening. Despite the fact that no sensitive WT cells are present, in the first 10 days of the experiment, the growth behaviour as a response to treatment is very similar to the case of the mixed tumours. In fact, two replicates, the initial response to treatment was even more drastic than that of the mixed tumours. After this point, once again, tumours grow at a rate similar to the vehicle, but it is this initial dip in tumour size that should be our main focus.

The hypotheses above are offered in good faith under a working assumption that the experiments were carried out correctly. However, the error made with gathering this data does adequately explain a number of the unusual observations. A dose of Osimertinib must be prepared with DMSO immediately before application – with time, the solution will separate. It is our understanding now that for the experiments above, the solution was prepared in large batches in advance of the experiments. As such, it is impossible to know what doses of osimertinib are actually being supplied at every point in time. Furthermore, it is possible that the cell lines were compromised, meaning that the resistant phenotypes may not have actually been resistant for some of the experiments. This could, for example, explain why we see some of the bizarre results in the *in vivo* experiments. In figure 2.26, we see a tumour which is fully resistant to treatment stabilising in size whilst being treated with a drug it should not respond to. We explored a number of hypotheses that might explain this, but the most likely explanation is that the cells we are treating are not, in fact, resistant to the drug at all – the resistant cell line being compromised would explain this. Likewise, in the MTD protocol in the same figure, several replicates see a large dip in tumour volume at the fourth time point. This is likely due to poor mixing of DMSO and osimertinib – the mice may have received a dose of drug which was much higher than intended at the third time point. If the cells which comprise this tumour are not fully (or even partially) drug resistant, then that too would explain this dip.

Since our data was compromised, any findings presented in this chapter which pertain to the mechanisms of resistance, the interactions between cell lines, any frequency-dependent effects observed, and the effects of both erlotinib and osimertinib on these cell lines, are purely hypothetical, because these findings largely arose from the data itself or analysis of it. Any treatment strategies attempted or proposed are likewise hypothetical, since they would necessarily depend on the behaviours of these cell lines *in vitro* and *in vivo*. The methodology presented in this chapter which we use to analyse data and parameterise models is still robust, and is used in the following chapter. HAL model simulation results are likewise still valid – a majority of the parameter values used in these simulations were chosen ad-hoc, independent of our experimental results. We would expect this model to be able to capture the qualitative behaviour of a similar system, likely even more successfully than it did above if given the

2910 data needed to accurately parameterise the model.

Chapter 3

Using Mathematical and Bayesian models to Predict Adaptive Therapy Outcomes

3.1 Introduction

In the literature review and the introduction to the previous chapter, we reviewed the rationale behind cancer adaptive therapy, a treatment in which we aim to maintain a constant tumour burden for as long as possible to prevent the emergence of treatment-resistant cell phenotypes. We also discuss the role that mathematical modelling can play in explaining phenomena which might not be immediately clear from experimental results, as well as its potential for predicting and optimising treatment.

In this chapter, we examine new experimental data provided by our collaborators Alexandre Deloupy, Dan Fisher and Liliana Krasinska (IGMM, CNRS, Montpellier). Looking at both *in vitro* monolayer culture data and *in vivo* murine data, we conduct a pre-clinical trial of adaptive therapy. Using our Bayesian hierarchical modelling workflow, we infer key growth parameters from our *in vivo* data, and discuss the differences in experimental outcomes between these experiments and the monolayer culture assays. We then use these results to examine whether an adaptive protocol would be a feasible treatment option, and predict the outcomes of a number of different adaptive protocols with an ODE model which has been calibrated to our *in vivo* experimental outcomes. Model predictions can then be compared to further *in vivo* experimental results, and used to inform clinical trials in future.

3.2 Methods

3.2.1 Monolayer Culture Growth Assays

To determine PC-9 WT and PC-9 Osimertinib Resistant (OR) growth rates in isolation, ten thousand cells were seeded in a well and allowed to grow to confluence over 10 days. Each day, the medium was replenished and the cell number was recorded. WT cells were expressed with green fluorescent protein (GFP) and OR cells were expressed with m-cherry red (a red fluorescent protein) to allow easy cell counting. Three replicates per cell line were prepared.

3.2.2 Co-culture Frequency Assays

A number of assays were prepared to determine how cell frequencies *in vitro* change both over time and in response to varying osimertinib concentration. Four different initial seeding ratios and four drug concentrations were examined. Initial seeding ratios of WT:OR cells were 99:1, 9:1, 1:1 and 1:9; drug doses used were 0nM, 0.5nM, 1nM and 2nM of osimertinib. Every 2-3 days, a proportion of the cells were replated with fresh medium, whilst the rest were used to determine the changing frequencies via FACS analysis (which we cover in the previous chapter). There were two replicates per experiment.

3.2.3 Calculating IC50 Values of Cell Lines

To quantify how cell viability changes with drug response, cells were grown in monolayer culture with an increasing dose of osimertinib. The initial dose started at $10^{-4}\mu\text{M}$ and was increased up to $10^1\mu\text{M}$. We then applied a best fit line using the function

$$\frac{1}{1 + \delta/\text{IC50}}$$

with drug dose δ to ascertain the IC50 values for the cell phenotypes.

3.2.4 In-vivo Mouse Experiments

A total of 15 six-week-old female nude mice were used in the mouse experiments, with 5 replicates per cell line. At the start of the experiment, mice were injected with 2 million PC-9 WT, PC-9 Erlotinib resistant or PC-9 Osimertinib resistant cells, suspended in a phosphate

PC9 Sensitive and PC9 Resistant cell lines in vivo:

2,000,000 Cells/mouse


We inject 100 μL : 50 μL (Cells + PBS) + 50 μL PBS

FIGURE 3.1: An overview of the experimental setup used in the *in vivo* mouse assays, provided by collaborator Alexandre Deloupy (IGMM, CNRS, Montpellier).

buffered saline solution to maintain a stable pH. Once the cells reached a measurable size, the tumour volume was recorded every 2-3 days for a total of 42 days. An overview of these methods can be found in figure 3.1.

Replicate 3 was omitted from the PC-9 WT cohort. Notes in the data showed that the tumour on this mouse reached a burst point abnormally early (around day 10 in the experiment). These issues persisted for the duration of the assay, and we opt to omit this replicate. Replicate 1 was omitted from the PC-9 Erlotinib resistant cohort because the tumour was only a detectable size at three time points – no tumour was detected on this mouse for any of the other time points during the experiment.

We include the data and model fits for the PC-9 Erlotinib resistant mouse cohort in the appendix B, since a majority of our experiments focus on the behaviours of, and interactions between, the PC-9 WT and PC-9 OR cell lines.

3.2.5 Calculation of Selection Coefficients

The method used to calculate selection coefficients is identical to the method presented in the previous chapter – we refer the reader to section 2.3.2 for the full explanation.




FIGURE 3.2: Results of the monolayer isolated growth experiments. Differently coloured lines on each plots each represent one replicate of the experiments.

3.2.6 Bayesian Hierarchical Model

Our Bayesian hierarchical model is identical to the model presented in the methods section of the previous chapter, with a few minor differences. We fit our model to the log-transformed data in this chapter, but this does not affect the design of the model, but rather the data we feed into it. Additionally, we use wide, weakly informative priors to infer parameters, rather than narrower ones, to allow the sampler to more effectively explore the sample space and to prevent our priors from biasing our inferred posteriors too much. We also set our carrying capacity prior to be smaller than we anticipated the true posterior to be, and rescaled the value during the inference by a factor of 100. This decision was made due to the way *Stan* works, and is a method of improving sampler efficiency and model geometry. All other aspects of the method used are the same as seen previously.

3.3 Results

3.3.1 Monolayer Isolation Growth Assays

We see in figure 3.2 the results of the monolayer growth assays, in which each cell line was allowed to grow to confluence in isolation. Both plots show that after an initial period of approximately exponential growth, cells reach confluence at around day 5.




FIGURE 3.3: Experimental results showing the changing cell viability of WT and OR phenotypes as drug dose is increased. IC50s of Osimertinib were approximated as 4.62nM for the WT cells and 1885.86nM for the OR Cells.

Given the low number of replicates in these assays, we opted to fit a logistic model to these data using a nonlinear least squares method rather than using our hierarchical Bayesian workflow. This model has the form

$$\frac{dX}{dt} = rX\left(1 - \frac{X}{K}\right),$$

where $X = S, R$, and r and K are the growth rate and carrying capacity, respectively. The result of this fitting showed that the WT cells had a daily growth rate of 1.96/day and a carrying capacity of 1.85×10^6 cells. The OR cells had a daily growth rate of 1.39/day with a carrying capacity of 1.28×10^6 cells.

3.3.2 Drug Response Analysis Shows that OR cells are Highly Resistant to Treatment with Osimertinib

Figure 3.3 shows the results of the drug response analysis assays, showing how cell viability varied in response to an increasing dose of osimertinib. The IC50 for WT cells was found to be 4.62nM and the IC50 for OR cells was found to be 1885.86nM.




FIGURE 3.4: Results of the drug-free co-culture frequency assays with varying initial ratios. First and third rows depict the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. Second and fourth rows compare the corresponding predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line). Only the results of the first replicates are shown, since there are no major qualitative differences between both replicates. The results of the second replicates are shown in appendix B.




FIGURE 3.5: Results of the co-culture frequency assays with a drug dose of 0.5nM. First and third rows depict the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. Second and fourth rows compare the corresponding predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line).




FIGURE 3.6: Results of the co-culture frequency assays with a drug dose of 1nM. First and third rows depict the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. Second and fourth rows compare the corresponding predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line).




FIGURE 3.7: Results of the co-culture frequency assays with a drug dose of 2nM. First and third rows depict the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. Second and fourth rows compare the corresponding predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line).

3.3.3 Co-culture Frequency Assays

We examine the results of the co-culture frequency assays in two ways to determine whether or not there are any frequency-dependent effects that are not immediately clear from analysing the data. The first method compares the differences in a selection coefficient we predict directly from the homogeneous growth rates of our two cell lines and a mean calculated selection coefficient which we calculate directly from the frequency data itself. This selection coefficient tells us how much of an advantage OR cells have at a given time point. If the selection coefficient is positive, then OR cells have an advantage; if it is negative, they are at a disadvantage.

The results of these analyses can be seen in figures 3.4, 3.5, 3.6 and 3.7. Calculated selection coefficients for the drug-free, 0.5nM and 1nM experiments show that OR cells are disadvantaged, and WT cells are selected for – this is even true of the case where WT cells have an initial seeding ratio of only 10%. For the drug-free and 0.5nM doses, we see that even when OR cells outnumber WT cells initially, they manage to overtake the OR cells in frequency by the end of the assay. Conversely, when the drug dose reaches 2nM, we see that OR cell frequency exceeds WT cell frequency when initial ratios are approximately equal, and when OR cells outnumber WT cells initially, WT cell frequency stays roughly constant under this increased drug burden. In the drug free, 0.5nM and 1nM experiments, we see that the mean calculated selection coefficient is negative, and so we predict that WT cells would still have a selective advantage in these cases. In the 2nM assays, however, the drug burden is large enough that OR cells have a slight selective advantage over WT cells.

3.3.4 Bayesian Hierarchical Modelling of In Vivo Data Yields Posterior Distributions of Parameter Values

In figures 3.8 and 3.9, we can see the results of the initial *in vivo* mouse experiments. We do not see any major qualitative behaviour of note that has not already been seen in the *in vitro* data here, but we note that after a period of growth, we see tumours reach a roughly stable size – these sizes differ between WT and OR cohorts. We present the data for the ER mouse experiments, and the resulting model fits, in appendix B.




FIGURE 3.8: Results of the *in vivo* mice experiments for the WT cohort. All 5 replicates are shown, with the top row showing the data on a linear scale, and the bottom row showing the data on a log scale.




FIGURE 3.9: Results of the *in vivo* mice experiments for the OR cohort. All 5 replicates are shown, with the top row showing the data on a linear scale, and the bottom row showing the data on a log scale.

To analyse these results, then, we implement a Bayesian Hierarchical model, which yields posterior distributions of parameter values. Given that tumours are all only comprised of one cell phenotype, and no treatment is applied, we can use a standard Gompertz growth model to infer tumour growth rates and carrying capacities. We use the following model for both cell phenotypes, with $X = S, R$:

$$\frac{dX}{dt} = \alpha X \log\left(\frac{K}{X}\right),$$

with baseline growth rate α and carrying capacity K . We set the prior distribution for our hyperparameter μ_α to be $N(0, 5)$, and the prior for K to be $N(10, 5)$, which we rescale in the *Stan* model so that it more closely resembles the carrying capacities seen in the data. This rescaling of a lower prior distribution is a facet of how *Stan* works internally, and helps the sampler explore the parameter space more quickly.

Parameter inference for both *in vivo* datasets was successful. We see the results of these inferences in figures 3.10, 3.11, 3.12 and 3.13. The sampler converged to clear and distinct posterior distributions for each replicate, with fewer than 0.01% of transitions during sampling diverging from the true trajectory (less than 50 per 10,000 for both inferences). $\hat{R}$ values were within acceptable bounds, and the number of effective samples for each parameter inferred was sufficiently high. The qualitative results of our inference tells us that *in vivo*, the growth rate of OR cells is generally higher than that of the WT cells, but also tells us that the OR cells show a lower carrying capacity than the WT cells, which contrasts our *in vitro* findings.

3.3.5 Inference Enables Parameterisation of a Mathematical Model which can be Used to Predict Outcomes

With the posterior distributions obtained from fitting our *in vivo* data to a Gompertz model, we can simulate the effect of different treatments to predict the outcomes of *in vivo* mouse trials which incorporate different treatments. To do this, we use the following mathematical model of two cell lines, seen in (24):

$$\frac{dS}{dt} = \alpha_S \log\left(\frac{K_S}{S + R}\right)S,$$




FIGURE 3.10: WT Model fit for each replicate used in the inference. 100 parameter sets were drawn from the posterior distributions and then used in the initial model (green lines) and compared to the data (black dots).




FIGURE 3.11: Histograms displaying the inferred posterior distributions for both hierarchical and group-level parameters for the WT model fit. Left column shows these distributions against their prior; right column shows them with the prior excluded, for clarity.




FIGURE 3.12: OR Model fit for each replicate used in the inference. 100 parameter sets were drawn from the posterior distributions and then used in the initial model (red lines) and compared to the data (black dots).




FIGURE 3.13: Histograms displaying the inferred posterior distributions for both hierarchical and group-level parameters for the OR model fit. Left column shows these distributions against their prior; right column shows them with the prior excluded, for clarity.

Parameter	Value	Source
α_S	0.066/day	Bayesian Hierarchical Model
α_R	0.082/day	Bayesian Hierarchical Model
K_S	4.86×10^{10} Cells	Bayesian Hierarchical Model
K_R	2.99×10^{10} Cells	Bayesian Hierarchical Model
$IC50_S$	4.62	Dose Response Data
θ	7	Ad hoc
δ_{MTD}	4.62	IC50 of WT Cells
δ_{low}	$\delta_{MTD}/5$	Ad hoc
δ_{AT}	Various	Ad hoc
$N(0) = N_0$	4×10^8	Approx. Number of Cells in a 4mm^3 Tumour
$S(0) = S_0$	$0.99N_0$	Ad Hoc
$R(0) = R_0$	$0.01N_0$	Ad Hoc

FIGURE 3.14: Parameter Values used for mathematical model simulations, and how these parameter values were obtained. We note that $N = S + R$, with $N_0 = S_0 + R_0$.

$$\frac{dR}{dt} = \alpha_R \log\left(\frac{K_R}{S+R}\right)R,$$

where S and R are the densities of our WT and OR cells, respectively, α_S and α_R are the respective growth rates and K_S and K_R are the respective carrying capacities. In this model, a resistance cost would correspond to a lower α_R or K_R – our previous inference results demonstrated that we have the latter but not the former. Figure 3.14 shows the parameter values used for the simulations performed.

A total of 6 different treatments were tested, along with a control simulation in which no treatments were applied. In our simulations, treatments were applied as a bolus every θ days, which we implemented by multiplying our WT cell subpopulation by

$$\frac{1}{1 + \frac{\delta}{\text{IC50}_S}},$$

where δ is the drug dose applied and IC50_S is the half-maximal inhibitory concentration of S cells. Given that the IC50 of our OR cells is approximately 450 times that of our WT cells, we opt to treat OR cells in our model, and simulations, as fully resistant to treatment. The 5 treatment protocols trialled were as follows:

- **Continuous MTD:** MTD dose of δ_{MTD} was applied every θ days.
- **Continuous Low Dose:** A dose of δ_{low} was applied every θ days.
- **Adaptive MTD:** Every θ days, we check if a cell number threshold of $10N_0$ is reached. If it is, an MTD dose is applied.
- **Containment at N_0 :** An initial dose of $0.5\delta_{\text{MTD}}$ is used. If cell number goes above N_0 , dose is increased by 20%; otherwise, it is decreased by 20%. Dose is not allowed to exceed δ_{MTD} .
- **Containment at $10N_0$:** As previous, but with dose alterations at a different threshold.

We trialled each treatment for 1000 days, and assume that OR cells make up only 1% of our tumour initially.

3.3.6 Mathematical Modelling Results Demonstrate that Adaptive Protocols can Delay the Onset of Resistance

The results of our simulation can be found in figure 3.15. In these simulations, we record the point at which the resistant cell population reaches a frequency of 25%. This point is somewhat arbitrary, but we would not consider resistant cells rare at this point, so it serves as a useful marker to give us some idea of how quickly resistance is emerging in the tumour population. The points at which these thresholds are reached, and the cumulative drug dose supplied for each treatment protocol, can be found in figure 3.16.

As we would predict, treatment with a continuous MTD dose sees the WT cell population suppressed, and allows OR cells to become abundant – they pass the 25% frequency threshold at day 87 in the simulation. A continuous low-dose protocol sees an immediate improvement in containing resistance, with the threshold being reached much later into the time period – the adaptive MTD protocol sees this threshold reached at a similar time point.

In the AT protocol where we attempt to contain at the initial size N_0 , we see that the resistant frequency threshold is reached almost as quickly in as it is in the continuous MTD protocol. In contrast, however, we see that in the AT protocol where we contain at $10N_0$, resistance emerges the slowest out of all of the protocols tested, with the 25% threshold being reached at day 633 – the latest of all of the protocols.

The cumulative drug doses for each protocol show that an adaptive containment protocol at the higher threshold requires the least amount of treatment cumulatively, whilst also boasting the highest capacity for delaying onset of resistance.

3.3.7 Mathematical Modelling Results Demonstrate that Adaptive Protocols can Delay the Onset of Resistance

An important consideration when implementing adaptive protocols which exploit competition between phenotypes is how much of a tumour is comprised of resistant cells initially. To investigate this, we repeated the above simulations, but with a higher initial ratio of OR cells. The results of these simulations can be found in figure 3.17 and 3.18.

We see no difference in which treatments have the greatest capacity for delaying onset of resistance here. The major qualitative difference we observe is that resistance emerges much




FIGURE 3.15: Model simulation results of 6 different treatment scenarios, with an initial cell ratio of 99% WT to 1% OR cells.. The plot on the top right shows the total cell number, N , throughout each simulation for each treatment. Remaining plots show dynamics of S (blue), R (red) and N (grey) for 6 different treatment scenarios. A vertical dashed line is shown to represent the point at which OR Cells reach a frequency of 25%.

Treatment Protocol	Time when R Frequency 25%	Cumulative Drug Dose
Continuous MTD	87	4624.62nM
Continuous Low Dose	444	924.92nM
Adaptive MTD	438	1011.18nM
AT: Containment at N_0	102	3988.34nM
AT: Containment at $10N_0$	633	624.48nM

FIGURE 3.16: Table presenting the times at which frequency of R reached 25% and cumulative drug dose for the 99:1 simulations.

earlier than in the 99:1 simulations, even with only a 4% increase in the initial proportion of OR cells. We also note that the adaptive MTD and containment at $10N_0$ protocols incur a much higher cumulative drug dose than they did previously.

3.4 Discussion

We notice in the monolayer isolation growth assays that cells reached confluence at different cell numbers. Given that resources are frequently replenished in the wells, we can assume that the lower carrying capacity of OR cells is most likely evidence that these cells are larger in size than their WT counterparts. This is because the carrying capacity here represents the point at which cells reach confluence – the upper limit on cell number here is just a result of there being no more space for them to proliferate. What we do note, however, is that OR cells have a lower growth rate than the sensitive WT cells, so we can assume, based on these results, that OR cells must pay a fitness cost to develop resistance to Osimertinib, at least *in vitro*.

During the analysis of our frequency data, we found that in the drug free case, the predicted and mean calculated selection coefficients were roughly the same for each initial seedling ratio, which suggests that there are no frequency-dependent effects which would influence selection here. In all cases where a drug dose is applied, however, we see that the predicted selection coefficient is consistently lower than the predicted one, which suggests that the OR cells are conferred less of a benefit *in vitro* than we predict when a drug dose is applied. This might suggest some sort of frequency-dependent effect which benefits WT cells when they outnumber OR cells, but we note that even in the 0.5nM and 1nM cases, OR cells are still at a disadvantage even when they initially have a greater frequency. This makes it difficult to




FIGURE 3.17: Model simulation results of 6 different treatment scenarios, with an initial cell ratio of 95% WT to 15% OR cells.. The plot on the top right shows the total cell number, N , throughout each simulation for each treatment. Remaining plots show dynamics of S (blue), R (red) and N (grey) for 6 different treatment scenarios. A vertical dashed line is shown to represent the point at which OR Cells reach a frequency of 25%.

Treatment Protocol	Time when R Frequency 25%	Cumulative Drug Dose
Continuous MTD	15	4624.62nM
Continuous Low Dose	84	924.92nM
Adaptive MTD	90	3234.00nM
AT: Containment at N_0	18	3988.34nM
AT: Containment at $10N_0$	123	2725.75nM

FIGURE 3.18: Table presenting the times at which frequency of R reached 25% and cumulative drug dose for the 95:5 simulations.

say with any surety whether such a frequency-dependent effect can be seen.

We note that some of the model fits from our Bayesian hierarchical modelling fit the data better than others. This is an unavoidable outcome when working with biological data. One replicate was excluded because tumours appeared to vanish and were impossible to measure; another was excluded because the tumour grew so large that the mouse suffered too many adverse side effects. We opted to only exclude these replicates, since our collaborators cited specific reasons to do so, but any other replicates which exhibited abnormal growth for no observable reason were used in our inference for the sake of transparency. Nevertheless, we demonstrate that we can infer clear and interpretable parameters for our model using our repeatable and reproducible Bayesian workflow, in well under 200 lines of code.

An important finding was that the qualitative results of our *in vitro* and *in vivo* assays differed in terms of quantifying the fitness cost of resistance for OR cells. In our *in vitro* assays, we found that OR cells proliferated slower than their WT counterparts. In our *in vivo* assays, however, we found the opposite – in fact, OR cells proliferated about 25% faster *in vivo* than they did *in vitro*. We also found, however, that OR cells had a carrying capacity which was considerably lower than WT cells *in vivo* (and we observed this, too, *in vitro*, but it is important to note that in this case, this was not a true carrying capacity). Whether this amounts to an overall fitness disadvantage is hard to say exactly. What it does highlight, however, is that we may have vastly overestimated the fitness advantage WT cells have over OR cells if we had opted to inform treatment decisions using *in vitro* data alone. This highlights the importance of performing comprehensive preclinical trials, ideally incorporating findings from both *in vitro* and *in vivo* experimental data to inform treatment decisions at a clinical level.

Our selection of treatment protocols was made with clinical viability in mind. We opted

not to trial any protocols which relied on certain resistant frequency thresholds being reached because this would require frequent biopsies of the tumour, which would be both invasive for patients and costly in clinical practice. Continuous protocols and protocols which adapt dose in response to tumour size progression only are much easier to implement clinically – although the latter would require scans on a regular basis, this would be less invasive than regular biopsies.

Using a model which was calibrated using *in vivo* mouse data, we presented simulation results which demonstrated that adaptive protocols can vastly outperform MTD protocols with respect to delaying onset to resistance. Containing at a high tumour size threshold and adapting the drug dose according to tumour progression not only delays onset of resistance longer than any of the other protocols trialled, but also requires the lowest cumulative drug dose. We must note, however, that doing so required us to maintain the tumour at a size which was close to its inferred carrying capacity. Allowing a tumour to grow to a near maximal size before attempting to treat it, and maintaining this size for lengthy periods, would be clinically infeasible. We also found, however, that trying to contain the tumour at too small of a size saw only minor gains over an MTD protocol. Striking a middle ground between these two extremes will be important in clinical applications.

Our simulation results demonstrate that the efficacy of adaptive protocols hinges on resistant phenotypes being rare within the tumour. Treatments trialled at the lower initial frequency of OR cells saw much better returns than the higher initial frequency of OR cells, both in terms of delaying onset of resistance and cumulative dose.

Although our simulation results are obtained using a model which is calibrated with mouse data, additional data is needed before our findings are clinic-ready. The fact that the qualitative growth behaviour *in vivo* differs from our *in vitro* findings shows that we require additional *in vivo* data which shows how it is that real tumours respond to treatment. Further data which shows how tumours which contain both WT and OR cells would be invaluable to validate our model findings – they might give us a better understanding of how competition between cells affects tumour dynamics, and we can analyse these results with the methods used previously.

Chapter 4

Mathematical Modelling of Bipolar Androgen Therapy

4.1 Introduction

In chapter 1, we discussed androgen deprivation therapy (ADT) is the current standard of care for advanced and metastatic prostate cancers, and how both continuous and intermittent protocols for androgen deprivation have been proposed. We reviewed the key mechanisms involved in the development of these cancers, and saw that current protocols involving androgen ablation may create an environment which selects for castrate resistant cell phenotypes within the tumour, reducing the effectiveness of ADT and eventually leading to the development of castrate resistant prostate cancer. Such cancers do not respond to ADT, leaving only adjunctive therapies like chemotherapy as the only option of slowing tumour progression.

Bipolar androgen therapy (BAT) was proposed by Denmeade and Isaacs in 2010 (181). The proposed method for BAT involves cycling between near-castrate levels of testosterone seen during ADT and supraphysiologic testosterone levels achieved via an intramuscular testosterone injection. After a period of ADT, an injection of testosterone cypionate or testosterone enanthate is administered, after which ADT is immediately resumed – this cycling between high and low testosterone levels continues until the treatment period ends. We discuss in chapter 1 a range of different clinical trials which utilise BAT protocols. A consistent feature across all of these trials is the scheduling of testosterone doses. Injections of testosterone cypionate/enanthate are consistently given at a 400mg every 28 days (234; 218; 188; 216; 187; 217). The choice of dose is likely owing to the fact that this is

Parameter	Value	Source
r_H	0.043	(219)
r_R	$0.8 * 0.043$	(219)
δ_H	0.027	(219)
δ_R	0.046	(219)
γ	0.098	See Caption.
ν_H	300	Ad-hoc
ν_R	50	Ad-hoc
κ_H	300	Ad-hoc
κ_R	500	Ad-hoc
ϕ	0.031	Ad-hoc
ψ	500	Ad-hoc
ρ_R	50	Ad-hoc
μ_R	500	Ad-hoc
ϵ	500	Ad-hoc
σ	0.0041	Ad-hoc

FIGURE 4.1: Model parameters used in simulations and their sources. We were able to calculate γ directly given that it in BAT protocols seen in the literature, it takes 28 days for testosterone to go from a supraphysiological level to a near-castrate one. All ad-hoc values are shape parameters, giving us the functions seen in figure 4.2.

the highest legal dose of intramuscular testosterone according to the FDA (192), allowing clinicians to quickly reach the high levels of supraphysiological testosterone needed for BAT to be effective.

Several notable mathematical models of ADT and IADT exist, but to our knowledge, no mathematical models of BAT have been published. As such, we present a model with minimal assumptions which aims to describe the BAT treatment protocols seen in the literature. We then simplify this model to make it more analytically tractable, and use our models to propose an alternative and improved approach to treatment scheduling seen in previous clinical trials.

4.2 Methods

4.2.1 Model Parameterisation

Clinical trials for BAT so far have been small and vary in objective from study to study, with few consistencies between them. What does remain consistent for each, however, is the scheduling of testosterone injections. The doses seem to be selected due mainly to FDA

drug restrictions – a 400mg of testosterone is the largest legal dose a patient can be given (235). Injections consistently take place every 28 days in BAT trials – when the injection is given, testosterone levels rise sharply to this desired supraphysiological level and then gradually decline to around twice to four times the castrate testosterone level of 50ng/dL before another injection is given. An objective, then, would be to investigate whether this schedule is optimal. It could well be that more or less frequent injections may yield better results, both in terms of delaying onset of treatment resistance and restoring sensitivity to antiandrogen treatments like enzalutamide or abiraterone.

In order to investigate this, we design a model which investigates how hormone-naïve and castrate resistant cells respond to varying concentrations of testosterone – the parameters used in this model can be found in figure 4.1. Detailed data on how cell growth and death rates vary with androgen concentration is not available, so we aim to try and at least capture the qualitative behaviour: for hormone-naïve cells, growth rate increases and death rate decreases as androgen increases; for castrate-resistant cells, the same happens until testosterone concentrations are high, at which point growth rates decrease and death rates increase. Parameters which describe baseline growth and death rates were chosen based on biologically relevant values, which were stated in a review of other mathematical models(219).

4.3 Results

4.3.1 Designing a model of Bipolar Androgen Therapy

We propose a mathematical model of two cell populations – hormone-naïve (HN), denoted by H , and castrate resistant (CR), denoted by R – with which we investigate the temporal dynamics of BAT on CRPC. A key assumption in the model is that the CR cell population expresses the androgen receptor (AR), a transcription factor in prostate cancer cells which, when bound to androgens, initiates DNA transcription and promotes proliferation. Mechanisms of resistance to ADT are varied, but a key feature in a number of these mechanisms is that androgens are still vital to cell proliferation. Hypersensitivity of the AR, for instance, has been found in 25 – 80% of CRPCs (199; 200; 198). Other mechanisms of resistance, such as downregulation of key co-repressors (209) or upregulation of key co-activators (210)




FIGURE 4.2: Plots demonstrating the shape and behaviour of the functions which govern growth and death rates for the model outlined in section 4.3.1.

in CRPC, aberrant activation of the AR via outlaw pathways (211), or mutations on the AR allowing transcription to occur via progesterone (202) or cortisone (201), all exhibit activity of the AR. For the sake of transparency, we assume that in the case of our model, the mechanism of resistance in our CR cells is upregulation of the AR leading to hypersensitivity, since this particular mechanism is common and so is likely a major cause of CR cell behaviour at high testosterone concentrations.

We model testosterone, whose concentration is denoted by A , as a limiting resource, and assume that testosterone dynamics are independent of both tumour size and the individual HN and CR population sizes. For hormone-naïve cells, we assume that cell growth and death depend linearly on testosterone levels, assuming that growth rate increases and death rate decreases as testosterone levels increase. This ensures that the model is neither too optimistic, nor does it make any assumptions about the effects of testosterone concentration on mutation between HN and CR cells, which we do not include in this model.

We assume that the androgen-dependent growth rate of HN cells, $g_H(A)$, is governed by an increasing Michaelis-Menten function, and that the androgen-dependent death rate, $d_H(A)$, follows a simplified decreasing Hill function. The same functions will be used for CR cell lines when in near-castrate to normal androgen concentration ranges, but when we

enter supraphysiological androgen levels, we switch the dynamics, so that growth rate follows a decreasing hill function and death rate follows an increasing Michaelis-Menten function, which are shifted along the x -axis so that the point at which the dynamics switch is as seamless as possible. This decision is motivated by *in vivo* clinical findings (185; 186).

Since BAT is intended as a treatment which is applied once ADT has failed – at which point it is assumed that resistant has already emerged – we shall assume in our model that there is always some initial population of CR cells. This also ensures that we avoid the trivial case in which we have a tumour comprised only of HN cells, which can simply be treated with ADT alone. Given that ADT is applied at all times bar when testosterone is injected, we assume that testosterone levels decrease at a constant rate at all times, and when testosterone injections are given, we assume an instantaneous increase in testosterone levels. We assume that all parameters, as well as the cell populations and androgen concentration, have positive initial values to ensure biological relevance.

$$\frac{dH}{dt} = g_H(A)H - d_H(A)H,$$

$$\frac{dR}{dt} = g_R(A)R - d_R(A)R,$$

$$\frac{dA}{dt} = -\gamma A,$$

$$A(t_B) = 1500,$$

where androgen concentration, A , is set to a fixed value of 1500 at times t_B when testosterone injections take place, and γ is the decay rate of testosterone as a result of ADT treatment during BAT. The functions which govern the androgen-dependent growth and death of both cell populations are defined as follows:

$$g_H(A) = \frac{r_H}{1 + \frac{\nu_H}{A}},$$

$$d_H(A) = \frac{\delta_H}{1 + \frac{A}{\kappa_H}},$$

$$g_R(A) = \begin{cases} \left(\frac{r_R}{1 + \frac{\nu_R}{A}} \right) + \phi & A \leq 500, \\ \frac{r_R}{1 + \frac{A-\psi}{\kappa_R}} & \text{Otherwise,} \end{cases} \quad (4.1)$$

$$d_R(A) = \begin{cases} \frac{\delta_R}{1 + \frac{A}{\rho_R}} & A \leq 500, \\ \frac{\delta_R}{1 + \frac{\mu_R}{A-\epsilon}} + \sigma & \text{Otherwise,} \end{cases} \quad (4.2)$$

$$d_R(A) = \begin{cases} \frac{\delta_R}{1 + \frac{A}{\rho_R}} & A \leq 500, \\ \frac{\delta_R}{1 + \frac{\mu_R}{A-\epsilon}} + \sigma & \text{Otherwise,} \end{cases} \quad (4.3)$$

$$d_R(A) = \begin{cases} \frac{\delta_R}{1 + \frac{A}{\rho_R}} & A \leq 500, \\ \frac{\delta_R}{1 + \frac{\mu_R}{A-\epsilon}} + \sigma & \text{Otherwise,} \end{cases} \quad (4.4)$$

where r_H and r_R are the respective maximal growth rates of H and R , δ_H and δ_R are the respective maximal death rates of H and R , and $\nu_h, \kappa_H, \nu_r, \kappa_R, \phi, \psi, \rho_R, \mu_R, \epsilon$ and σ are all shape parameters for these functions. We assume that the behaviours in the functions $g_R(A)$ and $d_R(A)$ change at a testosterone concentration of 500ng/dL because this is considered to be a normal physiological level, and as such, for the sake of our model, we treat anything above this concentration as a higher than normal concentration. When simulating our model, we halt the decay of testosterone at 100ng/dL . This is because the literature suggests that total androgen ablation is generally not achievable, since adrenal production of testosterone is not easily prevented. Were total androgen ablation possible, and assuming that both H and R cells require testosterone to proliferate, the optimal treatment strategy, trivially, would be to completely remove testosterone from the system. Plots of functions which govern the influence of varying androgen concentrations over growth and death rates for H and R can be found in figure 4.2.

4.3.2 Initial Model Simulations

An initial round of simulations was performed to demonstrate that the model performs as desired, and these are shown in figure 4.3. We simulate our model undergoing various treatments using the R package *deSolve* (236). In these simulations, initial proportions of 1:1 HN to CR cells were used, but it should be noted that since we can control either the HN or CR cell populations with ADT or high testosterone treatment, respectively, we could feasibly opt for a BAT protocol given any initial tumour composition. We can see that in the case of ADT alone, HN cells are suppressed completely. Our model does assume that CR cells will see some degree of inhibition during ADT, given that we assume that these cells are still




FIGURE 4.3: (Left) Simulated model results with continuous ADT. (Right) Simulated model results of a BAT treatment protocol which follows the same treatment protocol seen in previous literature (234; 218; 188; 216; 187; 217). Both simulations start with a 1:1 ratio of H:R cells, with a total initial population of 100 million cells. Hormone-naïve cell populations are represented by the blue lines, castrate-resistant populations are represented by the red lines, and the total cell numbers are represented by the dashed black lines.

Parameters used are shown in figure 4.1.

testosterone-dependent due to AR hypersensitivity. However, this is much less pronounced than the growth inhibition that we see in HN cells, and so we see CR cells thrive in the case where ADT alone is used. These results serve as a good motivation to try BAT in this scenario – proceeding with ADT alone, given our initial conditions, will select for the CR cell population. Adjunctive therapies would potentially delay progression for some time if this population is selected for, but with substantial adverse effects on quality of life.

The right panel of figure 1 demonstrates that BAT, when compared to ADT alone, has the potential to slow the growth of CR cells via an overloading of the hypersensitive AR with supraphysiological levels of testosterone. We immediately observe here that the H and R cell populations are growing at approximately similar rates whilst receiving testosterone injections. H cells are slightly more prevalent than R cells during this treatment phase, but the frequencies of H and R cells remain roughly constant during this treatment phase. When injections are ceased, and ADT resumed, we see a flat increase in R cell population and a flat decrease in H cell population. This immediately challenges the BAT protocol of 4 injections every 28 days, followed by ADT, that we have seen in the literature. If we were to inject testosterone on this schedule continuously, we would expect the model dynamics seen during the first four injections to continue. Although this would still result in a tumour which steadily increases in size, this tumour would still consist of more H cells than R cells,




FIGURE 4.4: A plot with the treatment response curve of our original model in black, and the line for g'_R drawn in red, representing the line we can travel along on average in an alternating treatment. Testosterone concentration increases from 100ng/dL to 1500ng/dL from left to right, so this plot is functionally a 2D projection of a 3D plot – here the x-axis contains the net growth rate of H cells, and the y-axis contains the effective growth rate for R cells.

and without the adverse QoL impact we see during continuous ADT. We also observe that the final size of the tumour over this fixed time period is similar in both cases. This alone would suggest that BAT is no worse than ADT with respect to final tumour size. We posit, however, that we can improve upon this final tumour size with a better scheduling of high and low testosterone treatment phases. These two treatment phase lengths, and their scheduling, is what we investigate and optimise.

4.3.3 Simplified model with Instantly Alternating Testosterone Concentrations

The model above captures the phenomena of interest sufficiently, but is analytically intractable. In order to gain a better understanding of how we might optimise this alternating treatment, we shall study a simplified model to obtain some concrete analytical results which we can use as a foundation for an improved BAT treatment protocol.

We consider a model in which we can instantaneously switch between the two extremes of testosterone concentration, omitting the testosterone decay from the previous model. With

this simplified model, we can investigate how we might decide how much time in a fixed treatment period would be dedicated to high testosterone treatment. Since we can instantaneously change the conditions between two testosterone doses in this model, we essentially are studying an alternating treatment in which we switch between ADT, where CR cells can grow steadily but HN cell populations decline, and BAT, where CR cell populations decline but HN cells grow rapidly.

To be more specific, we consider two cell populations, H and R , whose growth rates are dependent on the concentration of some limiting resource $a \in A$, where $A \subset \mathbb{R}^+$ is the set of plausible concentrations of this resource (which, in our case, is testosterone). We consider a treatment period of T days. We further define two functions, $g_H : A \rightarrow G_H, g_R : A \rightarrow G_R$, where $G_H, G_R \subset \mathbb{R}$ are the sets of possible resource-dependent growth rates (which do not necessarily have to be the growth functions defined in the model above, so as to keep the result as general as possible). Our model is as follows:

$$\begin{aligned}\frac{dH}{dt} &= g_H(a)H(t), \\ \frac{dR}{dt} &= g_R(a)R(t),\end{aligned}$$

where $H(t), R(t)$ are the respective population sizes of H and R at time $0 \leq t \leq T$, with initial conditions $H(0) = H_0, R(0) = R_0$. We further define $N(t) = H(t) + R(t)$ to be the total population size at time t . We specify that g_H is monotonically increasing in A .

4.3.4 Analytical Approach to Treatment Scheduling in the Instant Switching Model

Consider two treatment doses $\alpha, \beta \in A$, which are selected thusly:

$$\begin{aligned}g_H^\alpha &= g_H(\alpha) = g_H(\min(A)), \\ g_H^\beta &= g_H(\beta) = g_H(\max(A)), \\ g_R^\alpha &= g_R(\alpha) = g_R(\min(A)),\end{aligned}$$

$$g_R^\beta = g_R(\beta) = g_R(\max(A)).$$

α here corresponds to the near-castrate levels of testosterone which we see during ADT, and β corresponds to the supraphysiologic androgen levels we reach immediately following a testosterone injection. We assume that we can switch instantaneously between α and β whenever we wish. Consider a plot where the x -axis contains all values of $g_H(a)$ for all $a \in A$, and the y -axis contains all values of $g_R(a)$ for all $a \in A$, and define the line

$$g'_R = mg'_H + c$$

as the line which connects the two points (g_H^α, g_R^α) and (g_H^β, g_R^β) on this plot. Specifically, we define a function $g'_R : A \rightarrow G'_R$, where G'_R is the set of points on this line defined by $g'_R(a) = mg_H(a) + c$. In figure 4.4, we can see a comparison between such a line defined on our possible testosterone concentrations for the initial model and the varying growth rates given by the functions defined in our initial model. We further outline the following conditions to ensure that our results remain relevant to the problem at hand, with a view to explicitly calculating the optimal period of time, $T_\alpha \leq T$, which we spend undergoing treatment α in a fixed treatment period T .

1. $\forall a \in A$ where $a \notin \{\min(A), \max(A)\}$, $g_R(a) \geq g'_R(a)$ (the curve for $g_R(a)$ must lie above the line $g'_R(a)$ everywhere except the endpoints);
2. the gradient of the line satisfies $m < 0$ (a condition which ensures that our solutions are defined);
3. Neither (g_H^α, g_R^α) nor (g_H^β, g_R^β) are negative (A trivial case – if the growth rate for both cell types is negative at one dose, then this dose is optimal, and we shouldn't alternate at all);
4. $g_H^\alpha \neq g_H^\beta$ and $g_R^\alpha \neq g_R^\beta$ (Another trivial case. If $g_H^\alpha = g_H^\beta$ and $g_R^\alpha = g_R^\beta$, then alternating between the two does nothing. If $g_H^\alpha = g_H^\beta$ but $g_R^\alpha \neq g_R^\beta$ (or vice versa), we select the treatment with the smallest growth rate out of g_R^α and g_R^β and we stay at this dose for the duration of T);

- 3368 5. The per-cell growth rate of H must not be larger than the per-cell growth rate of R un-
 3369 der both α and β , and vice versa (this ensures that we obtain a value for the proportion
 3370 of time spend undergoing a dose of α which lies in the range $[0, 1]$);
- 3371 6. $\frac{-H_0}{R_0(\exp(T(g_R^\beta - g_H^\beta)))} \leq m \leq \frac{-H_0}{R_0(\exp(T(g_R^\alpha - g_H^\alpha)))}$ (another condition which ensures that we
 3372 obtain a value for the proportion of time spend undergoing a dose of α which lies in
 3373 the range $[0, 1]$ – the necessity of this condition and the previous one are made clearer
 3374 once we derive our proposed optimal proportion of time spent at α).

3375 Assuming these conditions are satisfied, we can begin to investigate how we might select
 3376 appropriate lengths of time to spend at the doses α and β . Firstly, we note that if $H_0 \gg R_0$ or
 3377 $R_0 \gg H_0$, then an alternating treatment will only be beneficial with respect to final tumour
 3378 size if T is sufficiently large. Suppose we have an arbitrarily short treatment period T , and
 3379 $H_0 \gg R_0$, and say for the sake of example that α reduces H while allowing R to grow,
 3380 while β does the opposite. The best strategy to minimise the final tumour size at final time T
 3381 would be to treat with a dose of α only for the duration of T . Generally speaking, we would
 3382 only want to switch to β if $R(t) > H(t)$ at some time t . As such, T would have to be large
 3383 enough to facilitate this in order for it to be beneficial to switch to any other treatment in this
 3384 case.

3385 Assume, then, that we have it that neither $H_0 \gg R_0$ nor $R_0 \gg H_0$. We have it, then
 3386 , that a treatment alternating only between the doses α and β will permit a lower average
 3387 growth rate for R than any of those found on the treatment curve $g_R(g_H(a))$. This is satisfied
 3388 immediately by condition 1 above, since if we alternate only between doses α and β , our
 3389 average growth rate $(\hat{g}_H, \hat{g}_R)$ will lie somewhere on the line defined by $g'_R = mg'_H + c$, and
 3390 this line acts as a lower bound for the treatment space bounded by the treatment curve and
 3391 the line itself.

3392 Consider now a treatment period of length T , and suppose that we wish to alternate
 3393 between α and β only. We desire the proportion of time spent at α , $P_\alpha = T_\alpha/T$, which
 3394 minimises the final tumour size $N(T)$. To find this, we start with the final tumour size, $N(T)$
 3395 – this is given by

$$N(T) = H_0 \exp(\hat{g}_H T) + R_0 \exp(\hat{g}_R T),$$

where $\hat{g}_H, \hat{g}_R$ are the average growth rates of H and R across the time period T . Since we are alternating between α and β , the average growth rate g_R can be expressed in terms of g_H by substituting in the expression previously defined for g'_R , yielding

$$N(T) = H_0 \exp(\hat{g}_H T) + R_0 \exp((m\hat{g}_H + c)T).$$

We wish to minimise this expression with respect to the growth rate, so we calculate the derivative of $N(T)$ with respect to $\hat{g}_H$:

$$\frac{dN}{d\hat{g}_H} = H_0 T \exp(\hat{g}_H T) + R_0 m T \exp(cT + m\hat{g}_H T).$$

Setting this equal to zero and solving for $\hat{g}_H$ yields the average growth rate which minimises $N(T)$,

$$\hat{g}_H^{\min} = \frac{\ln(\frac{-H_0}{R_0 m}) - cT}{T(m-1)}.$$

Now, we note that if we only alternate between α and β , then we can express the average growth rate $\hat{g}_H$ in terms of α and β :

$$\hat{g}_H = \frac{g_H^\alpha T_\alpha + g_H^\beta T_\beta}{T}.$$

Since T, g_H^α and g_H^β are fixed, we can find the time T_α which minimises the final tumour size $N(T)$ by equating this expression with $\hat{g}_H^{\min}$, which, after rearranging, yields

$$T_\alpha = \frac{\ln(\frac{-H_0}{R_0 m}) - T(c + g_H^\beta(m-1))}{(m-1)(g_H^\alpha - g_H^\beta)}.$$

If we divide this value by the total time, T , and substitute $g_R^\alpha = mg_H^\alpha + c$ and $g_R^\beta = mg_H^\beta + c$ into this expression where appropriate, we can find the proportion of time P_α which we should spend at dose α ,

$$P_\alpha = \frac{\frac{1}{T} \ln(\frac{-H_0}{R_0 m}) + (g_H^\beta - g_R^\beta)}{(g_R^\alpha - g_H^\alpha) + (g_H^\beta - g_R^\beta)}.$$

3410 We can extend this result in a few ways. Consider that we can alternatively express the
 3411 derivative $dN/d\hat{g}_H = 0$ as

$$\frac{dN}{d\hat{g}_H} = H_0 T \exp(\hat{g}_H T) + R_0 m T \exp(cT + m\hat{g}_H T).$$

3412 Since $\hat{g}_R = m\hat{g}_H + c$, equating this derivative to 0 and rearranging yields

$$\hat{g}_H - \hat{g}_R = \frac{1}{T} \ln\left(\frac{-mR_0}{H_0}\right).$$

3413 For large T , we have that $\hat{g}_H - \hat{g}_R \approx 0$, or equivalently, $\hat{g}_H \approx \hat{g}_R$. In other words, the point on
 3414 the line $g'_R = mg'_H + c$ which will yield the minimum tumour size is approximately the point
 3415 at which it intersects the line $g_H = g_R$. This corresponds to a testosterone concentration; as
 3416 such, we have a means of finding out the optimal average testosterone concentration for the
 3417 treatment period, and by extension the approximately optimal ratio of T_α to T_β , graphically.

3418 With this result derived, we can justify why conditions 5 and 6 are necessary. Suppose
 3419 we have it that H grows more quickly than R under both doses α and β . If we were to make a
 3420 plot similar to figure 4.4 in this instance, the points corresponding to the doses α and β would
 3421 lie to the right hand side of the line $g_H = g_R$. For an arbitrarily long treatment period T , we
 3422 would want to spend as much time as possible at the dose at which H grows the slowest to
 3423 minimise the final tumour size. The result we have presented above shows that we want to
 3424 aim for the point on the line $g'_R = mg'_H + c$, bounded by the points corresponding to α and
 3425 β , which intersects the line $g_H = g_R$. Here, both of our doses lie to the right of this line,
 3426 and so no such point exists. As such, we would need to aim for a point on $g'_R = mg'_H + c$
 3427 which lies further to the left than both α and β . This is not possible if we can only alternate
 3428 between α and β , unless we permit treatment times at either dose which are greater than the
 3429 total treatment time T . This would give us a value for P_α that lies outside $[0, 1]$, and so we
 3430 can see why condition 5 is required. A similar issue arises if R grows faster than H at both
 3431 α and β . To justify why we require condition 6, consider our expression for P_α if we take a
 3432 value for m which is arbitrarily close to 0. As $m \rightarrow 0$ from below, the numerator in P_α , and
 3433 thus P_α itself, becomes arbitrarily large. If we set $P_\alpha = 0$ and solve for m , we obtain the
 3434 lower bound in condition 6; the upper bound is obtained by setting $P_\alpha = 1$.

We can approximate the proportion P_α using the growth rates of H and R at doses α and β alone. If we consider the expression for P_α above in the limit $T \rightarrow \infty$, then we see that

$$P_\alpha \approx \frac{g_H^\beta - g_R^\beta}{g_H^\beta - g_R^\beta + g_R^\alpha - g_H^\alpha}.$$

These results tell us that the proportion of time spent undergoing both treatments, surprisingly, depends not only on the growth rates of both cell types at doses α and β , but also on both the initial proportions of H and R and, perhaps most surprisingly, the treatment period T itself. This means that the proportion of time we spend at the two doses will differ with two different total treatment times, even if all other conditions are the same. Corollary 2, however, demonstrates that this will matter less and less as we extend the treatment period. As such, if we have a very long BAT treatment period, we see that the ratios of high and low testosterone treatment will be roughly equal, and will predominantly depend on the growth rates of H and R at testosterone doses α and β .

4.3.5 Analytical Approach to Treatment Scheduling in a Model Approximating Continuous Testosterone Decay

In order to establish how close to optimal the results for our instant switching model are to a model with continuous testosterone decay, we define an analytically tractable model which approximates this continuous decay with a step function. We define the model in largely the same way as before, but make a few modifications. When we switch from the low dose α to the high dose β , we assume that we can still do so instantaneously, since in a clinical setting, testosterone is injected into the patient, which would rapidly increase the blood testosterone concentration. To approximate the decay of testosterone, we add an amendment to the model which stipulates that any time we switch from the high dose β to the low dose α , we must spend a period of d days at an intermediate concentration δ , at which we have that $(g_H, g_R) = (g_H^\delta, g_R^\delta)$. How we can select appropriate values for (g_H^δ, g_R^δ) is outlined in the methods. At this intermediate point, we would generally expect both an increase in both H and R cell populations.

Our model is as follows:

$$\begin{aligned}\frac{dH}{dt} &= g_H(a)H(t), \\ \frac{dR}{dt} &= g_R(a)R(t),\end{aligned}$$

3461 where

$$g_H(a) = \begin{cases} g_H^\alpha & \text{if } a = \alpha, \\ g_H^\beta & \text{if } a = \beta, \\ g_H^\delta & \text{if } a = \delta, \end{cases} \quad (4.5)$$

3462

$$g_R(a) = \begin{cases} g_R^\alpha & \text{if } a = \alpha, \\ g_R^\beta & \text{if } a = \beta, \\ g_R^\delta & \text{if } a = \delta, \end{cases} \quad (4.6)$$

3463 where $g_H^\alpha, g_H^\beta, g_R^\alpha$ and g_R^β are defined as before. We assume that the amount of time we spend
 3464 at the intermediate delay concentration δ is fixed – we would decide prior to treatment how
 3465 many times we are going to switch between α and β (which we do only for the sake of model
 3466 analysis – this would be inadvisable in a clinical setting, which we discuss later). As such,
 3467 what we need to find is the optimal amount of time we spend at the doses α and β after we
 3468 subtract from T the time spent at the delay concentration δ . We denote by $T_\delta \leq T$ the total
 3469 time spent at the dose δ . With T_α, T_β defined as before, we have that for a treatment period
 3470 $T, T = T_\alpha + T_\beta + T_\delta$.

3471 Assume that the conditions from the previous section are satisfied. Then, for a treatment
 3472 period T , we wish to alternate between α and β only, accounting for a delay of length d
 3473 (where $T_\delta = nd$, for some $n \in \mathbb{N}$) when switching from β to α . we wish, once again, to find
 3474 the proportion of time spent at dose α , $P_\alpha = T_\alpha/T$, which minimises the final tumour size
 3475 $N(T)$.

3476 We have that the total final tumour size $N(T)$ is

$$N(T) = H_0 \exp(\hat{g}_H T) + R_0 \exp(\hat{g}_R T).$$

3477 We can express $\hat{g}_H$ and $\hat{g}_R$ as weighted averages of their growth rates at α , β and δ . which
 3478 yields

$$N(T) = H_0 \exp(g_H^\alpha T_\alpha + g_H^\beta T_\beta + g_H^\delta T_\delta) + R_0 \exp(g_R^\alpha T_\alpha + g_R^\beta T_\beta + g_R^\delta T_\delta).$$

3479 We recall that T_δ is fixed, and that $T_\beta = T - T_\alpha - T_\delta$. Replacing T_β with this expression
 3480 and differentiating with respect to T_α yields

$$\begin{aligned} \frac{dN(T)}{dT_\alpha} &= H_0(g_H^\alpha - g_H^\beta) \exp(g_H^\alpha T_\alpha + g_H^\beta (T - T_\alpha - T_\delta) + g_H^\delta T_\delta) + \\ &\quad R_0(g_R^\alpha - g_R^\beta) \exp(g_R^\alpha T_\alpha + g_R^\beta (T - T_\alpha - T_\delta) + g_R^\delta T_\delta). \end{aligned}$$

3481 Equating the above expression to 0 yields

$$T_\alpha = \frac{\ln \frac{-H_0(g_H^\beta - g_H^\alpha)}{R_0(g_R^\alpha - g_R^\beta)} + T(g_H^\beta - g_R^\beta) + T_\delta(g_R^\beta - g_H^\beta + g_H^\delta - g_R^\delta)}{(g_R^\alpha - g_R^\beta) + (g_H^\beta - g_H^\alpha)}.$$

3482 Dividing this by T yields the proportion of time we should spend at α , P_α :

$$P_\alpha = \frac{\frac{1}{T} \ln \frac{-H_0(g_H^\beta - g_H^\alpha)}{R_0(g_R^\alpha - g_R^\beta)} + (g_H^\beta - g_R^\beta) + \frac{T_\delta}{T}(g_R^\beta - g_H^\beta + g_H^\delta - g_R^\delta)}{(g_R^\alpha - g_R^\beta) + (g_H^\beta - g_H^\alpha)}.$$

3483 We note that as $T_\delta \rightarrow 0$, then in the limit $T \rightarrow \infty$, we have it that

$$P_\alpha \approx \frac{g_H^\beta - g_R^\beta}{g_H^\beta - g_R^\beta + g_R^\alpha - g_H^\alpha},$$

3484 which is the same approximate value for P_α as in our instant switching model. If we have
 3485 a sufficiently small T_δ , this approximation will still hold, but as T_δ increases relative to the
 3486 size of T , this approximation will become increasingly less accurate. It follows, then, that
 3487 minimising the amount of time that we spend at the dose δ will yield a treatment strategy
 3488 which is approximately equal to the strategy we present for a model where switching between
 3489 doses is instant.

4.3.6 Analysis of Switching Frequency shows Potential Benefit of Switching More than Once

In a model where we can switch instantaneously, how and when we switch has no difference on the final cell number $N(T)$. What needs to be considered, however, is that treating with only α or β for a prolonged period of time will select for H or R for a long time, allowing one to grow unchecked for this time period. Say, for example, we start treatment with α for a time T_α , which is optimal with respect to $N(T)$. Assuming α selects for R cells, we will see a period of T_α days in which R becomes very large – with a large enough T , this may be several orders of magnitude larger than both $N(0)$ and $N(T)$. In a clinical setting, this would potentially lead to adverse effects on quality of life, and in extreme cases, this could lead to a tumour which is a potentially lethal size in the middle of the treatment period T . The solution to this problem, then, would be to break T_α up into several shorter treatment periods of equal length. These would be interspersed with periods in which we treat with β , selecting for H instead, and decreasing the overall tumour size throughout treatment, whilst still preserving the optimal final tumour count $N(T)$. If we wish to minimise tumour size throughout treatment, then we should switch as frequently as possible – this would minimise the amount of time we select for either H or R continuously.

Additional considerations need to be made, however, when we consider our model which includes a delay when switching doses. Firstly, we must consider that if we switch too frequently, or if our delay period is too long, then we arrive in a position where $T_\delta > T_\alpha$ or $T_\delta > T_\beta$. In other words, we may end up spending more time at the intermediate dose δ than at the desired doses α or β . This gives us a simple upper bound on the number of switches we apply, since ideally, as a bare minimum, we wish to spend more time undergoing the desired beneficial treatments than we want to spend at the delay concentration δ . It also provides a convenient indicator as to whether it is beneficial to switch at all for a treatment period. If we decide to switch only once, but we still find that T_δ is longer than T_α or T_β , then it follows that the best strategy may be one which does not switch from β to α at all (switching from α to β does not incur a delay, so it should be noted that T_δ being longer than T_α or T_β does not necessarily suggest that a protocol in which we do not switch at all is optimal).

We do note, however, that it may still be beneficial to spend more time at the delay dose

δ than at α or β . Consider, for example, the case in which the treatment curve is slightly concave, meaning that we do not incur much of a penalty during the transition between β and α , where spending time at either of these doses has a drastic and immediate effect on H and R . If the effect of these doses is so profound that spending even a day or two at either would have a huge effect, switching would be beneficial, even if T_δ is large relative to the total treatment period. When this is not the case, the previous condition of ensuring $T_\delta \leq T_\alpha$ or T_β gives us a convenient upper bound on the number of switches we should choose.

Consider a number of switches n , and assume our choices of T_α and T_β are optimal. If we split up T_α , T_β evenly, then we shall have n phases of low and high testosterone, with lengths of T_α/n and T_β/n , respectively. Generally, we should try to avoid a choice of n which is large enough that we have that the delay length, d , is longer than either T_α/n or T_β/n , since this leaves us in a position where we would spend more time at the delay dosage δ than at the desired doses of α or β . Furthermore, given that (g_H^δ, g_R^δ) lies above the line $g_R' = mg_H' + c$, we have it that $g_R^\delta > g_R^\alpha, g_R^\beta$. This means that every time we switch from β to α , we incur a growth penalty which increases the final cell number $N(T)$. The more often we switch, the greater this penalty on $N(T)$ will be.

This already conflicts with our findings from the instantaneous switching model, so we might assume that our best option would be to switch as little as possible, thus maximising the amount of time we spend undergoing treatment with doses α and β . This leads us to a dilemma – if we want to minimise the final cell number, we should switch as little as possible. This could, however, leave us in a situation where we have a very large tumour in the middle of treatment – in order to minimise this, we should switch more often. We need, then, to find some middle ground between switching doses as frequently as possible to minimise the tumour size throughout the treatment duration, and switching as little as possible to minimise the tumour size at the end of our treatment period.

In order to investigate this, we perform a number of simulations in which we gradually increase the number of switches to observe the effect this has on both the final cell number, $N(T)$, and the mean average cell number throughout treatment,

$$\frac{\sum_{t=0}^T N(t)}{T}.$$




FIGURE 4.5: **(Left)** Simulation result demonstrating how the number of switches effects the final cell number $N(T)$. **(Right)** Simulation result demonstrating how the number of switches effects the mean average cell number across treatment. In both simulations, the optimal times for T_α and T_β were calculated using the expression obtained in proposition 4. They were then split into evenly sized periods to accommodate the number of switches. Initial cell proportions were 1:1, with an initial cell number $N(0)$ of 100 million. The treatment period was set to 1000 days for these simulations, so that we could look at larger numbers of switches, giving us a better view of the relationships we are interested in.

We see the results of these simulations in the delay model in figure 4.5. The effect that the number of switches has on $N(T)$ is unsurprising: as we have stated, each switch effectively acts as a penalty on the final cell number. We would expect that increasing the number of switches would increase the final cell number in turn. The effect the number of switches has on mean cell number is also as we expected. Switching too infrequently means that the mean tumour size throughout treatment is larger than the final cell number – in the case where we only switch once, for example, we see that the mean cell number is almost two orders of magnitude larger than the final cell number. If we switch too frequently, we see that both the mean cell number and the final cell number are large. We can see, however, that if we switch only a handful of times, the mean cell number is at its lowest here – with these parameters, the minimum mean cell number is achieved if we switch 4 times, but the code used to find this optimum number of switches numerically is easily generalisable to other parameterisations.

4.3.7 Model Simulations Demonstrate Efficacy of Our Proposed Treatment Protocol

In order to determine whether or not our analytical results will hold in a model where testosterone decays over time, we simulate the result in R using the R package *deSolve* (236).




FIGURE 4.6: (Left) Results of simulating our proposed treatment strategy in our initial model, where H , R and N cells are represented as in figure 4.3. (Right) Comparison of the total cell numbers throughout treatment for an ADT protocol (cyan), the standard BAT protocol seen in the literature (green), and our proposed protocol (magenta). Parameters used for simulation can be found in figure 4.1.

Our model predicts that ADT and BAT are similar in terms of final cell number at the end of treatment, and in the short term, BAT is actually worse than ADT treatment in terms of tumour burden. We see in figure 4.6 that our proposed protocol outperforms both the standard BAT protocol and ADT alone in the long term, enabling a ten-fold reduction in final cell number. Our proposed protocol utilises the above result from the delay model, giving us an approximate proportion of time we should spend undergoing high- and low-testosterone treatments. If we compare our findings to the model simulations of BAT and ADT in figure 4.3, we can see that our model also maintains a higher proportion of hormone-naïve cells when compared to other protocols for the majority of the treatment period. We do find, however, that using the same optimal split of treatment as we do in the delay model selects for the H cells more than we would expect. This suggests our proposed protocol spends too much time in a high-testosterone state.

To address this, we make an adjustment to the proposed protocol, in which we manually reduce the number of days total we spend undergoing high-testosterone therapy, and increasing the number of days spent undergoing low-testosterone therapy accordingly. The results of this simulation, and a comparison between it and the results shown in figure 4.6, can be found in figure 4.7. We see that in the short term, the differences between our two proposed protocols is marginal, but as time progresses, the superiority of the manually adjusted protocol becomes clear, with total cell number reaching a lower plateau during a treatment phase




FIGURE 4.7: (Left) Results of simulating our proposed treatment strategy in our initial model, where H , R and N cells are represented as in figure 4.3 and 4.6. (Right) Comparison of the total cell numbers throughout treatment for an ADT protocol (cyan), the standard BAT protocol seen in the literature (green), the proposed protocol from figure 4.6 (magenta) and our new, manually adjusted proposed protocol (orange). In these simulations, we decrease the total time spent at a high-testosterone dose by 8 days, and increase the low-testosterone treatment period accordingly. Parameters used for simulation can be found in figure 4.1.

in the new protocol than it did in the original proposed protocol. Final outcomes for both protocols are similar, with the manually adjusted protocol having a slightly lower final cell number than the original one which we derived from our analytical results.

With both of our proposed strategies, after one cycle of high and low testosterone, net increase in tumour size when compared to the initial size is similar to that seen with an ADT protocol, and less than that seen in a standard BAT protocol. As the treatment period continues, however, we see that net increase in tumour size for ADT and standard BAT protocols becomes greater, and deviates further from our proposed protocols. We posit that the clinical gains from our proposed protocols would be even greater in the long term when compared to outcomes seen in ADT only protocols and standard BAT protocols on similar time scales.

4.4 Discussion

Our analytical and numerical results provide us with not only an understanding of the key mechanisms involved in BAT, but a general understanding of how one might approach designing a roughly optimal treatment for an alternating therapy where cells see a trade-off

between two differing drug doses or treatments, given sufficient information about the behaviours of the cells during said treatments. They brought us to a treatment strategy which outperforms other strategies seen in the literature. Furthermore, our model predicts that BAT and ADT will have similar outcomes at the end of a given treatment period. Given that one of the benefits of BAT is a reduction in adverse QoL effects, we posit that the BAT protocols seen in previous literature could be considered as an alternative to continuous ADT without any loss in clinical gains, but advocate for our own proposed protocols, which use key growth parameters of both hormone-naïve and castrate-resistant prostate cancer cells to derive proportions for high- and low-testosterone treatments that massively increase clinical gains.

A number of pre-existing models of CADT and IADT include PSA as a key factor in their own investigations (226; 221; 223; 224), but whether this will be a good decision for modelling BAT is up for debate. PSA responses in clinical trials of BAT are inconsistent, but not inconclusive – PSA response was only seen in two fewer patients than radiographic response in Schweizer *et al*'s (187) pilot clinical trial, but only 9 of 30 patients displayed PSA response after BAT alone in Teply *et al*'s clinical trial (188). This is not necessarily surprising, since PSA production is linked directly to testosterone levels, so we would presume that PSA levels in BAT patients would often be higher than in patients undergoing alternative treatments. We do note, however, that it may be an interesting question in and of itself to use a model to try and investigate the disparity in PSA responses between patients undergoing ADT and those undergoing BAT. Ideally, models could be fitted to patient PSA data for both treatments to determine precisely how effective PSA is as an indicator of disease progression during BAT. This, paired with the fact that we do not intend to match the model to any clinical data at present, resulted in our decision to exclude PSA production as a component of our model. A potential extension of this work could be to investigate a model which also captures how PSA dynamics vary during BAT, but clinical data on how PSA varies in response to testosterone injections would be needed to calibrate such a model.

Our findings indicate that the restoration of ADT sensitivity after BAT is owed to the fact that large amounts of testosterone in the system create a favourable environment for extant

sensitive cells, allowing them to grow rapidly in a short period, whilst simultaneously creating an environment which harms resistant cells within the tumour. Although this is in line with clinical findings in a holistic sense, the reality is likely more complex. We assume in our model that the mechanism of resistance to ADT is an upregulation of the androgen receptor, leading to hypersensitivity which facilitates androgen-dependent growth in low testosterone concentrations. Supraphysiological concentrations of testosterone have been shown *in vitro* to inhibit growth of CRPC cells where androgen receptor activity was observed (183; 184), but this activity could be owing to a number of different mechanisms. Although upregulation of the AR leading to hypersensitivity has been observed often in CRPC (199; 200; 198), different mechanisms can drive resistance. We would hypothesise that BAT would still be effective in these instances, but it would be more likely that its efficacy relied on high testosterone levels “tricking” castrate-resistant cells into abandoning their adaptive response to ADT. A different model would be required to investigate such a phenomenon, in which H and R cells are allowed to mutate back and forth according to concentrations of testosterone – this would be a meaningful avenue for further investigation. If we found that the success of BAT does indeed rely on hypersensitivity of the androgen receptor, we could see improved clinical gains by stratifying CRPC patient cohorts: patients presenting with CRPC which exhibits hypersensitivity of the androgen receptor would see benefits from BAT treatment; patients which do not would be assigned standard-of-care treatments such as androgen deprivation therapy. Even if this information was not known, a short initial trial of BAT could be attempted, and treatment could either be continued or accordingly adapted depending on patient response.

In a similar vein, our model assumes a linear relationship between testosterone and the growth and death rates of HN cells, but one of the clinical studies which inspired BAT found that hormone-naïve prostate cancer cells could grow optimally in the absence of any androgens, and saw growth inhibition when large concentrations of androgen were introduced. We opted to omit this as a factor in the model to avoid any overly optimistic predictions, since incorporating this could see HN cells behaving similarly to CR cells at high testosterone concentrations, meaning that one could trivially cure the disease by maintaining a constant high testosterone concentration. Since we do not see any cure of disease in any clinical trials, it is

likely that this HN cell behaviour, which was observed in an *in vitro* lab setting, may not be as pronounced in a real patient, if present at all. This is unlike CR cells, for which we have *in vivo* evidence of this non-trivial CR cell behaviour (185; 186).

A potential factor for the success of BAT could be that HN cells are conferred a competitive advantage during periods where testosterone is high and their numbers are replenished. In a real tumour, when HN cells are abundant relative to CR cells, we would assume that they are more likely to surround CR cells completely, depriving them of the space needed for them to divide and potentially containing them until the ADT treatment following BAT depletes the HN cell population again. We do not assume that such competition is taking place in our model since our primary focus is how cells are affected by varying testosterone dynamics, but this could be an interesting avenue to pursue, both experimentally and with further modelling efforts.

We see at the end of the results section that adjusting the treatment cycles by only a few day has a notable effect on the qualitative outcomes of our treatment. We conclude, then, that the optimal treatment schedule we derive from our simplified model is highly sensitive to even minor differences in the lengths of high- and low-testosterone treatment periods, and optimal treatment schedules will likely vary greatly if different parameter values are chosen for the model. As such, we cannot propose a specific treatment schedule based on our results, since we would assume that any differences in parameter values, even minor ones, could majorly impact derived treatment times. We instead propose the analytical results above as a framework for a patient-specific treatment strategy. We do note, however, that the potential clinical implications of this sensitivity to minor perturbations will likely not be of major concern, so long we only overshoot or undershoot treatment periods by a small amount. These small differences will add up for long overall treatment periods (or, similarly, if we have it that high testosterone or low testosterone treatment periods very rapidly increase the populations of H or R , respectively), so care would have to be taken in these instances so that we do not deviate too far from our derived optimal treatment schedule.

One major consideration for our model, in terms of its biological relevance, is the frequency with which we administer injections of testosterone into the system. Our findings suggest that approximately optimal treatments should maintain high levels of testosterone

3687 for a fixed period, rather than constantly cycling between low and high doses using monthly
3688 injections followed by ADT. Although the consistency in scheduling of testosterone in clini-
3689 cal trials was a motivating factor in developing the model, injecting high doses of testosterone
3690 more frequently than every 28 days may exacerbate the negative side effects associated with
3691 testosterone injections, such as high blood pressure, potentially putting the patient at risk of
3692 a severe cardiac event or stroke (237). Collaboration with clinicians would ensure that the
3693 findings of the model stay within realistic clinical bounds in future, and would be essential in
3694 further modelling efforts if findings are to be clinically actionable.

Chapter 5

Bayesian Modelling of Cancer

Incidence Data to Investigate

Increased Early Onset Cancer

5.1 Foreword

The results and analyses in this chapter have been gathered in collaboration with David Robert Grimes (TCD Biostatistics Unit, Trinity College Dublin), David Basanta (Moffitt Cancer Center, Department of Integrated Mathematical Oncology), and Robert Noble (City St. George's, University of London, Department of Mathematics), in preparation for a publication. Analysis and modelling is still currently ongoing. David Robert Grimes proposed the investigation itself, and suggested the use of Weibull-type models to investigate pan-cancer data. David Basanta and Robert Noble assisted with the background reading and analysis of statistical results, and contributed the introduction to this chapter. All four authors contributed to the discussion. The Bayesian statistical modelling of the Pan-cancer data and the data pertaining to specific cancers, as well as the analysis of the results of this modelling, was done by myself.

5.2 Introduction

Cancer is primarily a disease of ageing, primarily occurring in older adults. In the US, the median age of cancer diagnosis is 66 years of age (238) consistent with data from England

where it has stayed steady with median diagnosis age in the 65-69 age bracket (239), Ireland where the median age at diagnosis for all cancer combined (excluding non-melanoma skin cancers) was 69 years for men and 67 years for women between 1994 and 2020 (240) and 63-64 years in Spain (241; 242). Worldwide, the burden of cancer incidence and mortality has grown rapidly, reflecting both an ageing and growing world population as well as changes in the distribution and prevalence of risk factors for cancer, many of which are closely associated with socioeconomic development (243; 244).

While much of this increase in incidence is understood to be the consequence of an aging population and cancer remains primarily a disease of older adults, there is emerging evidence that incidence of cancer in younger adults (< 50 years) is increasing in high and middle social demographic regions (245; 246). This trend has been seen in many cancer types but it is more pronounced in certain cancers such as colorectal, GI, and breast (247; 245). Many factors have been proposed as potential explanations of this apparent trends, including lifestyle (changing diet, higher levels of obesity), and environmental (high exposure to unidentified environmental carcinogens).

An alternative consideration that has the potential to explain these trend is improved detection. Over the past decades, advances in diagnostic imaging, national screening programmes, and increased awareness of cancer have markedly improved detection times for many common cancers. breast cancer screening in the United States, for example, hugely increased the number of early stage cancers detected by up to a factor of 10 between 1977 and 2012 (248). Depending on age group, Mean Sojourn Time (MST) for breast cancer, the average time spent in the preclinical screen-detectable phase of a disease is one measure of how much earlier a cancer may be detected by screening. For breast cancer, this varies by age group with historical estimates for MST ranging from 2-7 years (249; 250). This also varies by period since the introduction of screening; one Dutch study over a 40 year period suggested MST increasing from approximately 2.5 years to 4.6 years (251).

While breast cancer is the most common cancer diagnosed in women, MST has also increased for a number of other cancers due largely to national and regional screening programmes. Screening programmes for colorectal cancers lead to an estimated MST of 3-6 years, while cervical cancer screening is an unqualified success, detecting not only early

stage cancers but abnormalities that could become cancer many years before they progress (252; 253). In HDI countries where apparent cancer incidence has increased, age-adjusted mortality from cancer continues to decline. These factors together raise an important question: could the apparent rise in young cancers be explained in part or completely by improved early detection of cancer?

Despite the impact those advances could have, early detection has been discarded on account that mortality rates (for instance in colorectal cancer) have increased consistently on a variety of regions regardless of changes in screening policy (254). While the answer is likely to be that both improved detection and lifestyle and environmental change drive the dramatic increase in early onset cancer incidence, it is not clear which of these two factors is the main driver of this. The motivating question of this work is an important one. If the answer is that environmental and lifestyle changes are the most important, it would suggest a true rise in cancer rates and require a deeper mechanistic understanding of novel factors in younger people or their environment affecting cancer induction. If improved detection, the seemingly worrying increase could actually be indicative of improved diagnostic ability. In this work, we derive a simple mathematical model for two competing scenarios: the first where there is a genuine change in the rate of cancer onset and the second where the rate remains constant but the MST of cancers increases with time. The mathematical model proposed is intentionally simple, with a minimum number of parameters to simply observe the consequences of both competing assumptions and what their fit to data might imply.

Steven Frank gives a brief historical overview of age-specific cancer incidence modelling in his book on the dynamics of cancer (255). Two phenomena are observed in cancer incidence data – firstly, incidence increases rapidly with age; secondly, that patterns emerge at a population level despite the fact that cancers are highly stochastic at the patient level.

In (256), Charles and Luce-Clausen investigate skin tumours on mice exposed to carcinogenic paint. They are among the first to find a relationship between the length of time exposed to a carcinogen and the probability of mutations occurring in single alleles, where they assume that the same allele must mutate twice before they transform into an observable skin cancer. Whittemore and Keller (257) present an explanation for the exponential increase of incidence with age which hinge on two different model assumptions. The first assumption

assumes that carcinogenic mutation events happen to approximately n cells simultaneously, leading to an age-specific incidence proportional to t^{n-1} , with age t . The second assumption assumes that changes to cancer cells in a tissue region happen sequentially, similar to the relationship proposed by Charles and Luce-Clausen who assumed two sequential mutations. Whittemore and Keller assume n sequential mutations, which Muller also discusses in (258). Nordling introduced log-log plots which could be used to approximate the number of sequential mutations in (259). Armitage and Doll (260) extended this result by looking at distinct cancers, allowing different rates for different steps in the accumulation of sequential mutations. If we presume n mutations must occur before cancer is observable, then Armitage and Doll showed that if mutation rate increases with carcinogen dose d , then incidence is proportional to d^n . Whittemore and Keller expanded upon this result, showing that either mutations to multiple cells simultaneously or sequential mutations on one cell both yield models which fit incidence cancer well. What we can conclude from this is that we can explain the trends in cancer incidence rates through the lens of multi-stage events which drive progression of cancer. At a cellular level, these events either happen to multiple cells simultaneously or must affect individual cells multiple times.

In Frank's book, he discusses the notion of modelling cancer incidence using reliability theory, a topic reviewed by Gavrillov and Gavrillova in (261). Reliability considers time to failure (which, in our case, would be the death of a patient) based on assumptions we make about failure rates of individual components and the relations between these components. The multi-stage theories presented by the authors above can be viewed as instances of component failure – in this case, cells accumulating cancerous mutations. Frank discusses how Weibull models are often used to describe the relationship between failure rates and age. Making no assumptions about underlying processes, they allow us to investigate incidence data and compare datasets looking at only a few parameters which describe trends in the data. To examine these competing explanations for changing trends in cancer incidence, a robust and reliable database of cancer incidence by age cohort over time is vital. Accordingly, the simple model was contrasted with data from the American SEER Surveillance, Epidemiology, and End Results (SEER) Program conducted by the National Cancer Institute database,

extracted with the DevCan (NCI) software package which applies demographic and statistical models to SEER data to compute the probability of being diagnosed or dying of cancer from birth or conditional on a certain age, adjusted for competing risks. This yields probabilities of developing or dying from cancer for a hypothetical population based on different SEER data periods, typically reported every three years from 1977 onwards. This gives an evolving picture of cancer rates by age cohort since 1977, and a laboratory to examine the possible consequences of both scenarios and their fit to DevCan data.

5.3 Methods

5.3.1 Data Sources for Initial Analysis of Diagnostic Trends

Given that the changes in cancer trends should be fairly consistent across the globe, we examine cancer data sets taken from the US Surveillance, Epidemiology and End Results (SEER) database. Using DevCan 6.7.5, we extracted the data from SEER for cumulative probability of developing cancer by age cohort. In particular, we extract the US SEER 22 Incidence and Mortality data, which covers the years 2000 to 2021 inclusive (which excludes the year 2020, the data for which is incomplete because of the COVID pandemic), and the SEER8 database which covers the years 1977-1999. Datasets were exported covering all races and genders without stratification, and cover all in-situ and malignant cancer subtypes. From these datasets, cumulative incidence for $t = 50$ years of age for each year group was fit using a linear regression model to confirm that the slope of the regression line was positive, showing that incidence for this year group has indeed increased over time – this is shown in figure 5.2. A linear least-squares algorithm was then used to investigate whether or not fixing the parameter k (explained below) was reasonable, with the general form of the following model equation:

$$p(t, n) = 1 - \exp(-(\beta(n)t)^k).$$

The resulting k values were then used for model fitting. Fixed values used in any inference are stated where appropriate.

5.3.2 Cancer Incidence Data and Fitting

The SEER data sets we analyse record cumulative cancer incidence every three years. As such, we model the data using a Weibull-type model with a rate parameter β and shape parameter k . The latter dictates time dynamics of cancer onset, and if we have it that $k > 1$, cancer rates will increase with age. We assume in our modelling that k is fixed for a given dataset. This should be a reasonable assumption given that cancer induction should increase as age increases and that cancer induction dynamics, such as hazard exposure and eventual development of cancer, should remain roughly consistent over time. Our rate parameter β , however, relates to how early cancers are induced, and it is this parameter which we wish to investigate.

We propose that an increase in incidence for patients aged 50 could potentially be explained by two mechanisms. The first mechanism would be a true increase in cancer incidence in patients under 50: this would result in $\beta(n)$ varying with time (in this case, the year to which the data corresponds, n). In our analyses, we model this potential mechanism as follows:

$$p(t, n) = 1 - \exp(-(\beta(n)t)^k), \quad (5.1)$$

Where k and $\beta(n)$ are as above and t is the age cohort for a given year n . The alternative mechanism which explains the increase in early onset cancer would be an increase in mean sojourn time. This would correspond to a decrease in lag detection time, which we denote Δ_t . This represents the case where cancers are detected earlier over time, which would also increase the early onset cancer rate. In this case, we fix both k and β , and investigate the changes in Δ_t relative to the first year in our SEER dataset, 1977. We model this potential mechanism as

$$p(t, n) = 1 - \exp(-(\beta(t + \Delta_t(n)))^k). \quad (5.2)$$

Our data analysis of the incidence data showed that incidence trends for ages 0-25 and 85-90 were highly variable. This is not surprising, given that we would not expect to see the same trends in adult cancers that we would in younger or elderly age groups. As such, model analysis was done looking only at ages 30-80.

Bayesian Hierarchical Model Outline and Implementation

We implement a non-centered hierarchical Bayesian model in *Stan* (56), an open source inference software to estimate these parameters via a repeatable Hamiltonian Monte Carlo Markov chain simulation. The benefit of this approach is that we obtain a distribution of plausible values for each fitted parameter, rather than a point estimate with a confidence interval, allowing us to simulate a range of possible model trajectories. The hierarchical structure of the model allows us to obtain additional information by pooling the data from each year group together. Bayesian model fits were performed on log-transformed data, which was rescaled so that incidence was measured per 10,000 people to aid with inference.

In a Bayesian model, we first define a joint distribution which links our observed data for a given year group, p , and the parameter value we wish to infer β ,

$$P(p, \beta).$$

Stan evaluates the log density of $P(p, \beta)$, which we can rewrite as the product of the prior density $P(\beta)$ and the likelihood $P(p|\beta)$. We seek to find the posterior distribution $P(\beta|p)$, which is the set of potential parameter values given the data we have, which we construct using Bayes' rule using the prior density and likelihood,

$$P(\beta|p) \propto P(\beta)P(p|\beta).$$

Introducing a hierarchical structure to our model allows us to obtain additional information by pooling the data from each year group together. To do this, we introduce a new parameter μ_β , which we call the hyperparameter, using this parameter to estimate the posterior distributions of β for each year group individually. We store these group-level parameters in a vector, $\underline{\beta} = \{\beta(1977), \dots, \beta(2021)\}$. Our joint prior distribution for our hierarchical model, incorporating new hyperparameter μ_β , is

$$P(\mu_\beta, \underline{\beta}) = P(\mu_\beta)P(\underline{\beta}|\mu_\beta),$$

and the corresponding joint posterior distribution, which depends on all of our data, denoted $\underline{p}$, is

$$P(\mu_\beta, \beta | \underline{p}) \propto P(\mu_\beta, \underline{\beta}) P(\underline{p} | \mu_\beta, \underline{\beta}).$$

3877 Formulation of the detection change model is analogous, as are the formulations for the
3878 models examining individual cancers.

3879 We need to select prior distributions to feed into our model that will allow sufficient
3880 exploration of the parameter space. Our previous best-fit estimations of the parameter value
3881 β and Δ will serve as a good indication of where we should place the mean of weakly
3882 informative prior for the hyperparameter, allowing us to infer the group-level parameters we
3883 need. We denote the means of these priors $\bar{\beta}$ and $\bar{\Delta}$ respectively. Our hyperparameter μ_β
3884 will require a variance, τ , and we will require a hierarchical parameter to infer the group-
3885 level variances, $\bar{\sigma}$ for each replicate as well. This parameter will be denoted μ_σ and will have
3886 variance ξ . We place relatively wide, weakly informative priors on these variances.

3887 When implementing a non-centered hierarchical model, rather than sampling the group-
3888 level values directly from the hyperparameters, we recover them using a scaling and a trans-
3889 lation with a latent Gaussian variable. This does not change the interpretation of the model
3890 – rather, it is a simple reparameterisation of the model which improves model geometry and
3891 helps the sampler to converge on the posterior distributions without issue. We need one each
3892 for the recovery of β and σ values, which we denote $\underline{\theta}_\beta$ and $\underline{\theta}_\sigma$. Collectively, we use the
3893 following priors;

$$\mu_\beta \sim N^+(0, 5),$$

$$\tau \sim N^+(0, 5),$$

$$\mu_\sigma \sim N^+(0, 5),$$

$$\xi \sim N^+(0, 5),$$

$$\underline{\theta}_\beta \sim N(0, 1),$$

$$\underline{\theta}_\sigma \sim N(0, 1),$$

3894 and we recover the group level parameters as follows:

$$\underline{\beta} = \mu_{\beta} + \tau \underline{\theta}_{\beta},$$

$$\underline{\sigma} = \mu_{\sigma} + \xi \underline{\theta}_{\sigma},$$

3895 which imply, respectively, that $\underline{\beta} \sim N(\mu_{\beta}, \tau)$ and $\underline{\sigma} \sim N(\mu_{\sigma}, \xi)$. We then test whether our
3896 sampled values fit the model, where

$$p(t, n) \sim N(1 - \exp(-(\beta(n)t)^k), \sigma(n))$$

3897 tests the fit of the model where $\beta(n)$ varies in time. The formulation of the model with
3898 change in detection time is analogous, save for a difference in the step in which we test the
3899 model against our sampled parameters.

$$p(t, n) \sim N(1 - \exp(-(\beta(t + \Delta_t(n)))^k), \sigma(n))$$

3900 tests the fit of the detection lag model.

3901 5.4 Results

3902 5.4.1 SEER Trends with Age and Initial Analysis

3903 In figure 5.2, we see the results of the linear regression for the pan-cancer SEER Cumulative
3904 Cancer incidence data for age $t = 50$. The results here show an increase in cumulative cancer
3905 incidence for $t = 50$ of $3.049 \times 10^{-4} \%$ per year, which differs significantly from the null
3906 hypothesis of a zero slope ($p = 1.97 \times 10^{-8}$). We therefore conclude that there has indeed
3907 been a significant measurable increase in diagnosed cancer incidence over time for this age
3908 group.

3909 5.4.2 Pan-cancer Model Investigations

3910 We tested both hypotheses for all 16 pan-cancer SEER datasets using our hierarchical Bayesian
3911 framework. We test the true rate change hypothesis by fitting the data to equation 5.1, and the
3912 detection change hypothesis by using 5.2. Inference for both the rate and detection change




FIGURE 5.1: Incidence data for All Cancers from 1977-2021.




FIGURE 5.2: Results of linear regression line of best fit for the pan-cancer SEER Cumulative Cancer incidence data for age $t = 50$. The blue points are the data points for each year group; the black line shows the linear regression line, with the shaded regions surrounding this line showing the 95% confidence interval.

model was successful. We can see in figure 5.3 that our inference converged to a clear posterior distribution for every year group for both β_n and Δ_n . To test how well the model fits the data, we drew 100 random β parameters from the posterior distributions and then plotted the resulting model with the randomly drawn values. Parameters inferred for every year group are shown, and the model fit for 2021 is presented to demonstrate how the model fits the data. Fits for every year group can be found in appendix C.

5.4.3 Testing our hypotheses on specific cancers

The results shown in figure 5.3 demonstrate that increasing cancer incidence rates could feasibly be explained by either a true rate increase or detection improvement. We go on to test the same hypotheses on four different cancer types – breast cancer, colorectal cancer, stomach cancer and pancreatic cancer. Regular screening for breast and colorectal cancers has been commonplace for some time (262; 263), so we might expect that improved detection times would be a more likely cause for a supposed rate increase over time for under 50s.




FIGURE 5.3: Results of Bayesian modelling of the pan-cancer SEER data. **(Top Left)** Inferred group-level posterior distributions for the rate model, with the posterior distribution of the hyperparameter, denoted μ_β . **(Bottom Left)** The same for the detection improvement model. **(Top Right)** Model data for 2021 (black dots) compared to rate model fit, which was calculated using parameters drawn randomly from the posterior distribution (red lines). **(Bottom Right)** The same for the detection improvement model. For the rate model, we fix $k = 4.06$. For the detection model, we additionally fix $\beta = 0.094$.




FIGURE 5.4: Incidence data for Breast Cancer from 1977-2021.




FIGURE 5.5: Incidence data for Colorectal Cancer from 1977-2021.




FIGURE 5.6: Incidence data for Pancreatic Cancer from 1977-2021.




FIGURE 5.7: Incidence data for Stomach Cancer from 1977-2021.

Conversely, stomach and pancreatic cancers are not screened for, so we would expect that a true rate increase here is the more likely explanation for the observed rate increases.

For these data sets, we performed an initial exploratory data analysis to confirm that we saw a incidence rate increase for these cancers in the $t = 50$ age cohort, and then fixed the model parameters as we did for the pan cancer data before carrying out our Bayesian modelling. The datasets shown in figures 5.4, 5.5, 5.6 and 5.7 display the trends in cancer incidence with age from 1977 to 2021 for ages 30-80. We can observe that cancer incidence in the breast, colorectal and pancreatic cancer datasets is generally higher for more recent years than it was in 1977 for the $t = 50$ age cohort. This trend is not observed in the stomach cancer data, but we fit our models to it regardless, to investigate whether or hypothesis that either a true decrease in incidence rate or poor detection can explain the trends in this data.

What these datasets demonstrate is that unlike the pan-cancer data, the shape of the curves for each year group changes across the years for individual cancers, and that the nature of this change varies depending on cancer type. We find that with the breast cancer data, rates are similar for younger ages, but as the years progress, rates increase for older ages. With pancreatic cancer, we find that there is a notable increase in incidence rate for ages 30-40 in more recent years, but only a marginal increase in incidence for older ages. We find that colorectal and stomach cancers exhibit similar trends. Incidence rates for younger people increase across years, and decrease for older age groups – we note that these changes are more pronounced for the colorectal cancers, whereas they are only marginal for stomach cancers. We now attempt to fit models using equation 5.1 to test the rate change hypothesis and equation 5.2 to test the detection change hypothesis.

The results of the breast cancer modelling are shown in figure 5.8. Internal checks in *Stan* all passed, so we can be confident that the inference was successful, and our model does converge to the true posterior (this was the case for the remaining results as well). We can see, however, that neither model is a particularly good fit for the data. In the detection model, the model fits for ages 50-80, but not for earlier ages; the rate model does not fit the data at all. We also observe that unlike with the pan-cancer data, posteriors have wide variances, making it difficult to glean any qualitative information from these inferred distributions. We see that the 1980 and 2021 model fits are both a better fit to the 1980 data than to the 2021




FIGURE 5.8: Results of Bayesian modelling of the breast cancer SEER data. **(Top Left)** Inferred group-level posterior distributions for the rate model, with the posterior distribution of the hyperparameter, denoted μ_β . **(Bottom Left)** The same for the detection improvement model. **(Top Right)** Model data compared to rate model fit for 1980 (red points with magenta lines) and 2021 (blue points with cyan lines). These were calculated using parameters drawn randomly from the posterior distribution. **(Bottom Right)** The same for the detection improvement model. For the rate model, we fix $k = 2.91$. For the detection model, we additionally fix $\beta = 0.0054$.

data, meaning that the model fails to capture the observed changes in shape in the breast cancer data.

In figure 5.9, we have the modelling results for the colorectal cancer data. In this instance, both models are an adequate fit for the data. The same issue with inferred parameters observed in the breast cancer modelling is seen for the rate model, although we note that in the detection model, parameter values for the years 2016, 2019 and 2021 suggest an improvement in detection, and overall, we do see distinct differences between posterior distributions for each year group. This is not significant enough, however, to assert that one model explains our hypothesis better than the other. We observe that for both models, the model fits for both the 1980 and 2021 data are closer to the 2021 data than the 1980 data, meaning that again, our models fail to capture the changing shapes of the data as years progress.

We see the results of our Bayesian modelling on stomach cancer data in figure 5.10. Both models fit the data in this instance, and we see similar, modelling outcomes that we did in the colorectal cancer modelling – wide variances on inferred parameter posterior distributions for the rate model, but clearer distinct posterior distributions for the detection model. We see a similar phenomenon to previous fits with both 1980 and 2021 model fits being a better fit for the 2021 data than the 1980 data, failing to capture changing shapes of the data across years.

The pancreatic cancer incidence modelling results are shown in figure 5.11. Models both fit the data, and posterior distributions show that inferences converged to distinct posterior distributions. We do note, however, that the variances on these posteriors are wide. We again observe that model fits for both years are closer to the 2021 data than the 1980 data, again failing to capture the changing shapes of the data across year replicates.

5.5 Discussion

This analysis suggests that while in American data there seems to be a notable increase in early onset cancers since 1977, this data may be equally explained by either a true increase in incidence under 50, by a true improvement in detection time for cancer or, more likely, by some combination of both mechanisms. While there has been much speculation the potential reasons for an apparently cancer rate increase in younger people, it is also possible that the




FIGURE 5.9: Results of Bayesian modelling of the colorectal cancer SEER data. **(Top Left)** Inferred group-level posterior distributions for the rate model, with the posterior distribution of the hyperparameter, denoted μ_β . **(Bottom Left)** The same for the detection improvement model. **(Top Right)** Model data compared to rate model fit for 1980 (red points with magenta lines) and 2021 (blue points with cyan lines). These were calculated using parameters drawn randomly from the posterior distribution. **(Bottom Right)** The same for the detection improvement model. For the rate model, we fix $k = 4.15$. For the detection model, we additionally fix $\beta = 0.0063$.




FIGURE 5.10: Results of Bayesian modelling of the stomach cancer SEER data. **(Top Left)** Inferred group-level posterior distributions for the rate model, with the posterior distribution of the hyperparameter, denoted μ_β . **(Bottom Left)** The same for the detection improvement model. **(Top Right)** Model data compared to rate model fit for 1980 (red points with magenta lines) and 2021 (blue points with cyan lines). These were calculated using parameters drawn randomly from the posterior distribution. **(Bottom Right)** The same for the detection improvement model. For the rate model, we fix $k = 4.21$. For the detection model, we additionally fix $\beta = 0.0040$.




FIGURE 5.11: Results of Bayesian modelling of the pancreatic cancer SEER data. **(Top Left)** Inferred group-level posterior distributions for the rate model, with the posterior distribution of the hyperparameter, denoted μ_{β} . **(Bottom Left)** The same for the detection improvement model. **(Top Right)** Model data compared to rate model fit for 1980 (red points with magenta lines) and 2021 (blue points with cyan lines). These were calculated using parameters drawn randomly from the posterior distribution. **(Bottom Right)** The same for the detection improvement model. For the rate model, we fix $k = 4.56$. For the detection model, we additionally fix $\beta = 0.0043$.

true reason for this ostensible rise is improved detection. This analysis merely suggests that under reasonable assumptions, both models generate comparable predictions that align with current observations when investigating pan-cancer data.

During our analysis, for both the pan-cancer and individual model, we trialled a model which was intended to test both hypotheses simultaneously, with the goal of determining if such a model could indicate whether rate or detection changes had a greater influence over incidence trends if both parameters were allowed to vary. The models we trialled were identical to those seen in equation 5.2, but rather than fixing β , we allow it to vary during inference. The issue with these models is our interpretation of the results they provide. In the original detection models, the parameter Δ represents a change in detection time relative to a fixed time point: this fixed time point, in the model, is represented by a fixed rate parameter of β for the year 1977. The detection parameter, then, acts as a proxy for the varying rates we observe across year groups time. If we were to allow both β and Δ to vary, then we can no longer interpret Δ as a changing detection time relative to a fixed time point (represented by a previously fixed rate parameter for β). We intend, in future, to develop a better model which allows both parameters to vary without affecting what parameters represent.

We cannot confidently say that the findings we discussed above hold for individual cancers. Although models did fit the data in several cases, they consistently failed to capture the changing shape of difference cancer incidence trends which we observed in (figures 5.4, 5.5, 5.7 and 5.6). It may be that fixing the shape parameter, k , is not as sensible of a choice as it was in the case of the pan-cancer data, where data shape was relatively consistent across years. Allowing this parameter to vary would, however, make comparison between years and models more difficult.

We note that the posterior distributions for the Δ parameter in the stomach and colorectal cancer modelling were negative. We posit that the reason for this is that the data for these cancers exhibits a different pattern to the other individual cancers in terms of how trends change across the years. For both cancers, as years progressed, incidence rates for younger people increased, while incidence rates for older people decreased – although we note that for the $t = 50$ age cohort, we still see an increase in cancer rates in recent years for these cases. For the pan-cancer modelling, positive values that were inferred reflected that increasing rates

relative to 1977 could be explained by a positive value for Δ . Likewise, the decreasing rates relative to 1977 for the stomach and colorectal cancer data sets can be explained by a negative value for Δ .

Failing to infer meaningful parameter distributions for individual cancers does not necessarily invalidate our findings for the pan-cancer data, and our hypothesis that either a true rate increase or an improved detection can both explain the trends in pan-cancer incidence data was still shown to be true. Failing to match similar models to individual cancers was not wholly unexpected. Frank's book on the dynamics of cancer (255) examines incidence data on an even wider variety of cancers than we investigate here. Although there are some general consistencies, such as an increase in incidence with age, different cancers exhibit different incidence trends, and trying to capture the behaviours of different cancers with a one-size-fits-all model may prove difficult. It may be prudent to have ready an array of different models, each with a variant that examines both rate and detection changes, so that we can choose appropriate models to fit the incidence data for each individual cancer. This would potentially allow us to more accurately examine the phenomena of interest. Given the time and resources, we would perform these analyses on a larger range of individual cancer incidence data sets – as many as we could gain access to. Testing four individual ones is useful for validation, but a comprehensive analysis of a wider range of cancer types would allow us to see whether our findings for the pan-cancer data are an artefact of compiling all of these data sets together, or if these findings are indicative of a wider trend that can be seen by observing the qualitative trends of individual cancers.

We found that our models showed that either a true increase in rate or improved detection time could explain the trends in pan-cancer incidence data, but that these models generally failed to validate either hypothesis for individual cancers. We postulate that this is down to the fact that the Weibull models we use to test our hypotheses are too simplistic to explain trends in individual cancers, since they cannot account for specific factors which might drive variations in incidence rates over time. Colorectal and breast cancers, for example, can occur more frequently in obese and overweight people (?), and it has been found that since 1990, obesity rates worldwide have more than doubled, with adolescent obesity quadrupling in the same timeframe (?). Although we do not investigate cervical cancer in this study, it is

another useful example of a cancer whose incidence rates exhibit more complex dynamics over time. The human papillomavirus has been linked to 83% of cervical cancer mortalities (?), but a vaccine was made available in the mid 2000s – this would likely have a massive impact on incidence rates, which should be observable. For the pan-cancer data, factors relating to specific cancers will have a minimum impact, since the confounding factors which influence these individual cancers are all aggregated. For individual cancers, however, these effects will be much more prevalent – as such, trying to fit simple models which do not account for these effects will be less successful. Developing models which account for these factors for each individual cancer would be a monumental task, but it would be useful to investigate models which examine factors that drive incidence rates in more common cancers. For the examples above, we might add a component to our models (or a similar model) which accounts for incidences of breast and colorectal cancers linked to obesity, and observe how much this effects incidence over time. In the case of cervical cancers, we might add a component to the model (perhaps some sigmoidal function) which accounts for the introduction of the HPV vaccine – it might do nothing until this vaccine is made widely available, and then curb the incidence rates in our in accordance with the vaccine’s effects on incidence rates over time.

The motivating question of what is behind the observed rise in cancer rates in younger people remains unclear. While this analysis suggests two alternative reasons for what is observed, ascertaining which is correct (or if either have merit) is a far from trivial question to answer. Answering such questions would take dedicated epidemiological investigations and modelling approaches to fully elucidate, but this work suggests that an ostensible rise in cancer incidence might be an artefact not of increasing disease, but of better detection. If this were the case, this seemingly negative trend in cancer incidence rates would in reality be a testament of a more positive story. There are many instances of precisely this in cancer research - the oft-misunderstood fact that cumulative lifetime cancer incidence has risen for decades appears largely due to the fact that cancer is primarily a disease of ageing, the increased cancer burden a positive sign of a longer-living population (264; 265).

This is an important aspect to consider, as too are the significant limitations of this work. The simple Weibull model was chosen as it is basic and fits the general trends of lifetime cancer incidence well, but its simplicity also means we need be cautious of over-inference and

4075 over-extrapolation. There is also the possibility that even if generally suitable, its parameters
4076 may vary in non-mutually exclusive ways. Nor does the model considered here factor in other
4077 potential processes or competing factors. These are all weaknesses that mean its predictions
4078 need to be treated as phenomenological at best, rather than mechanistic (266), and interpreted
4079 cautiously. Nevertheless, it is useful at exploring the dynamics of cancer incidence under
4080 different assumptions.

Chapter 6

Conclusion

Our investigations into adaptive therapy found that containment strategies can be superior to standard-of-care MTD protocols in terms of delaying onset of resistance. We analysed the results of a variety of different experiments in a number of different ways, but perhaps most importantly with our Bayesian hierarchical modelling workflow. Obtaining posterior distributions for key biological model parameters, rather than point estimates, allows us to make predictions which are more likely to account for potential variability in experimental outcomes. We used our workflow to calibrate a mathematical model which describes competition between cell lines. By using this model calibrated with *in vivo* murine data, we were able to trial different treatment protocols and compare them to existing ones. In doing so, we validated the adaptive therapy hypothesis and provided a theoretical basis that will inform both clinical and pre-clinical trials of these new and improved treatment strategies.

We developed a mathematical model of Bipolar androgen therapy, which has never previously been done. We developed simplified versions of this model which we could study analytically, which allowed us to develop new alternatives to current treatment protocols which we predict will drastically improve treatment outcomes when compared both to current standard-of-care androgen deprivation therapy and existing bipolar androgen therapy protocols. These model results generalise to similar treatment models of alternating therapies.

We also demonstrated the versatility of our Bayesian modelling workflow by using it to try and explain trends in cancer incidence data. We found that increasing incidence rates across all cancers could be explained either by a true increase in rate since 1977, or an improvement in cancer detection time relative to 1977. We used the same model workflow to

test these same hypotheses on individual cancers. While this result may hold across all cancers, the model was found to be too simple for individual cancers where rate dynamics were more variable across the years.

Altogether, we demonstrated that Bayesian hierarchical models can be powerful tools which have multiple uses when studying problems in the wider sphere of mathematical oncology. We demonstrated that we can use this workflow to quickly and efficiently analyse large sets of data, capturing variability between replicates and producing posterior distributions of interpretable parameters that describe cell dynamics and predict treatment outcomes. This same framework can be used to test hypotheses for large datasets by applying simple mathematical models whose parameters have a clear interpretation. We utilised mathematical models to proposed treatment strategies which improve on current standards of care. The results obtained from both analytical and Bayesian analysis of these models can now be used to inform both clinical and pre-clinical trials, both *in vitro* and *in vivo*.

Bibliography

- [1] HE Skipper, FM Schabel Jr, LB Mellett, JA Montgomery, LJ Wilkoff, HH Lloyd, and RW Brockman. Implications of biochemical, cytokinetic, pharmacologic and toxicologic relationships in the design of optimal therapeutic schedules. *Cancer Chemother. Rep.*, 54:431–450, 1970.
- [2] Philip Savage, Justin Stebbing, Mark Bower, and Tim Crook. Why does cytotoxic chemotherapy cure only some cancers? *Nature Clinical Practice Oncology*, 6(1):43–52, 2009.
- [3] Robert A Gatenby, Ariosto S Silva, Robert J Gillies, and B Roy Frieden. Adaptive therapy. *Cancer research*, 69(11):4894–4903, 2009.
- [4] RC Axtell and JJ Arends. Ecology and management of arthropod pests of poultry. *Annual review of entomology*, 35(1):101–126, 1990.
- [5] Michael Renton, Roberto Busi, Paul Neve, David Thornby, and Martin Vila-Aiub. Herbicide resistance modelling: past, present and future. *Pest management science*, 70(9):1394–1404, 2014.
- [6] M Campbell. Genotype vs phenotype: Examples and definitions. <https://www.technologynetworks.com/genomics/articles/genotype-vs-phenotype-examples-and-definitions-318446>.
- [7] Charles Darwin. On the origin of species, 1859, 2016.
- [8] Joanna Masel. Genetic drift. *Current Biology*, 21(20):R837–R838, 2011.
- [9] Montgomery Slatkin. Gene flow in natural populations. *Annual review of ecology and systematics*, 16(1):393–430, 1985.

- 4140 [10] S. Malmquist and K. Prescott. *Human Biology*. University of Minnesota, 2nd edition,
4141 2021.
- 4142 [11] Susan Elmore. Apoptosis: a review of programmed cell death. *Toxicologic pathology*,
4143 35(4):495–516, 2007.
- 4144 [12] Weijun Huang, LaTonya J. Hickson, Alfonso Eirin, James L. Kirkland, and Lilach O.
4145 Lerman. Cellular senescence: the good, the bad and the unknown. *Nature Reviews*
4146 *Nephrology*, 18(10):611–627, 2022.
- 4147 [13] Frances R Balkwill, Melania Capasso, and Thorsten Hagemann. The tumor microen-
4148 vironment at a glance. *Journal of cell science*, 125(23):5591–5596, 2012.
- 4149 [14] Robert A Gatenby. A change of strategy in the war on cancer. *Nature*, 459(7246):508–
4150 509, 2009.
- 4151 [15] Paolo Di Giamberardino, Daniela Iacoviello, and Muhammad Zubair. Optimal therapy
4152 for hiv infection containment and virions inhibition. *Infectious Disease Modelling*,
4153 9(1):10–26, 2024.
- 4154 [16] German Advisory Committee Blood et al. Human immunodeficiency virus (hiv).
4155 *Transfusion Medicine and Hemotherapy*, 43(3):203, 2016.
- 4156 [17] P. Andrew Futreal, Lachlan Coin, Mhairi Marshall, Thomas Down, Timothy Hubbard,
4157 Richard Wooster, Nazneen Rahman, and Michael R. Stratton. A census of human
4158 cancer genes. *Nature Reviews Cancer*, 4(3):177–183, 2004.
- 4159 [18] Ivo P Touw and Stefan J Erkeland. Retroviral insertion mutagenesis in mice as a
4160 comparative oncogenomics tool to identify disease genes in human leukemia. *Mol*
4161 *Ther*, 15(1):13–19, Jan 2007.
- 4162 [19] Michael R. Stratton, Peter J. Campbell, and P. Andrew Futreal. The cancer genome.
4163 *Nature*, 458(7239):719–724, 2009.
- 4164 [20] Pedro M Enriquez-Navas, Jonathan W Wojtkowiak, and Robert A Gatenby. Applica-
4165 tion of evolutionary principles to cancer therapy. *Cancer research*, 75(22):4675–4680,
4166 2015.

- 4167 [21] Jeffrey West, Fred Adler, Jill Gallaher, Maximilian Strobl, Renee Brady-Nicholls, Joel
4168 Brown, Mark Roberson-Tessi, Eunjung Kim, Robert Noble, Yannick Viossat, et al.
4169 A survey of open questions in adaptive therapy: Bridging mathematics and clinical
4170 translation. *Elife*, 12:e84263, 2023.
- 4171 [22] Ivana Bozic, Johannes G Reiter, Benjamin Allen, Tibor Antal, Krishnendu Chatterjee,
4172 Preya Shah, Yo Sup Moon, Amin Yaqubie, Nicole Kelly, Dung T Le, et al. Evolution-
4173 ary dynamics of cancer in response to targeted combination therapy. *elife*, 2:e00747,
4174 2013.
- 4175 [23] Xuan Wang, Haiyun Zhang, and Xiaozhuo Chen. Drug resistance and combating drug
4176 resistance in cancer. *Cancer Drug Resistance*, 2(2):141, 2019.
- 4177 [24] Yannick Viossat and Robert Noble. A theoretical analysis of tumour containment.
4178 *Nature ecology & evolution*, 5(6):826–835, 2021.
- 4179 [25] J Schneider, M Bak, TH Efferth, M Kaufmann, J Mattern, and M Volm. P-glycoprotein
4180 expression in treated and untreated human breast cancer. *British journal of cancer*,
4181 60(6):815–818, 1989.
- 4182 [26] Philip Gerlee. The model muddle: in search of tumor growth laws. *Cancer research*,
4183 73(8):2407–2411, 2013.
- 4184 [27] A K LAIRD. Dynamics of tumor growth. *Br J Cancer*, 13(3):490–502, Sep 1964.
- 4185 [28] Tom E Wheldon. Mathematical models in cancer research. 1988.
- 4186 [29] Ludwig Von Bertalanffy. Problems of organic growth. *Nature*, 163(4135):156–158,
4187 1949.
- 4188 [30] Alfred J Lotka. Contribution to the theory of periodic reactions. *The Journal of*
4189 *Physical Chemistry*, 14(3):271–274, 2002.
- 4190 [31] Vito Volterra. *Variazioni e fluttuazioni del numero d'individui in specie animali con-*
4191 *viventi*, volume 2. Società anonima tipografica" Leonardo da Vinci", 1927.
- 4192 [32] Vito Volterra. Variations and fluctuations of the number of individuals in animal
4193 species living together. *Animal ecology*, pages 412–433, 1931.

- 4194 [33] Archibald Vivian Hill. The possible effects of the aggregation of the molecules of
4195 hemoglobin on its dissociation curves. *j. physiol.*, 40:iv–vii, 1910.
- 4196 [34] Leonor Michaelis, Maud L Menten, et al. Die kinetik der invertinwirkung. *Biochem.*
4197 *z.*, 49(333-369):352, 1913.
- 4198 [35] RB Martin, ME Fisher, RF Minchin, and KL Teo. Optimal control of tumor size used
4199 to maximize survival time when cells are resistant to chemotherapy. *Mathematical*
4200 *biosciences*, 110(2):201–219, 1992.
- 4201 [36] R B Martin, M E Fisher, R F Minchin, and K L Teo. Low-intensity combination
4202 chemotherapy maximizes host survival time for tumors containing drug-resistant cells.
4203 *Math Biosci*, 110(2):221–252, Jul 1992.
- 4204 [37] Katarina Bacevic, Robert Noble, Ahmed Soffar, Orchid Wael Ammar, Benjamin
4205 Boszonyik, Susana Prieto, Charles Vincent, Michael E Hochberg, Liliana Krasinska,
4206 and Daniel Fisher. Spatial competition constrains resistance to targeted cancer therapy.
4207 *Nature communications*, 8(1):1995, 2017.
- 4208 [38] Jill A Gallaher, Pedro M Enriquez-Navas, Kimberly A Luddy, Robert A Gatenby, and
4209 Alexander RA Anderson. Spatial heterogeneity and evolutionary dynamics modulate
4210 time to recurrence in continuous and adaptive cancer therapies. *Cancer research*,
4211 78(8):2127–2139, 2018.
- 4212 [39] Ariosto S Silva and Robert A Gatenby. A theoretical quantitative model for evolution
4213 of cancer chemotherapy resistance. *Biology direct*, 5:1–17, 2010.
- 4214 [40] Ariosto S Silva, Yoonseok Kam, Zayar P Khin, Susan E Minton, Robert J Gillies, and
4215 Robert A Gatenby. Evolutionary approaches to prolong progression-free survival in
4216 breast cancer. *Cancer research*, 72(24):6362–6370, 2012.
- 4217 [41] Jingsong Zhang, Jessica J Cunningham, Joel S Brown, and Robert A Gatenby. Inte-
4218 grating evolutionary dynamics into treatment of metastatic castrate-resistant prostate
4219 cancer. *Nature communications*, 8(1):1816, 2017.

- [42] Inna Smalley, Eunjung Kim, Jiannong Li, Paige Spence, Clayton J Wyatt, Zeynep Eroglu, Vernon K Sondak, Jane L Messina, Nalan Akgul Babacan, Silvy Stuchi Maria-Engler, et al. Leveraging transcriptional dynamics to improve braf inhibitor responses in melanoma. *EBioMedicine*, 48:178–190, 2019.
- [43] Renee Brady-Nicholls, Jingsong Zhang, Tian Zhang, Andrew Z Wang, Robert Butler, Robert A Gatenby, and Heiko Enderling. Predicting patient-specific response to adaptive therapy in metastatic castration-resistant prostate cancer using prostate-specific antigen dynamics. *Neoplasia*, 23(9):851–858, 2021.
- [44] Maximilian AR Strobl, Jeffrey West, Yannick Viossat, Mehdi Damaghi, Mark Robertson-Tessi, Joel S Brown, Robert A Gatenby, Philip K Maini, and Alexander RA Anderson. Turnover modulates the need for a cost of resistance in adaptive therapy. *Cancer research*, 81(4):1135–1147, 2021.
- [45] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1), 2017.
- [46] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.
- [47] Charles J Geyer. Introduction to markov chain monte carlo. *Handbook of markov chain monte carlo*, 20116022(45):22, 2011.
- [48] Galin L Jones. On the markov chain central limit theorem. 2004.
- [49] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- [50] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970.
- [51] JD Cook. Effective sample size for mcmc. <https://www.johndcook.com/blog/2017/06/27/effective-sample-size-for-mcmc/>.

- [52] ERIC Shang. Introduction to markov chains and markov chain mixing. *IEEE Access*, 6:9700–9710, 2018.
- [53] Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.
- [54] Radford M Neal et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- [55] Matthew D Hoffman, Andrew Gelman, et al. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, 2014.
- [56] Stan Development Team. Stan modeling language user’s guide and reference manual (version 2.21.2). <https://mc-stan.org>, 2022.
- [57] Michael J Betancourt. Generalizing the no-u-turn sampler to riemannian manifolds. *arXiv preprint arXiv:1304.1920*, 2013.
- [58] Blair Colyer, Maciej Bak, David Basanta, and Robert Noble. A seven-step guide to spatial, agent-based modelling of tumour evolution. *Evolutionary Applications*, 17(5):e13687, 2024. e13687 EVA-2024-024-RS.R1.
- [59] Matias Casás-Selves and James DeGregori. How cancer shapes evolution and how evolution shapes cancer. *Evolution: Education and outreach*, 4:624–634, 2011.
- [60] Lauren MF Merlo, John W Pepper, Brian J Reid, and Carlo C Maley. Cancer as an evolutionary and ecological process. *Nature reviews cancer*, 6(12):924–935, 2006.
- [61] Jing Yang, Tong-Jun Zhao, Chang-Qing Yuan, Jing-Hui Xie, and Fang-Fang Hao. A nonlinear competitive model of the prostate tumor growth under intermittent androgen suppression. *Journal of theoretical biology*, 404:66–72, 2016.
- [62] David Basanta, Jacob G Scott, Russ Rockne, Kristin R Swanson, and Alexander RA Anderson. The role of idh1 mutated tumour cells in secondary glioblastomas: an evolutionary game theoretical view. *Physical biology*, 8(1):015016, 2011.

- 4273 [63] Kaveh Danesh, Rick Durrett, Laura J Havrilesky, and Evan Myers. A branching pro-
4274 cess model of ovarian cancer. *Journal of theoretical biology*, 314:10–15, 2012.
- 4275 [64] Rick Durrett, Jasmine Foo, Kevin Leder, John Mayberry, and Franziska Michor. Evo-
4276 lutionary dynamics of tumor progression with random fitness values. *Theoretical pop-*
4277 *ulation biology*, 78(1):54–66, 2010.
- 4278 [65] Jeffrey West, Zaki Hasnain, Paul Macklin, and Paul K Newton. An evolutionary
4279 model of tumor cell kinetics and the emergence of molecular heterogeneity driving
4280 gompertzian growth. *SIAM review*, 58(4):716–736, 2016.
- 4281 [66] Richard Durrett, Jasmine Foo, and Kevin Leder. Spatial moran models, ii: cancer
4282 initiation in spatially structured tissue. *Journal of mathematical biology*, 72:1369–
4283 1400, 2016.
- 4284 [67] Serhii Aif, Nico Appold, Lucas Kampman, Oskar Hallatschek, and Jona Kayser. Evo-
4285 lutionary rescue of resistant mutants is governed by a balance between radial ex-
4286 pansion and selection in compact populations. *Nature Communications*, 13(1):7916,
4287 2022.
- 4288 [68] Daniel K Saha, Alexander RA Anderson, Luis Cisneros, and Carlo C Maley. In silico
4289 investigations of adaptive therapy using a single cytotoxic or a single cytostatic drug.
4290 *bioRxiv*, pages 2023–05, 2023.
- 4291 [69] Joshua A Bull and Helen M Byrne. Quantification of spatial and phenotypic hetero-
4292 geneity in an agent-based model of tumour-macrophage interactions. *PLOS Compu-*
4293 *tational Biology*, 19(3):e1010994, 2023.
- 4294 [70] Robert Noble, Dominik Burri, Cécile Le Sueur, Jeanne Lemant, Yannick Viossat,
4295 Jakob Nikolas Kather, and Niko Beerenwinkel. Spatial structure governs the mode
4296 of tumour evolution. *Nature ecology & evolution*, 6(2):207–217, 2022.
- 4297 [71] Robert Noble, John T Burley, Cécile Le Sueur, and Michael E Hochberg. When,
4298 why and how tumour clonal diversity predicts survival. *Evolutionary applications*,
4299 13(7):1558–1568, 2020.

- [72] Maciej Bak, Blair Colyer, Veselin Manojlović, and Robert Noble. Warlock: an automated computational workflow for simulating spatially structured tumour evolution. *arXiv preprint arXiv:2301.07808*, 2023.
- [73] Maximilian AR Strobl, Jill Gallaher, Jeffrey West, Mark Robertson-Tessi, Philip K Maini, and Alexander RA Anderson. Spatial structure impacts adaptive therapy by shaping intra-tumoral competition. *Communications medicine*, 2(1):46, 2022.
- [74] Joel L Schiff. *Cellular automata: a discrete view of the world*. John Wiley & Sons, 2011.
- [75] Pierre-Yves Louis and Francesca R Nardi. Probabilistic cellular automata. *Emergence, complexity, computation*, 27, 2018.
- [76] Murray Eden. A two-dimensional growth process. *Dynamics of fractal surfaces*, 4:223–239, 1961.
- [77] R Jullien and R Botet. Scaling properties of the surface of the eden model in $d=2, 3, 4$. *Journal of Physics A: Mathematical and general*, 18(12):2279, 1985.
- [78] Haider Tari, Ketty Kessler, Nick Trahearn, Benjamin Werner, Maria Vinci, Chris Jones, and Andrea Sottoriva. Quantification of spatial subclonal interactions enhancing the invasive phenotype of pediatric glioma. *Cell Reports*, 40(9), 2022.
- [79] Sara Hamis, James Yates, Mark AJ Chaplain, and Gibin G Powathil. Targeting cellular dna damage responses in cancer: an in vitro-calibrated agent-based model simulating monolayer and spheroid treatment responses to atr-inhibiting drugs. *Bulletin of Mathematical Biology*, 83:1–21, 2021.
- [80] Bartłomiej Waclaw, Ivana Bozic, Meredith E Pittman, Ralph H Hruban, Bert Vogelstein, and Martin A Nowak. A spatial model predicts that dispersal and cell turnover limit intratumour heterogeneity. *Nature*, 525(7568):261–264, 2015.
- [81] Nick Jagiella, Benedikt Müller, Margareta Müller, Irene E Vignon-Clementel, and Dirk Drasdo. Inferring growth control mechanisms in growing multi-cellular spheroids

- 4326 of nslc cells from spatial-temporal image data. *PLoS computational biology*,
4327 12(2):e1004412, 2016.
- 4328 [82] Ketevan Chkhaidze, Timon Heide, Benjamin Werner, Marc J Williams, Weini Huang,
4329 Giulio Caravagna, Trevor A Graham, and Andrea Sottoriva. Spatially constrained tu-
4330 mour growth affects the patterns of clonal selection and neutral drift in cancer genomic
4331 data. *PLoS computational biology*, 15(7):e1007243, 2019.
- 4332 [83] Richard Durrett. Random graph dynamics, 2007.
- 4333 [84] Trevor Williams and Rolf Bjerknes. Stochastic model for abnormal clone spread
4334 through epithelial basal layer. *Nature*, 236(5340):19–21, 1972.
- 4335 [85] François Graner and James A Glazier. Simulation of biological cell sorting using a
4336 two-dimensional extended potts model. *Physical review letters*, 69(13):2013, 1992.
- 4337 [86] Nicholas J Savill and Paulien Hogeweg. Modelling morphogenesis: from single cells
4338 to crawling slugs. *Journal of theoretical biology*, 184(3):229–235, 1997.
- 4339 [87] András Szabó and Roeland MH Merks. Cellular potts modeling of tumor growth,
4340 tumor invasion, and tumor evolution. *Frontiers in oncology*, 3:87, 2013.
- 4341 [88] Jonathan F Li and John Lowengrub. The effects of cell compressibility, motility and
4342 contact inhibition on the growth of tumor cell clusters using the cellular potts model.
4343 *Journal of theoretical biology*, 343:79–91, 2014.
- 4344 [89] Andreas Deutsch and Sabine Dormann. *Mathematical modeling of biological pattern*
4345 *formation*. Springer, 2005.
- 4346 [90] Andreas Deutsch, Josué Manik Nava-Sedeño, Simon Syga, and Haralampos
4347 Hatzikirou. Bio-igca: a cellular automaton modelling class for analysing collective
4348 cell migration. *PLoS computational biology*, 17(6):e1009066, 2021.
- 4349 [91] Pahini Pandya, Jose L Orgaz, and Victoria Sanz-Moreno. Modes of invasion during
4350 tumour dissemination. *Molecular oncology*, 11(1):5–27, 2017.

- [92] Alessandro Lugli, Inti Zlobec, Martin D Berger, Richard Kirsch, and Iris D Nagtegaal. Tumour budding in solid cancers. *Nature Reviews Clinical Oncology*, 18(2):101–115, 2021.
- [93] Sean L Preston, Wai-Man Wong, Annie On-On Chan, Richard Poulson, Rosemary Jeffery, Robert A Goodlad, Nikki Mandir, George Elia, Marco Novelli, Walter F Bodmer, et al. Bottom-up histogenesis of colorectal adenomas: origin in the monocryptal adenoma and initial expansion by crypt fission. *Cancer research*, 63(13):3819–3825, 2003.
- [94] Patrick Alfred Pierce Moran. Random processes in genetics, 1958.
- [95] Andrea Sottoriva, Haeyoun Kang, Zhicheng Ma, Trevor A Graham, Matthew P Salomon, Junsong Zhao, Paul Marjoram, Kimberly Siegmund, Michael F Press, Darryl Shibata, et al. A big bang model of human colorectal tumor growth. *Nature genetics*, 47(3):209–216, 2015.
- [96] Kimberly D Siegmund, Paul Marjoram, Yen-Jung Woo, Simon Tavaré, and Darryl Shibata. Inferring clonal expansion and cancer stem cell dynamics from dna methylation patterns in colorectal cancers. *Proceedings of the National Academy of Sciences*, 106(12):4828–4833, 2009.
- [97] Niko Beerenwinkel, Roland F Schwarz, Moritz Gerstung, and Florian Markowetz. Cancer evolution: mathematical models and computational inference. *Systematic biology*, 64(1):e1–e25, 2015.
- [98] Benjamin Franz, Chuan Xue, Kevin J Painter, and Radek Erban. Travelling waves in hybrid chemotaxis models. *Bulletin of mathematical biology*, 76:377–400, 2014.
- [99] Paul Macklin, Mary E Edgerton, Alastair M Thompson, and Vittorio Cristini. Patient-calibrated agent-based modelling of ductal carcinoma in situ (dcis): from microscopic measurements to macroscopic predictions of clinical progression. *Journal of theoretical biology*, 301:122–140, 2012.

- 4377 [100] Philip Gerlee and Alexander RA Anderson. Modelling evolutionary cell behaviour
4378 using neural networks: application to tumour growth. *Biosystems*, 95(2):166–174,
4379 2009.
- 4380 [101] Keith S Hoek, Ossia M Eichhoff, Natalie C Schlegel, Udo Döbbeling, Nikita Kobert,
4381 Leo Schaerer, Silvio Hemmi, and Reinhard Dummer. In vivo switching of human
4382 melanoma cells between proliferative and invasive states. *Cancer research*, 68(3):650–
4383 656, 2008.
- 4384 [102] Mark Robertson-Tessi, Robert J Gillies, Robert A Gatenby, and Alexander RA An-
4385 derson. Impact of metabolic heterogeneity on tumor growth, invasion, and treatment
4386 outcomes. *Cancer research*, 75(8):1567–1579, 2015.
- 4387 [103] Motoo Kimura. The number of heterozygous nucleotide sites maintained in a finite
4388 population due to steady flux of mutations. *Genetics*, 61(4):893, 1969.
- 4389 [104] Ryan O Schenck, Eunjung Kim, Rafael R Bravo, Jeffrey West, Simon Leedham, Dar-
4390 ryl Shibata, and Alexander RA Anderson. Homeostasis limits keratinocyte evolution.
4391 *Proceedings of the National Academy of Sciences*, 119(35):e2006487119, 2022.
- 4392 [105] Jeffrey West. Hybrid automata library. <https://halloworld.org>.
- 4393 [106] Robert Noble. An eden model in r (with optional mutation of division rates). Available
4394 from <https://github.com/robjohnnoble/EdenModel>, 2018.
- 4395 [107] Daniel T Gillespie. A general method for numerically simulating the stochastic
4396 time evolution of coupled chemical reactions. *Journal of computational physics*,
4397 22(4):403–434, 1976.
- 4398 [108] Daniel T Gillespie. Approximate accelerated stochastic simulation of chemically re-
4399 acting systems. *The Journal of chemical physics*, 115(4):1716–1733, 2001.
- 4400 [109] Takayuki Osogami and Mor Harchol-Balter. Closed form solutions for mapping
4401 general distributions to quasi-minimal ph distributions. *Performance Evaluation*,
4402 63(6):524–552, 2006.

- 4403 [110] Giulia Belluccini, Martín López-García, Grant Lythe, and Carmen Molina-París.
4404 Counting generations in birth and death processes with competing erlang and expo-
4405 nential waiting times. *Scientific Reports*, 12(1):11289, 2022.
- 4406 [111] Nicole M Anderson and M Celeste Simon. The tumor microenvironment. *Current*
4407 *Biology*, 30(16):R921–R925, 2020.
- 4408 [112] MA Chaplain and AR Anderson. Mathematical modelling, simulation and prediction
4409 of tumour-induced angiogenesis. *Invasion & metastasis*, 16(4-5):222–234, 1996.
- 4410 [113] An-Shen Qi, Xiang Zheng, Chan-Ying Du, and Bao-Sheng An. A cellular automaton
4411 model of cancerous growth. *Journal of theoretical biology*, 161(1):1–12, 1993.
- 4412 [114] Aalpen A Patel, Edward T Gawlinski, Susan K Lemieux, and Robert A Gatenby. A
4413 cellular automaton model of early tumor growth and invasion: the effects of native
4414 tissue vascularity and increased anaerobic tumor metabolism. *Journal of theoretical*
4415 *biology*, 213(3):315–331, 2001.
- 4416 [115] William H Press. *Numerical recipes 3rd edition: The art of scientific computing*.
4417 Cambridge university press, 2007.
- 4418 [116] Ahmadreza Ghaffarizadeh, Samuel H Friedman, and Paul Macklin. Biofvm: an effi-
4419 cient, parallelized diffusive transport solver for 3-d biological simulations. *Bioinfor-*
4420 *matics*, 32(8):1256–1258, 2016.
- 4421 [117] Rafael R Bravo, Etienne Baratchart, Jeffrey West, Ryan O Schenck, Anna K Miller, Jill
4422 Gallaher, Chandler D Gatenbee, David Basanta, Mark Robertson-Tessi, and Alexan-
4423 der RA Anderson. Hybrid automata library: A flexible platform for hybrid modeling
4424 with real-time visualization. *PLoS computational biology*, 16(3):e1007635, 2020.
- 4425 [118] Joseph J Casciari, Stratis V Sotirchos, and Robert M Sutherland. Mathematical mod-
4426 elling of microenvironment and growth in emt6/ro multicellular tumour spheroids.
4427 *Cell proliferation*, 25(1):1–22, 1992.

- 4428 [119] Yangjin Kim, Magdalena A Stolarska, and Hans G Othmer. A hybrid model for tumor
4429 spheroid growth in vitro i: theoretical development and early results. *Mathematical*
4430 *Models and Methods in Applied Sciences*, 17(supp01):1773–1798, 2007.
- 4431 [120] David Robert Grimes, Catherine Kelly, Katarzyna Bloch, and Mike Partridge. A
4432 method for estimating the oxygen consumption rate in multicellular tumour spheroids.
4433 *Journal of The Royal Society Interface*, 11(92):20131124, 2014.
- 4434 [121] Joseph J Casciari, Stratis V Sotirchos, and Robert M Sutherland. Variations in tumor
4435 cell growth rates and metabolism with oxygen concentration, glucose concentration,
4436 and extracellular ph. *Journal of cellular physiology*, 151(2):386–394, 1992.
- 4437 [122] David Robert Grimes, Alexander G Fletcher, and Mike Partridge. Oxygen con-
4438 sumption dynamics in steady-state tumour models. *Royal Society open science*,
4439 1(1):140080, 2014.
- 4440 [123] Arturo Araujo and David Basanta. Hybrid discrete-continuum cellular automaton
4441 (hca) model of prostate to bone metastasis. *bioRxiv*, page 043620, 2016.
- 4442 [124] Arturo Araujo, Leah M Cook, Conor C Lynch, and David Basanta. Size matters:
4443 Metastatic cluster size and stromal recruitment in the establishment of successful
4444 prostate cancer to bone metastases. *Bulletin of mathematical biology*, 80:1046–1058,
4445 2018.
- 4446 [125] Arturo Araujo, Leah M Cook, Conor C Lynch, and David Basanta. An integrated
4447 computational model of the bone microenvironment in bone-metastatic prostate can-
4448 cer. *Cancer research*, 74(9):2391–2401, 2014.
- 4449 [126] Andrea Sottoriva and Simon Tavaré. Integrating approximate bayesian computation
4450 with complex agent-based models for cancer research. In *Proceedings of COMP-*
4451 *STAT’2010: 19th International Conference on Computational Statistics Paris France,*
4452 *August 22-27, 2010 Keynote, Invited and Contributed Papers*, pages 57–66. Springer,
4453 2010.
- 4454 [127] Colin G Cess and Stacey D Finley. Calibrating agent-based models to tumor images
4455 using representation learning. *PLOS Computational Biology*, 19(4):e1011070, 2023.

- 4456 [128] Il'ya Meerovich Sobol'. On the distribution of points in a cube and the approximate
4457 evaluation of integrals. *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*,
4458 7(4):784–802, 1967.
- 4459 [129] Michael D McKay, Richard J Beckman, and William J Conover. A comparison of
4460 three methods for selecting values of input variables in the analysis of output from a
4461 computer code. *Technometrics*, 42(1):55–61, 2000.
- 4462 [130] Sara Hamis, Stanislav Stratiev, and Gibin G Powathil. Uncertainty and sensitivity
4463 analyses methods for agent-based mathematical models: An introductory review. *The*
4464 *Physics of Cancer: Research Advances*, pages 1–37, 2021.
- 4465 [131] Ahmadreza Ghaffarizadeh, Randy Heiland, Samuel H Friedman, Shannon M Mumenthaler,
4466 and Paul Macklin. Physicell: An open source physics-based cell simulator for
4467 3-d multicellular systems. *PLoS computational biology*, 14(2):e1005991, 2018.
- 4468 [132] Robert Noble. ggmmuller: Create muller plots of evolutionary dynamics. r package
4469 version 0.5.3. <https://CRAN.R-project.org/package=ggmmuller>, 2019.
- 4470 [133] Chandler D Gatenbee, Ryan O Schenck, Rafael R Bravo, and Alexander RA Anderson.
4471 Evofreq: visualization of the evolutionary frequencies of sequence and model
4472 data. *BMC bioinformatics*, 20:1–4, 2019.
- 4473 [134] Odd Erik Gundersen. The fundamental principles of reproducibility. *Philosophical*
4474 *Transactions of the Royal Society A*, 379(2197):20200210, 2021.
- 4475 [135] Geir Kjetil Sandve, Anton Nekrutenko, James Taylor, and Eivind Hovig. Ten simple
4476 rules for reproducible computational research. *PLoS computational biology*,
4477 9(10):e1003285, 2013.
- 4478 [136] Michael A Heroux and James M Willenbring. Barely sufficient software engineering:
4479 10 practices to improve your cse software. In *2009 ICSE workshop on software*
4480 *engineering for computational science and engineering*, pages 15–21. IEEE, 2009.
- 4481 [137] Zenodo. <https://zenodo.org/>, 2023.

- 4482 [138] Laura Wratten, Andreas Wilm, and Jonathan Göke. Reproducible, scalable, and share-
4483 able analysis pipelines with bioinformatics workflow managers. *Nature methods*,
4484 18(10):1161–1168, 2021.
- 4485 [139] Felix Mölder, Kim Philipp Jablonski, Brice Letcher, Michael B Hall, Christopher H
4486 Tomkins-Tinch, Vanessa Sochat, Jan Forster, Soohyun Lee, Sven O Twardziok,
4487 Alexander Kanitz, et al. Sustainable data analysis with snakemake. *F1000Research*,
4488 10, 2021.
- 4489 [140] Paolo Di Tommaso, Maria Chatzou, Evan W Floden, Pablo Prieto Barja, Emilio
4490 Palumbo, and Cedric Notredame. Nextflow enables reproducible computational work-
4491 flows. *Nature biotechnology*, 35(4):316–319, 2017.
- 4492 [141] Silas Kieser, Joseph Brown, Evgeny M Zdobnov, Mirko Trajkovski, and Lee Ann
4493 McCue. Atlas: a snakemake workflow for assembly, annotation, and genomic binning
4494 of metagenome sequence data. *BMC bioinformatics*, 21:1–8, 2020.
- 4495 [142] Martin Hölzer and Manja Marz. Poseidon: a nextflow pipeline for the detection of
4496 evolutionary recombination events and positive selection. *Bioinformatics*, 37(7):1018–
4497 1020, 2021.
- 4498 [143] Qi Zhao, Yu Sun, Dawei Wang, Hongwan Zhang, Kai Yu, Jian Zheng, and Zhixiang
4499 Zuo. Lncpipe: A nextflow-based pipeline for identification and analysis of long non-
4500 coding rnas from rna-seq data. *J Genet Genomics*, 45(7):399–401, 2018.
- 4501 [144] MacIntosh Cornwell, Mahesh Vangala, Len Taing, Zachary Herbert, Johannes Köster,
4502 Bo Li, Hanfei Sun, Taiwen Li, Jian Zhang, Xintao Qiu, et al. Viper: Visualization
4503 pipeline for rna-seq, a snakemake workflow for efficient and complete rna-seq analy-
4504 sis. *BMC bioinformatics*, 19:1–14, 2018.
- 4505 [145] David Moreau, Kristina Wiebels, and Carl Boettiger. Containers for computational
4506 reproducibility. *Nature Reviews Methods Primers*, 3(1):50, 2023.
- 4507 [146] Dirk Merkel et al. Docker: lightweight linux containers for consistent development
4508 and deployment. *Linux j*, 239(2):2, 2014.

- [147] Björn Grüning, Ryan Dale, Andreas Sjödin, Brad A Chapman, Jillian Rowe, Christopher H Tomkins-Tinch, Renan Valieris, Johannes Köster, and Bioconda Team. Bioconda: sustainable and comprehensive software distribution for the life sciences. *Nature methods*, 15(7):475–476, 2018.
- [148] Gary R Mirams, Christopher J Arthurs, Miguel O Bernabeu, Rafel Bordas, Jonathan Cooper, Alberto Corrias, Yohan Davit, Sara-Jane Dunn, Alexander G Fletcher, Daniel G Harvey, et al. Chaste: an open source c++ library for computational physiology and biology. *PLoS computational biology*, 9(3):e1002970, 2013.
- [149] Maciej H Swat, Gilberto L Thomas, Julio M Belmonte, Abbas Shirinifard, Dimitrij Hmeljak, and James A Glazier. Multi-scale modeling of tissues using compucell3d. In *Methods in cell biology*, volume 110, pages 325–366. Elsevier, 2012.
- [150] Robert Noble. demon. https://github.com/robjohnnoble/demon_model, 2019.
- [151] Fabrizio Angaroni, Alessandro Guidi, Gianluca Ascolani, Alberto d’Onofrio, Marco Antoniotti, and Alex Graudenzi. J-space: a julia package for the simulation of spatial models of cancer evolution and of sequencing experiments. *BMC bioinformatics*, 23(1):269, 2022.
- [152] Adam Streck, Tom L Kaufmann, and Roland F Schwarz. Smith: spatially constrained stochastic model for simulation of intra-tumour heterogeneity. *Bioinformatics*, 39(3):btad102, 2023.
- [153] Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2020, with forecasts from 2021 to 2025. <https://www.statista.com/statistics/871513/worldwide-data-created/>, August 2023.
- [154] Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. The fair guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9, 2016.
- [155] The european open science cloud (eosc). <https://eosc-portal.eu/about/eosc>.

- 4537 [156] European life-science infrastructure for biological information. [https://elixir-](https://elixir-europe.org)
4538 [europe.org](https://elixir-europe.org).
- 4539 [157] National institutes of health. <https://www.nih.gov>.
- 4540 [158] Global alliance for genomics and health. <https://www.ga4gh.org>.
- 4541 [159] Alan C Hindmarsh, Peter N Brown, Keith E Grant, Steven L Lee, Radu Serban,
4542 Dan E Shumaker, and Carol S Woodward. Sundials: Suite of nonlinear and differen-
4543 tial/algebraic equation solvers. *ACM Transactions on Mathematical Software (TOMS)*,
4544 31(3):363–396, 2005.
- 4545 [160] Jörn Starruß, Walter De Back, Lutz Brusch, and Andreas Deutsch. Morpheus: a
4546 user-friendly modeling environment for multiscale and multicellular systems biology.
4547 *Bioinformatics*, 30(9):1331–1332, 2014.
- 4548 [161] Kimberly D Miller, Rebecca L Siegel, Chun Chieh Lin, Angela B Mariotto, Joan L
4549 Kramer, Julia H Rowland, Kevin D Stein, Rick Alteri, and Ahmedin Jemal. Can-
4550 cer treatment and survivorship statistics, 2016. *CA: a cancer journal for clinicians*,
4551 66(4):271–289, 2016.
- 4552 [162] Fritz H Schröder, Jonas Hugosson, Monique J Roobol, Teuvo LJ Tammela, Stefano
4553 Ciatto, Vera Nelen, Maciej Kwiatkowski, Marcos Lujan, Hans Lilja, Marco Zappa,
4554 et al. Prostate-cancer mortality at 11 years of follow-up. *New England Journal of*
4555 *Medicine*, 366(11):981–990, 2012.
- 4556 [163] Charles Huggins and Clarence V Hodges. Studies on prostatic cancer: I. the effect of
4557 castration, of estrogen and of androgen injection on serum phosphatases in metastatic
4558 carcinoma of the prostate. *The Journal of urology*, 168(1):9–12, 2002.
- 4559 [164] Yoichi Arai, Tatsuhiro Yoshiki, and Osamu Yoshida. Prognostic significance of
4560 prostate specific antigen in endocrine treatment for prostatic cancer. *The Journal of*
4561 *urology*, 144(6):1415–1419, 1990.
- 4562 [165] Brian J Feldman and David Feldman. The development of androgen-independent
4563 prostate cancer. *Nature Reviews Cancer*, 1(1):34–45, 2001.

- 4564 [166] Takeshi Sasaki and Yoshiaki Sugimura. The importance of time to prostate-specific
4565 antigen (psa) nadir after primary androgen deprivation therapy in hormone-naïve
4566 prostate cancer patients. *Journal of clinical medicine*, 7(12):565, 2018.
- 4567 [167] Thenappan Chandrasekar, Joy C Yang, Allen C Gao, and Christopher P Evans. Mech-
4568 anisms of resistance in castration-resistant prostate cancer (crpc). *Translational an-
4569 drology and urology*, 4(3):365, 2015.
- 4570 [168] Marika J Linja, Kimmo J Savinainen, Outi R Saramäki, Teuvo LJ Tammela, Robert L
4571 Vessella, and Tapio Visakorpi. Amplification and overexpression of androgen receptor
4572 gene in hormone-refractory prostate cancer. *Cancer research*, 61(9):3550–3555, 2001.
- 4573 [169] Rajal B Shah, Rohit Mehra, Arul M Chinnaiyan, Ronglai Shen, Debashis Ghosh, Ming
4574 Zhou, Gary R MacVicar, Soorynarayana Varambally, Jason Harwood, Tarek A Bismar,
4575 et al. Androgen-independent prostate cancer is a heterogeneous group of diseases:
4576 lessons from a rapid autopsy program. *Cancer research*, 64(24):9209–9216, 2004.
- 4577 [170] Paulette R Dillard, Ming-Fong Lin, and Shafiq A Khan. Androgen-independent
4578 prostate cancer cells acquire the complete steroidogenic potential of synthesizing
4579 testosterone from cholesterol. *Molecular and cellular endocrinology*, 295(1-2):115–
4580 120, 2008.
- 4581 [171] Federica Recine and Cora N Sternberg. Hormonal therapy and chemotherapy in
4582 hormone-naïve and castration resistant prostate cancer. *Translational Andrology and
4583 Urology*, 4(3):355, 2015.
- 4584 [172] Philip Cornford, Joaquim Bellmunt, Michel Bolla, Erik Briers, Maria De Santis, To-
4585 bias Gross, Ann M Henry, Steven Joniau, Thomas B Lam, Malcolm D Mason, et al.
4586 Eau-estro-siog guidelines on prostate cancer. part ii: treatment of relapsing, metastatic,
4587 and castration-resistant prostate cancer. *European urology*, 71(4):630–642, 2017.
- 4588 [173] Nicholas D James, Melissa R Spears, and Matthew R Sydes. Abiraterone in metastatic
4589 prostate cancer. *The New England journal of medicine*, 377(17):1696–1697, 2017.
- 4590 [174] Richard J Lee. Hormone therapy for prostate cancer. *Cancer Chemotherapy, Im-
4591 munotherapy and Biotherapy*, 12:64, 2011.

- 4592 [175] National comprehensive cancer network. [https://](https://www.nccn.org/patients/guidelines/content/PDF/prostate-advanced-patient.pdf)
4593 [www.nccn.org/patients/guidelines/content/PDF/](https://www.nccn.org/patients/guidelines/content/PDF/prostate-advanced-patient.pdf)
4594 [prostate-advanced-patient.pdf](https://www.nccn.org/patients/guidelines/content/PDF/prostate-advanced-patient.pdf).
- 4595 [176] Laurence H Klotz, Harry W Herr, Michael J Morse, and Willet F Whitmore Jr. Inter-
4596 mittent endocrine therapy for advanced prostate cancer. *Cancer*, 58(11):2546–2550,
4597 1986.
- 4598 [177] Juanita M Crook, Christopher J O’Callaghan, Graeme Duncan, David P Dearnaley,
4599 Celestia S Higano, Eric M Horwitz, Eliot Frymire, Shawn Malone, Joseph Chin, Ab-
4600 denour Nabid, et al. Intermittent androgen suppression for rising psa level after radio-
4601 therapy. *New England Journal of Medicine*, 367(10):895–903, 2012.
- 4602 [178] Maha Hussain, Catherine M Tangen, Donna L Berry, Celestia S Higano, E David
4603 Crawford, Glenn Liu, George Wilding, Stephen Prescott, Subramanian Kanaga Sun-
4604 daram, Eric Jay Small, et al. Intermittent versus continuous androgen deprivation in
4605 prostate cancer. *New England Journal of Medicine*, 368(14):1314–1325, 2013.
- 4606 [179] Sindy Magnan, Ryan Zarychanski, Laurie Pilote, Laurence Bernier, Michele Shemilt,
4607 Eric Vigneault, Vincent Fradet, and Alexis F Turgeon. Intermittent vs continuous an-
4608 drogen deprivation therapy for prostate cancer: a systematic review and meta-analysis.
4609 *JAMA oncology*, 1(9):1261–1269, 2015.
- 4610 [180] Manish S Bhandari, Juanita Crook, and Maha Hussain. Should intermittent andro-
4611 gen deprivation be used in routine clinical practice? *Journal of clinical oncology*,
4612 23(32):8212–8218, 2005.
- 4613 [181] Samuel R Denmeade and John T Isaacs. Bipolar androgen therapy: the rationale for
4614 rapid cycling of supraphysiologic androgen/ablation in men with castration resistant
4615 prostate cancer. *The Prostate*, 70(14):1600–1607, 2010.
- 4616 [182] JP Michiel Sedelaar and John T Isaacs. Tissue culture media supplemented with 10%
4617 fetal calf serum contains a castrate level of testosterone. *The Prostate*, 69(16):1724–
4618 1729, 2009.

- 4619 [183] Marie-Odile Joly-Pharaboz, Alain Ruffion, Anne-Marie Roch, Laurence Michel-
4620 Calemard, Jean André, Jacqueline Chantepie, Brigitte Nicolas, and Geneviève Panaye.
4621 Inhibition of growth and induction of apoptosis by androgens of a variant of Incap cell
4622 line. *The Journal of steroid biochemistry and molecular biology*, 73(5):237–249, 2000.
- 4623 [184] Ruoqian Shen, Makoto Sumitomo, Jie Dai, Adam Harris, David Kaminetzky, Min
4624 Gao, Kerry L Burnstein, and David M Nanus. Androgen-induced growth inhibition
4625 of androgen receptor expressing androgen-independent prostate cancer cells is medi-
4626 ated by increased levels of neutral endopeptidase. *Endocrinology*, 141(5):1699–1704,
4627 2000.
- 4628 [185] Yoshihisa Umekita, Richard A Hiipakka, John M Kokontis, and Shutsung Liao. Hu-
4629 man prostate tumor growth in athymic mice: inhibition by androgens and stimula-
4630 tion by finasteride. *Proceedings of the National Academy of Sciences*, 93(21):11802–
4631 11807, 1996.
- 4632 [186] Chih-pin Chuu, Richard A Hiipakka, Junichi Fukuchi, John M Kokontis, and Shutsung
4633 Liao. Androgen causes growth suppression and reversion of androgen-independent
4634 prostate cancer xenografts to an androgen-stimulated phenotype in athymic mice. *Cancer research*, 65(6):2082–2084, 2005.
- 4636 [187] Michael T Schweizer, Hao Wang, Brandon Lubber, Rosa Nadal, Avery Spitz, D Marc
4637 Rosen, Haiyi Cao, Emmanuel S Antonarakis, Mario A Eisenberger, Michael A Car-
4638 ducci, et al. Bipolar androgen therapy for men with androgen ablation naïve prostate
4639 cancer: results from the phase ii batman study. *The Prostate*, 76(13):1218–1226, 2016.
- 4640 [188] Benjamin A Teply, Hao Wang, Brandon Lubber, Rana Sullivan, Irina Rifkind, Ashley
4641 Bruns, Avery Spitz, Morgan DeCarli, Victoria Sinibaldi, Caroline F Pratz, et al. Bipolar
4642 androgen therapy in men with metastatic castration-resistant prostate cancer after
4643 progression on enzalutamide: an open-label, phase 2, multicohort study. *The Lancet Oncology*, 19(1):76–86, 2018.
- 4645 [189] Punit Saraon, Andrei P Drabovich, Keith A Jarvi, and Eleftherios P Diamandis. Mech-
4646 anisms of androgen-independent prostate cancer. *Ejifcc*, 25(1):42, 2014.

- [190] Raymond J Andersen, Nasrin R Mawji, Jun Wang, Gang Wang, Simon Haile, Jae-Kyung Myung, Kate Watt, Teresa Tam, Yu Chi Yang, Carmen A Banuelos, et al. Regression of castrate-recurrent prostate cancer by a small-molecule inhibitor of the amino-terminus domain of the androgen receptor. *Cancer cell*, 17(6):535–546, 2010.
- [191] E David Crawford and Amy H Hou. The role of lhrh antagonists in the treatment of prostate cancer. *Oncology*, 23(7):626, 2009.
- [192] National Cancer Institute. Hormone therapy for prostate cancer. <https://www.cancer.gov/types/prostate/prostate-hormone-therapy-fact-sheet>, February 2021.
- [193] Fred Saad and Sebastien J Hotte. Guidelines for the management of castrate-resistant prostate cancer. *Canadian Urological Association Journal*, 4(6):380, 2010.
- [194] R Bruce Montgomery, Elahe A Mostaghel, Robert Vessella, David L Hess, Thomas F Kalhorn, Celestia S Higano, Lawrence D True, and Peter S Nelson. Maintenance of intratumoral androgens in metastatic prostate cancer: a mechanism for castration-resistant tumor growth. *Cancer research*, 68(11):4447–4454, 2008.
- [195] Tapio Visakorpi, Eija Hyytinen, Pasi Koivisto, Minna Tanner, Riitta Keinänen, Christian Palmberg, Aarno Palotie, Teuvo Tammela, Jorma Isola, and Olli-P Kallioniemi. In vivo amplification of the androgen receptor gene and progression of human prostate cancer. *Nature genetics*, 9(4):401–406, 1995.
- [196] RSD Brown, J Edwards, A Dogan, H Payne, SJ Harland, JMS Bartlett, and JRW Masters. Amplification of the androgen receptor gene in bone metastases from hormone-refractory prostate cancer. *The Journal of pathology*, 198(2):237–244, 2002.
- [197] J Edwards, N _ S Krishna, KM Grigor, and JMS Bartlett. Androgen receptor gene amplification and protein expression in hormone refractory prostate cancer. *British journal of cancer*, 89(3):552–556, 2003.
- [198] Pasi Koivisto, Juha Kononen, Christian Palmberg, Teuvo Tammela, Eija Hyytinen, Jorma Isola, Jan Trapman, Kitty Cleutjens, Arjan Noordzij, Tapio Visakorpi, et al.

- 4674 Androgen receptor gene amplification: a possible molecular mechanism for androgen
4675 deprivation therapy failure in prostate cancer. *Cancer research*, 57(2):314–319, 1997.
- 4676 [199] Barry S Taylor, Nikolaus Schultz, Haley Hieronymus, Anuradha Gopalan, Yonghong
4677 Xiao, Brett S Carver, Vivek K Arora, Poorvi Kaushik, Ethan Cerami, Boris Reva,
4678 et al. Integrative genomic profiling of human prostate cancer. *Cancer cell*, 18(1):11–
4679 22, 2010.
- 4680 [200] Wennuan Liu, Chunmei Carol Xie, Yi Zhu, Tao Li, Jishan Sun, Yu Cheng, Charles M
4681 Ewing, Sue Dalrymple, Aubrey R Turner, Jielin Sun, et al. Homozygous deletions and
4682 recurrent amplifications implicate new genes involved in prostate cancer. *Neoplasia*,
4683 10(8):897–IN37, 2008.
- 4684 [201] Karine Steketee, Leon Timmerman, Angelique CJ Ziel-van der Made, Paul Doesburg,
4685 Albert O Brinkmann, and Jan Trapman. Broadened ligand responsiveness of androgen
4686 receptor mutants obtained by random amino acid substitution of h874 and mutation hot
4687 spot t877 in prostate cancer. *International journal of cancer*, 100(3):309–317, 2002.
- 4688 [202] John S Sack, Kevin F Kish, Chihuei Wang, Ricardo M Attar, Susan E Kiefer, Yongmi
4689 An, Ginger Y Wu, Julie E Scheffler, Mark E Salvati, Stanley R Krystek Jr, et al. Cry-
4690 tallographic structures of the ligand-binding domains of the androgen receptor and its
4691 t877a mutant complexed with the natural agonist dihydrotestosterone. *Proceedings of
4692 the National Academy of Sciences*, 98(9):4904–4909, 2001.
- 4693 [203] KH Chang, CE Ercole, and N Sharifi. Androgen metabolism in prostate cancer:
4694 from molecular mechanisms to clinical consequences. *British journal of cancer*,
4695 111(7):1249–1254, 2014.
- 4696 [204] Kai-Hsiung Chang, Rui Li, Mahboubeh Papari-Zareei, Lori Watumull, Yan Daniel
4697 Zhao, Richard J Auchus, and Nima Sharifi. Dihydrotestosterone synthesis bypasses
4698 testosterone to drive castration-resistant prostate cancer. *Proceedings of the National
4699 Academy of Sciences*, 108(33):13728–13733, 2011.
- 4700 [205] Michael Stanbrough, Glenn J Bubley, Kenneth Ross, Todd R Golub, Mark A Rubin,
4701 Trevor M Penning, Phillip G Febbo, and Steven P Balk. Increased expression of genes

- 4702 converting adrenal androgens to testosterone in androgen-independent prostate cancer.
4703 *Cancer research*, 66(5):2815–2825, 2006.
- 4704 [206] Mark A Titus, Michael J Schell, Fred B Lih, Kenneth B Tomer, and James L Mohler.
4705 Testosterone and dihydrotestosterone tissue levels in recurrent prostate cancer. *Clini-
4706 cal cancer research*, 11(13):4653–4657, 2005.
- 4707 [207] Muralimohan Yepuru, Zhongzhi Wu, Anand Kulkarni, Feng Yin, Christina M Bar-
4708 rett, Juhyun Kim, Mitchell S Steiner, Duane D Miller, James T Dalton, and Ramesh
4709 Narayanan. Steroidogenic enzyme akr1c3 is a novel androgen receptor-selective coac-
4710 tivator that promotes prostate cancer growth. *Clinical Cancer Research*, 19(20):5613–
4711 5625, 2013.
- 4712 [208] Hannelore V Heemers and Donald J Tindall. Androgen receptor (ar) coregulators:
4713 a diversity of functions converging on and regulating the ar transcriptional complex.
4714 *Endocrine reviews*, 28(7):778–808, 2007.
- 4715 [209] Alejandro S Godoy, Paula C Sotomayor, Marcelo Villagran, Rami Yacoub, Viviana P
4716 Montecinos, Eileen M McNerney, Michael Moser, Barbara A Foster, and Sergio A
4717 Onate. Altered corepressor smrt expression and recruitment to target genes as a mech-
4718 anism that change the response to androgens in prostate cancer progression. *Biochem-
4719 ical and biophysical research communications*, 423(3):564–570, 2012.
- 4720 [210] Christopher W Gregory, Bin He, Raymond T Johnson, O Harris Ford, James L Mohler,
4721 Frank S French, and Elizabeth M Wilson. A mechanism for androgen receptor-
4722 mediated prostate cancer recurrence after androgen deprivation therapy. *Cancer re-
4723 search*, 61(11):4315–4319, 2001.
- 4724 [211] Qianben Wang, Wei Li, Yong Zhang, Xin Yuan, Kexin Xu, Jindan Yu, Zhong Chen,
4725 Rameen Beroukhim, Hongyun Wang, Mathieu Lupien, et al. Androgen receptor reg-
4726 ulates a distinct transcription program in androgen-independent prostate cancer. *Cell*,
4727 138(2):245–256, 2009.

- [212] Joanne Edwards and John MS Bartlett. The androgen receptor and signal-transduction pathways in hormone-refractory prostate cancer. part 2: Androgen-receptor cofactors and bypass pathways. *BJU international*, 95(9):1327–1335, 2005.
- [213] Alex Zhu, Juan Andino, Stephanie Daignault-Newton, Zoey Chopra, Aruna Sarma, and James M Dupree. What is a normal testosterone level for young men? rethinking the 300 ng/dl cutoff for testosterone deficiency in men 20-44 years old. *The Journal of urology*, 208(6):1295–1302, 2022.
- [214] Gianmarco Leone, Consuelo Buttiglieri, Chiara Pisano, Rosario Francesco Di Stefano, Fabrizio Tabbò, Fabio Turco, Francesca Vignani, Giorgio Vittorio Scagliotti, Massimo Di Maio, and Marcello Tucci. Bipolar androgen therapy in prostate cancer: Current evidences and future perspectives. *Critical Reviews in Oncology/Hematology*, 152:102994, 2020.
- [215] Samuel Denmeade, Emmanuel S Antonarakis, and Mark C Markowski. Bipolar androgen therapy (bat): A patient’s guide. *The Prostate*, 82(7):753–762, 2022.
- [216] Samuel R Denmeade, Hao Wang, Neeraj Agarwal, David C Smith, Michael T Schweizer, Mark N Stein, Vasileios Assikis, Przemyslaw W Twardowski, Thomas W Flaig, Russell Z Szmulewitz, et al. Transformer: a randomized phase ii study comparing bipolar androgen therapy versus enzalutamide in asymptomatic men with castration-resistant metastatic prostate cancer. *Journal of Clinical Oncology*, 39(12):1371–1382, 2021.
- [217] Mark Christopher Markowski, Mary-Ellen Taplin, Rahul Raj Aggarwal, Hao Wang, Aliya Lalji, Channing Judith Paller, Catherine Handy Marshall, Michael Anthony Carducci, Mario A Eisenberger, Angelo M De Marzo, et al. Combat-crpc: Concurrent administration of bipolar androgen therapy (bat) and nivolumab in men with metastatic castration-resistant prostate cancer (mcrpc)., 2021.
- [218] Mark C Markowski, Hao Wang, Rana Sullivan, Irina Rifkind, Victoria Sinibaldi, Michael T Schweizer, Benjamin A Teply, Nduku Ngomba, Wei Fu, Michael A Carducci, et al. A multicohort open-label phase ii trial of bipolar androgen therapy in

- men with metastatic castration-resistant prostate cancer (restore): a comparison of post-abiraterone versus post-enzalutamide cohorts. *European urology*, 79(5):692–699, 2021.
- [219] Tin Phan, Sharon M Crook, Alan H Bryce, Carlo C Maley, Eric J Kostelich, and Yang Kuang. Mathematical modeling of prostate cancer and clinical application. *Applied Sciences*, 10(8):2721, 2020.
- [220] Trachette L Jackson. A mathematical model of prostate tumor growth and androgen-independent relapse. *Discrete and Continuous Dynamical Systems-B*, 4(1):187–201, 2003.
- [221] Aiko Miyamura Ideta, Gouhei Tanaka, Takumi Takeuchi, and Kazuyuki Aihara. A mathematical model of intermittent androgen suppression for prostate cancer. *Journal of nonlinear science*, 18:593–614, 2008.
- [222] Takashi Shimada and Kazuyuki Aihara. A nonlinear model with competition between prostate tumor cells and its application to intermittent androgen suppression therapy of prostate cancer. *Mathematical biosciences*, 214(1-2):134–139, 2008.
- [223] Yoshito Hirata, Nicholas Bruchovsky, and Kazuyuki Aihara. Development of a mathematical model that predicts the outcome of hormone therapy for prostate cancer. *Journal of theoretical biology*, 264(2):517–527, 2010.
- [224] Travis Portz, Yang Kuang, and John D Nagy. A clinical data validated mathematical model of prostate cancer growth under intermittent androgen suppression therapy. *Aip Advances*, 2(1), 2012.
- [225] Michaël R Droop. Vitamin b12 and marine ecology. iv. the kinetics of uptake, growth and inhibition in *monochrysis lutheri*. *Journal of the Marine Biological Association of the United Kingdom*, 48(3):689–733, 1968.
- [226] Javier Baez and Yang Kuang. Mathematical models of androgen resistance in prostate cancer patients under intermittent androgen suppression therapy. *Applied Sciences*, 6(11):352, 2016.

- [227] Jeffrey B West, Mina N Dinh, Joel S Brown, Jingsong Zhang, Alexander R Anderson, and Robert A Gatenby. Multidrug cancer therapy in metastatic castrate-resistant prostate cancer: an evolution-based strategy. *Clinical Cancer Research*, 25(14):4413–4421, 2019.
- [228] Robert J Noble, Viola Walther, Christian Roumestand, Michael E Hochberg, Urszula Hibner, and Patrice Lassus. Paracrine behaviors arbitrate parasite-like interactions between tumor subclones. *Frontiers in ecology and evolution*, 9:675638, 2021.
- [229] Francesco Pappalardo, Giulia Russo, Flora Musuamba Tshinanu, and Marco Viceconti. In silico clinical trials: concepts and early adoptions. *Briefings in bioinformatics*, 20(5):1699–1708, 2019.
- [230] Luis-Miguel Chevin. On measuring selection in experimental evolution. *Biology letters*, 7(2):210–213, 2011.
- [231] Stan Development Team. Cmdstand command line interface user’s guide and reference manual (version 2.28.2). <https://mc-stan.org>.
- [232] Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 585(7825):357–362, 2020.
- [233] Jonah J Klowss, Alexander P Browning, Ryan J Murphy, Elliot J Carr, Michael J Plank, Gency Gunasingh, Nikolas K Haass, and Matthew J Simpson. A stochastic mathematical model of 4d tumour spheroids with real-time fluorescent cell cycle labelling. *Journal of the Royal Society Interface*, 19(189):20210903, 2022.
- [234] Michael T Schweizer, Emmanuel S Antonarakis, Hao Wang, A Seun Ajiboye, Avery Spitz, Haiyi Cao, Jun Luo, Michael C Haffner, Srinivasan Yegnasubramanian, Michael A Carducci, et al. Effect of bipolar androgen therapy for asymptomatic men with castration-resistant prostate cancer: results from a pilot clinical study. *Science translational medicine*, 7(269):269ra2–269ra2, 2015.

- [235] FDA. Food and drug administration: Testosterone cypionate injection.
https://www.accessdata.fda.gov/drugsatfda_docs/label/2022/216318s0001b1.pdf, June 2022.
- [236] Karline Soetaert, Thomas Petzoldt, and R Woodrow Setzer. Solving differential equations in r: package desolve. *Journal of statistical software*, 33:1–25, 2010.
- [237] MedlinePlus25. Testosterone injection. <https://medlineplus.gov/druginfo/meds/a614041.html>, March 2019.
- [238] Mary C White, Dawn M Holman, Jennifer E Boehm, Lucy A Peipins, Melissa Grossman, and S Jane Henley. Age and cancer risk: a potentially modifiable relationship. *American journal of preventive medicine*, 46(3):S7–S15, 2014.
- [239] National Health Service (England). Cancer registration statistics, england, 2022 national statistics, accredited official statistics. *NHS Digital*, 2024.
- [240] National Cancer Registry of Ireland (NCRI). Cancers in ireland 1994-2020 (annual statistical report 2022). *NCRI*, 2022.
- [241] Sofía Ruiz-Medina, Silvia Gil, Begoña Jimenez, Pablo Rodriguez-Brazzarola, Tamara Diaz-Redondo, Mireya Cazorla, Marta Muñoz-Ayllon, Inmaculada Ramos, Carmen Reyna, María José Bermejo, Ana Godoy, Esperanza Torres, Manuel Cobo, Laura Galvez, Antonio Rueda, Emilio Alba, and Nuria Ribelles. Significant decrease in annual cancer diagnoses in spain during the covid-19 pandemic: A real-data study. *Cancers*, 13(13), 2021.
- [242] Mariano Provencio, Enric Carcereny, Delvys Rodríguez-Abreu, Rafael López-Castro, María Guirado, Carlos Camps, Joaquim Bosch-Barrera, Rosario García-Campelo, Ana Laura Ortega-Granados, José Luis González-Larriba, Joaquín Casal-Rubio, Manuel Domine, Bartomeu Massutí, María Ángeles Sala, Reyes Bernabé, Juana Oramas, and Elvira del Barco. Lung cancer in spain: information from the thoracic tumors registry (ttr study). *Translational Lung Cancer Research*, 8(4), 2019.
- [243] Hyuna Sung, Jacques Ferlay, Rebecca L Siegel, Mathieu Laversanne, Isabelle Soerjomataram, Ahmedin Jemal, and Freddie Bray. Global cancer statistics 2020: Globocan

- 4837 estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA:*
4838 *a cancer journal for clinicians*, 71(3):209–249, 2021.
- 4839 [244] Freddie Bray, Mathieu Laversanne, Hyuna Sung, Jacques Ferlay, Rebecca L Siegel,
4840 Isabelle Soerjomataram, and Ahmedin Jemal. Global cancer statistics 2022: Globocan
4841 estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA:*
4842 *a cancer journal for clinicians*, 74(3):229–263, 2024.
- 4843 [245] Jianhui Zhao, Liying Xu, Jing Sun, Mingyang Song, Lijuan Wang, Shuai Yuan, Ying-
4844 shuang Zhu, Zhengwei Wan, Susanna Larsson, Konstantinos Tsilidis, Malcolm Dun-
4845 lop, Harry Campbell, Igor Rudan, Peige Song, Evropi Theodoratou, Kefeng Ding, and
4846 Xue Li. Global trends in incidence, death, burden and risk factors of early-onset cancer
4847 from 1990 to 2019. *BMJ oncology*, 2(1):e000049, 2023.
- 4848 [246] Hyuna Sung, Chenxi Jiang, Priti Bandi, Adair Minihan, Miranda Fidler-Benaoudia,
4849 Farhad Islami, Rebecca L Siegel, and Ahmedin Jemal. Differences in cancer rates
4850 among adults born between 1920 and 1990 in the usa: an analysis of population-based
4851 cancer registry data. *The Lancet Public Health*, 9(8):e583–e593, 2024.
- 4852 [247] Benjamin Koh, Darren Jun Hao Tan, Cheng Han Ng, Clarissa Elysia Fu, Wen Hui
4853 Lim, Rebecca Wenling Zeng, Jie Ning Yong, Jia Hong Koh, Nicholas Syn, Wang
4854 Meng, Karn Wijarnpreecha, Ken Liu, Choon Seng Chong, Mark Muthiah, Hung N.
4855 Luu, Arndt Vogel, Siddharth Singh, Khay Guan Yeoh, Rohit Loomba, and Daniel Q.
4856 Huang. Patterns in cancer incidence among people younger than 50 years in the us,
4857 2010 to 2019. *JAMA Network Open*, 6(8):e2328171–e2328171, 08 2023.
- 4858 [248] Archie Bleier and H Gilbert Welch. Effect of three decades of screening mammog-
4859 raphy on breast-cancer incidence. *New England Journal of Medicine*, 367(21):1998–
4860 2005, 2012.
- 4861 [249] Stephen W Duffy, Hsiu-Hsi Chen, Laszlo Tabar, and Nicholas E Day. Estimation of
4862 mean sojourn time in breast cancer screening using a markov chain model of both entry
4863 to and exit from the preclinical detectable phase. *Statistics in medicine*, 14(14):1531–
4864 1543, 1995.

- [250] Harald Weedon-Fekjær, Bo H Lindqvist, Lars J Vatten, Odd O Aalen, and Steinar Tretli. Estimating mean sojourn time and screening sensitivity using questionnaire data on time since previous screening. *Journal of Medical Screening*, 15(2):83–90, 2008.
- [251] AMWM Aarts, SW Duffy, SME Geurts, DP Vulkan, JDM Otten, CY Hsu, THH Chen, ALM Verbeek, and MJM Broeders. Test sensitivity of mammography and mean sojourn time over 40 years of breast cancer screening in nijmegen (the netherlands). *Journal of medical screening*, 26(3):147–153, 2019.
- [252] David Robert Grimes, Edward MA Corry, Talía Malagón, Ciaran O’Riain, Eduardo L Franco, and Donal J Brennan. Modeling cervical cancer screening strategies with varying levels of human papillomavirus vaccination. *JAMA network open*, 4(6):e2115321–e2115321, 2021.
- [253] David Robert Grimes. Impact of human papillomavirus age-related prevalence and vaccination levels on interpretation of cervical screening modalities: a modelling study. *BMJ open*, 14(2):e078551, 2024.
- [254] Gianluca Mauri, Giorgio Patelli, Andrea Sartore-Bianchi, Sergio Abrignani, Beatrice Bodega, Silvia Marsoni, Vincenzo Costanzo, Angela Bachi, Salvatore Siena, and Alberto Bardelli. Early-onset cancers: Biological bases and clinical implications. *Cell Reports Medicine*, 5(9), 2024.
- [255] Steven A Frank. Dynamics of cancer: incidence, inheritance, and evolution. 2018.
- [256] Donald R Charles and Ethel M Luce-Clausen. The kinetics of papilloma formation in benzpyrene-treated mice. *Cancer research*, 2(4):261–263, 1942.
- [257] Alice Whittemore and Joseph B Keller. Quantitative theories of carcinogenesis. *Siam Review*, 20(1):1–30, 1978.
- [258] Hermann J Muller. Radiation damage to the genetic material. ii. effects manifested mainly in the exposed individuals. *Am. Scientist*, 38, 1950.

- 4891 [259] CO2007872 Nordling. A new theory on the cancer-inducing mechanism. *British*
4892 *journal of cancer*, 7(1):68, 1953.
- 4893 [260] Peter Armitage and Richard Doll. The age distribution of cancer and a multi-stage
4894 theory of carcinogenesis. *British journal of cancer*, 91(12):1983–1989, 2004.
- 4895 [261] Leonid A Gavrilov and Natalia S Gavrilova. Reliability theory of aging and longevity.
4896 *Handbook of the Biology of Aging*, pages 3–42, 2005.
- 4897 [262] American Cancer Society. History of acs recommendations
4898 for the early detection of cancer in people without symptoms.
4899 [https://www.cancer.org/health-care-professionals/
4900 american-cancer-society-prevention-early-detection-guidelines/
4901 overview/chronological-history-of-ac-s-recommendations.
4902 html](https://www.cancer.org/health-care-professionals/american-cancer-society-prevention-early-detection-guidelines/overview/chronological-history-of-ac-s-recommendations.html).
- 4903 [263] Miroslav Zavoral, Stepan Suchanek, Ondrej Majek, Premysl Fric, Petra Minarikova,
4904 Marek Minarik, Bohumil Seifert, and Ladislav Dusek. Colorectal cancer screening: 20
4905 years of development and recent progress. *World journal of gastroenterology: WJG*,
4906 20(14):3825, 2014.
- 4907 [264] Carla Parry, Erin E Kent, Angela B Mariotto, Catherine M Alfano, and Julia H Row-
4908 land. Cancer survivors: a booming population. *Cancer epidemiology, biomarkers &*
4909 *prevention*, 20(10):1996–2005, 2011.
- 4910 [265] Sophie Pilleron, Diana Sarfati, Maryska Janssen-Heijnen, Jérôme Vignat, Jacques Fer-
4911 lay, Freddie Bray, and Isabelle Soerjomataram. Global cancer incidence in older
4912 adults, 2012 and 2035: a population-based study. *International journal of cancer*,
4913 144(1):49–58, 2019.
- 4914 [266] David Robert Grimes and Alexander G Fletcher. Close encounters of the cell kind:
4915 The impact of contact inhibition on tumour growth and cancer models. *Bulletin of*
4916 *mathematical biology*, 82:1–13, 2020.

4917 Appendix A

4918 Appendices for Chapter 2




FIGURE A.1: Frequency Data and Selection Coefficients for the 19:1 assays for both WT-OR replicates. The top three rows show the results of the first replicate, and the bottom three the second. For each group of plots for each replicate, the top row shows WT frequencies (green) and OR frequencies (red) on a linear scale, the middle row shows the selection coefficient of the WT cells relative to the OR cells (Dashed line represents estimated selection coefficient, black points show calculated selection coefficients from timepoint $t - 1$ to timepoint t , and bold line is the mean average of these points) and the bottom row shows the same frequencies as the top row in logit scale.




FIGURE A.2: Frequency Data and Selection Coefficients for the 9:1 assays for both WT-OR replicates. The top three rows show the results of the first replicate, and the bottom three the second. For each group of plots for each replicate, the top row shows WT frequencies (green) and OR frequencies (red) on a linear scale, the middle row shows the selection coefficient of the WT cells relative to the OR cells (Dashed line represents estimated selection coefficient, black points show calculated selection coefficients from timepoint $t - 1$ to timepoint t , and bold line is the mean average of these points) and the bottom row shows the same frequencies as the top row in logit scale.




FIGURE A.3: Group-level posterior Histograms for the 0nM (top), 1nM (middle) and 10nM (bottom) WT Inferences.




FIGURE A.4: Comparison between model fit and data for 0nM (top), 1nM (middle) and 10nM (bottom) WT homogeneous assays. Model fit is represented by a solid line with 95% confidence intervals, and data is represented by points. Information pertaining to WT cells is in green and OR cells is in Red. Both replicates shown.




FIGURE A.5: Group-level posterior Histograms for the 0nM (top), 1nM (middle) and 10nM (bottom) OR Inferences.




FIGURE A.6: Comparison between model fit and data for 0nM (top), 1nM (middle) and 10nM (bottom) OR homogeneous assays. Model fit is represented by a solid line with 95% confidence intervals, and data is represented by points. Information pertaining to WT cells is in green and OR cells is in Red. Both replicates shown.




FIGURE A.7: Group-level posterior Histograms for the 0nM (top 2 rows), 1nM (middle 2 rows) and 10nM (bottom 2 rows) 99:1 competitive inferences.




FIGURE A.8: Comparison between model fit and data for 0nM (top), 1nM (middle) and 10nM (bottom) 99:1 competitive assays. Model fit is represented by a solid line with 95% confidence intervals, and data is represented by points. Information pertaining to WT cells is in green and OR cells is in Red. Both replicates shown.




FIGURE A.9: Group-level posterior Histograms for the 0nM (top 2 rows), 1nM (middle 2 rows) and 10nM (bottom 2 rows) 1:1 Inferences. Replicate 2 was lost for these assays.




FIGURE A.10: Comparison between model fit and data for 0nM (top), 1nM (middle) and 10nM (bottom) 1:1 competitive assays. Model fit is represented by a solid line with 95% confidence intervals, and data is represented by points. Information pertaining to WT cells is in green and OR cells is in Red. Replicate 2 was lost for these assays.




FIGURE A.11: Homogeneous WT Area data – 1nM, 10nM OSM. Experiment 1. Replicates with missing data points or lost replicates are omitted. Linear scale on top row, log scale on bottom row.




FIGURE A.12: 99:1 WT:OR Area data – 0nM, 1nM, 10nM OSM. Experiment 1. Replicates with missing data points or lost replicates are omitted. Linear scale on top row, log scale on bottom row.




FIGURE A.13: Homogeneous OR Area data – 0nM, 1nM, 10nM OSM. Experiment 1. Replicates with missing data points or lost replicates are omitted. Linear scale on top row, log scale on bottom row.




FIGURE A.14: Homogeneous WT Area data – 0nM, 1nM, 10nM OSM. Experiment 1. Replicates with missing data points or lost replicates are omitted. Linear scale on top row, log scale on bottom row.




FIGURE A.15: 99:1 WT:OR Area data – 1nM, 10nM OSM. Experiment 1. Replicates with missing data points or lost replicates are omitted. Linear scale on top row, log scale on bottom row.




FIGURE A.16: 19:1 WT:OR Area data – 1nM, 10nM OSM. Experiment 1. Replicates with missing data points or lost replicates are omitted. Linear scale on top row, log scale on bottom row.

4919 Appendix B

4920 Appendices for Chapter 3




FIGURE B.1: Results of the second replicates of the drug-free co-culture frequency assays with varying initial ratios. Top row depicts the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. The Bottom row compares the predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line).




FIGURE B.2: Results of the second replicates of the co-culture frequency assays with a drug dose of 0.5nM. Top row depicts the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. The Bottom row compares the predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line).




FIGURE B.3: Results of the second replicates of the co-culture frequency assays with a drug dose of 1nM. Top row depicts the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. The Bottom row compares the predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line).




FIGURE B.4: Results of the second replicates of the co-culture frequency assays with a drug dose of 2nM. Top row depicts the frequency dynamics over time on a logit scale, with WT cells shown in green, and OR cells shown in red. The Bottom row compares the predicted selection coefficient (solid black line) with the calculated selection coefficients at every point (black dots). The mean of these calculated selection coefficient is also shown (dashed green line).




FIGURE B.5: Results of the *in vivo* mice experiments for the ER cohort. All 5 replicates are shown, with the top row showing the data on a linear scale, and the bottom row showing the data on a log scale.




FIGURE B.6: ER Model fit for each replicate used in the inference. 100 parameter sets were drawn from the posterior distributions and then used in the initial model (red lines) and compared to the data (black dots).




FIGURE B.7: Histograms displaying the inferred posterior distributions for both hierarchical and group-level parameters. Left column shows these distributions against their prior; right column shows them with the prior excluded, for clarity.

4921 **Appendix C**

4922 **Appendices for Chapter 5**




FIGURE C.1: Rate change model fits compared to pan-cancer data for every year.




FIGURE C.2: Detection change model fits compared to pan-cancer data for every year.




FIGURE C.3: Rate change model fits compared to Breast cancer data for every year.




FIGURE C.4: Detection change model fits compared to Breast cancer data for every year.




FIGURE C.5: Rate change model fits compared to colorectal Cancer data for every year.




FIGURE C.6: Detection change model fits compared to colorectal cancer data for every year.




FIGURE C.7: Rate change model fits compared to stomach cancer data for every year.




FIGURE C.8: Detection change model fits compared to stomach cancer data for every year.




FIGURE C.9: Rate change model fits compared to pancreatic cancer data for every year.




FIGURE C.10: Detection change model fits compared to pancreatic cancer data for every year.