Mathematical Modelling of Extinction Therapy: Preventing Evolutionary Rescue in Cancer Populations

A Thesis

Submitted towards the partial fulfilment of the requirements

for the BS-MS Dual Degree programme

by

Srishti Patil

Date: April, 2023

Under the guidance of

Dr. Robert Noble

City, University of London

From May 2022 to Mar 2023

INDIAN INSTITUTE OF SCIENCE EDUCATION AND RESEARCH

PUNE - 411 008, INDIA

Certificate

This is to certify that this dissertation entitled “Modelling Extinction Therapy: Preventing Evolutionary Rescue in Cancer Populations” submitted towards the partial fulfillment of the BS-MS degree at the Indian Institute of Science Education and Research, Pune represents original research carried out by Srishti Patil at City, University of London, under the supervision of Dr Robert Noble during academic year May 2022 to March 2023.

Dr. Robert Noble

Project Supervisor

Lecturer in Mathematics,

City, University of London

Srishti Patil

BS-MS, IISER Pune

20181110

Date: May 10, 2023

Declaration

I, hereby declare that the matter embodied in the report titled “Modelling Extinction Therapy: Preventing Evolutionary Rescue in Cancer Populations” is the results of the investigations carried out by me at City, University of London from the period May 2022 to March 2023 under the supervision of Dr Robert Noble and the same has not been submitted elsewhere for any other degree.

Srishti Patil

BS-MS, IISER Pune

20181110

Date: May 10, 2023

Acknowledgements

I thank the Arizona Cancer Evolution Center (ACE) for funding on part of this project, which allowed me to travel to the UK and work in-person at City, University of London. I would especially like to thank my supervisor, Dr. Robert Noble for believing in me and giving me this wonderful opportunity to work in one of the biggest cities in the world and to interact with amazing researchers from all over the world. Most of all, thank you for being an awesome mentor, for teaching me so much throughout the last year, for guiding and motivating me to do great research.

To all my mentors throughout my journey at IISER, thank you for teaching me, identifying the potential in me and helping me channel my skills in the best possible manner. Specifically, I thank Dr. Sourabh Dube for my first foray into research Dr. Chaitanya Gokhale for guiding me on my first extended, independent research project in my current field. They motivated me to explore my interests and I will forever be grateful for that. I thank IISER Pune for providing an amazing research and study environment in which I could freely learn, apply and challenge myself.

I thank my parents and family for supporting me and sharing my happiness in all things small and big. Their complete love and faith in me keeps me going. This work is dedicated to them. I would also like to thank my friends Misaal, Sultan and Aharna for their constant love and companionship since the beginning of my journey into research. I cannot imagine reaching where I am today without them. Lastly, I would like to thank my partner Lokesh for his love, support, feedback and encouragement. The last few months have been challenging but he kept pushing me forward. This work is what it is today because of all these people and I thank them from the bottom of my heart.

Contributions

Contributor name	Contributor role
Srishti P. and Robert N.	Conceptualization Ideas
Srishti P. and Robert N.	Methodology
Srishti P.	Software
Srishti P.	Validation
Srishti P.	Formal analysis
Srishti P.	Investigation
Robert N.	Resources
–	Data Curation
Srishti P.	Writing - original draft preparation
Srishti P., Robert N. and Yannick V.	Writing - review and editing
Srishti P.	Visualization
Robert N.	Supervision
Robert N.	Project administration
Robert N.	Funding acquisition

Abstract

Evolutionary therapies for cancer understand malignancies as adapting populations under Darwinian selection. They use concepts from ecology and evolutionary biology to deal with the emergence of resistance in these malignancies – a big problem in cancer treatments. Extinction Therapy (ET) is an evolutionary therapy that aims for complete eradication of the tumour. It fights the emergence of resistance with the smart and effective use of drugs/treatments to exploit the vulnerability of a small or declining population using multiple strikes (in the form of drugs, surgery, etc). In other words, extinction therapy "kicks the tumour while it's down". In this thesis, we model ET analytically using evolutionary rescue theory and run stochastic simulations to understand the behaviour of a cancer population undergoing ET. We also perform predictive mathematical modelling to aid the design and analysis of future experiments in ET. We find that the timing of subsequent strikes (after the primary therapy) is a very important determinant of the extinction probability. We calculate the optimal timing for these strikes and show how it changes with other model parameters. This work is one of the first few models of ET and sets the stage for future analytical and computational work in the field.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Evolutionary Therapies in Cancer

Cancer populations, in many ways, are akin to an ecosystem consisting of different species. Just like species in an ecosystem interact, compete for resources, adapt to changing environmental conditions and undergo natural selection, so do cancer cells in a tumour's ecosystem. Darwinian principles guide the evolutionary trajectories of the many heterogeneous cancer subpopulations, which inevitably affects their response to therapies (Iwasa, Nowak, and Michor 2006). Understanding the evolutionary processes that drive cancer growth can provide valuable insights into developing more effective therapies. (Korolev, Xavier, and Gore 2014; Enriquez-Navas, Wojtkowiak, and Gatenby 2015).

Cancer therapy has come a long way in recent years, but resistance to treatment remains a significant hurdle in the fight against the disease. Despite initial success in reducing tumour size and symptoms, cancer cells often evolve and develop resistance to the therapies used against them. This can lead to a relapse of the disease, making it difficult to achieve long-term survival rates. Despite the critical role that resistance plays in cancer outcomes, very few conventional treatments account for the evolutionary dynamics of cancer populations which can be leveraged to the patient's benefit (Aktipis et al. 2011).

One of the challenges of targeting specific molecular mechanisms is that cancer cells can use a variety of adaptive strategies to develop resistance (Pressley et al. 2021). Cancer cells can utilize a range of genetic mutations and alterations to evade therapies. This means that targeting a single mechanism may be ineffective in the long term, as cancer cells can quickly adapt and find alternative ways to survive (Greaves and

Maley 2012). Therefore, a more comprehensive understanding of the diverse adaptive strategies used by cancer cells is necessary to develop effective long-term solutions.

Mathematical modelling of clonal dynamics and emergence of resistance in cancer can facilitate our study of fundamental principles governing the system, which is critical for optimizing the response and outcome of clinical treatments based on evolutionary principles, specifically designed to fight resistance in cancer systems (Gatenby and Brown 2020). Historical development of evolutionary therapies has followed a trajectory that begins with exploration of eco-evolutionary concepts through theoretical disciplines such as mathematics (West, Adler, et al. 2023). These concepts are subsequently built upon by various studies to achieve a complete understanding of the phenomenon. Laboratory experiments serve to apply these concepts and help elucidate the process in a real-world setting. To further translate the findings to clinical practice, modelling is used to calculate the effects of specific drugs and compounds on the relevant system under controlled settings. Finally, the treatment undergoes clinical trials. Therefore, a thorough comprehension of the eco-evolutionary dynamics of a cancer population is the foundation of an evolutionary therapy.

1.1.1 The development of Adaptive Therapy

Adaptive therapy is a great example of an idea that began as a mathematical model and developed into an evolutionary therapy in clinical trials. Contemporary cancer treatments adhere to the notion that the ultimate objective is to achieve the maximal feasible reduction in the size of the malignant cell population. This approach is typically achieved through the implementation of Maximum Tolerated Dose (MTD), a therapeutic strategy in which toxicity due to a treatment is the primary limit to the amount of drug administered. Essentially, the MTD approach seeks to administer the maximum possible dose that can eliminate the greatest number of cancer cells while avoiding excessive toxic side-effects to the patient. However, investigation into the evolutionary principles underlying cancer population dynamics challenges this idea (West, Adler, et al. 2023; Enriquez-Navas, Wojtkowiak, and Gatenby 2015; Gatenby and Brown 2020).

Adaptive therapy is a therapeutic strategy aimed at controlling tumour burden by exploiting the competition between sensitive and resistant clones within a heterogeneous cancer population (Monro and Gaffney 2009; Gatenby, Silva, et al. 2009). The phenomenon of competitive release, where sensitive cells are killed off, leaving resistant cells to grow without competition, is a major contributor to relapse in conventional cancer therapies (West, Ma, and Newton 2018). Unlike conventional cancer treatments, which aim to reduce tumour size as much as possible (often leading to competitive release), adaptive therapy aims to maintain a controlled tumour size while delaying the

emergence of resistance. It does so by altering the dose of the treatment in response to the changing tumour burden. Therefore, the treatment dose in adaptive therapy is dynamically determined using evolving population dynamics with the aim of long-term control of tumour burden.

Historically, the term 'adaptive therapy' was termed by Gatenby, Silva, et al. (2009), but the concept was introduced much earlier in 1992 (R. B. Martin et al. 1992). Following the work of Gatenby, Silva, et al. (2009), there was much theoretical work on the conceptual exploration of adaptive therapy, employing methods like optimal control (J. J. Cunningham et al. 2018; R. B. Martin et al. 1992), agent based models (You et al. 2017; J. A. Gallaher, Brown, and Anderson 2019; J. Gallaher et al. 2022), Lotka-Volterra models (Zhang et al. 2017; Strobl et al. 2020) and frequency/density dependent models (Viossat and Noble 2021; Brady-Nicholls et al. 2021) among others. The theory of adaptive therapy was generalized and made more robust with recent progress in the area (Viossat and Noble 2021; West, Adler, et al. 2023). Several *in-vivo* and *in-vitro* experimental studies further analysed adaptive therapy in more specific contexts (Gatenby, Silva, et al. 2009; Zhang et al. 2017; Enriquez-Navas, Kam, et al. 2016; Kavran et al. 2022). This led to the first few clinical trials of adaptive therapy (Moffitt; NCT02415621, Moffitt; NCT03543969, etc). Further clinical trials are planned while more mathematical, computational and experimental studies continue to explore the field.

What we learn from the development of adaptive therapy is that collaborative research across experimental, clinical, theoretical and computational domains is best for a holistic progress of any evolutionary therapy for cancer. As theoretical foundations are laid by works like Viossat and Noble (2021), clinical trials (ANZadapt; NCT05393791) continue to test the feasibility of such therapies in practice. An interdisciplinary view goes a long way to ensure that all fundamental aspects of any treatment-motivated phenomenon are well understood and implemented.

1.2 Drawing inspiration from background extinctions

Large-scale eco-evolutionary interactions of competition and population control inspire adaptive therapy. Similarly, our focus, Extinction Therapy (ET) draws inspiration from large-scale eco-evolutionary dynamics of species extinction events. There are broadly two types of extinction events – mass extinctions and background extinctions (Walther et al. 2015). Mass extinctions occur when a single impact or strike drives entire populations and communities to extinction. In these events, a single strike is effective enough that the population is unable to bounce back and eventually goes extinct. This can be

due to high environmental stress or reduced adaptive capabilities among other factors. Background extinctions are not as dramatic, and involve multiple events that eventually lead to a species' demise (Walther et al. 2015). They are much more common than mass extinctions because a single strike can often fail to lead to extinction due to the presence of resistance in the population. These are the cases that lead to evolutionary rescue, in which a population destined to go extinct undergoes adaptive evolutionary changes and escapes extinction. Different events in the process of background extinction typically have different mechanisms of working. An individual surviving one event and becoming resistant to it could still be susceptible to another strike. Even though we talk about evolutionary rescue in terms of conservation biology, with the motive of avoiding extinction, the same theory is applicable in cases where we want a population to go extinct, like bacterial populations or cancer (Alexander et al. 2014).

1.3 The concept of Extinction Therapy

Anthropocene extinctions in recent times, like that of the Galapagos goat (Gatenby, Artzy-Randrup, et al. 2020), demonstrate how knowledge of eco-evolutionary processes in the target population can help us predict and control future evolutionary trajectories of the population. Similarly in the case of cancer populations, since the oncologist has the power of controlling the tumour environment, it provides them an opportunity to anticipate evolutionary trajectories of cancer cells and treat the tumour accordingly. Under the right conditions, with the right treatments, an extinction vortex can be created, leading to complete cure (Gatenby, Zhang, and Brown 2019). This is what we call Extinction Therapy (ET), which aims for a complete eradication of cancer meta-populations as opposed to long-term maintenance and control, as in Adaptive Therapy (AT).

How do multiple strikes lead to extinction? Even if single-strike treatments fail to completely eradicate cancer populations due to resistant phenotypes, they succeed in leaving the population small and fragmented. The small population size is highly vulnerable to changes due to drift, and it is less capable of adapting to environmental changes owing to loss in heterogeneity (Alexander et al. 2014). Individual growth rates may also be limited due to Allee effects (Dennis et al. 2016). Subsequent strikes (therapies) can take advantage of the stochasticity of small populations to initiate an extinction vortex. This means that the population is driven below the Minimum Viable Population (MVP) threshold, after which it is highly unlikely for the cancer to survive, due to demographic/environmental stochasticity and further reduced adaptability.

Subsequently, there are few important factors to consider when attempting to study

ET. First, how effective is the first strike and does it make the population vulnerable enough for further strikes to work? How do we characterise the threshold below which a population is considered appropriate for a subsequent strike? Second, we need to consider the emergence of resistant clones in the population, both pre-existing mutants and those generated during the treatment. What is the probability that a population is rescued by these resistant mutants? To calculate that, we must find the probability of a resistant clone escaping stochastic extinction and crossing the "establishment threshold", beyond which selection dominates over drift. Lastly, we must take into account effects on growth rates of clonal lineages in the population,and especially cost of resistance and density dependence (including Allee effects).

With the aim of complete eradication, ET draws from eco-evolutionary concepts of population extinctions and exploits them in the context of cancer treatment. To do this, multiple "strikes" are employed at different times, again based on the population dynamics of a continuously evolving population.

1.3.1 How is ET different from Combination Therapy?

Combination therapies in cancer use multiple drugs or multiple treatment strategies in order to target a larger section of the cancer cell population (Meille et al. 2016). These combinations are generally chosen such that cells resistant to one treatment/drug are sensitive to the other. The concept of such collateral sensitivity is important for ET as well. Multiple strikes will only be effective if they can kill cells that are resistant to the first strike. The difference between ET and combination therapies is the timing of the second strike and the use of evolutionary principles to guide treatment strategy.

In combination therapy, or sequential therapy, second or subsequent treatments are usually given when the first treatment shows sign of failure and the population begins to relapse. In other cases, multiple drugs with collateral sensitivities are simultaneously administered at the beginning of the treatment to achieve a maximum response (Chakrabarti and Michor 2017). Conversely, in the case of extinction therapy, the objective is to attack the cancer population while it is at its most vulnerable stage, likely resulting in the population falling below the No Evidence of Disease (NED) threshold, rendering it clinically undetectable. Moreover, previous models of ET have hypothesized that it may be necessary to give the second strike *while* the tumour is still decreasing in size in response to the first therapy.

The success of extinction therapy is expected to be highly sensitive to the timing of the second strike. It needs to be administered in a time window within which the population is vulnerable and susceptible to it. If it is given before the population reaches that

stage, or after it has relapsed and crossed the establishment threshold again, then the second strike will not lead to extinction (Gatenby, Artzy-Randrup, et al. 2020). Sometimes, the timing may seem non-intuitive but in it lies the strength of ET – its basis in eco-evolutionary theory. Essentially, empirical instances in which a combination or sequence of treatments are used in the clinic conventionally follow a fixed treatment regimen with fixed cycles of interventions, mostly independent of tumour response (Reed et al. 2020). A hallmark of evolutionary cancer therapies is that treatment is given *in response* to the dynamic tumour burden. Extinction therapy is different from combination therapy or sequential therapy in the same way adaptive therapy is different from intermittent dosing and metronomic therapy.

1.4 Existing models of Extinction Therapy

One of the first models of extinction therapy was proposed by Gatenby, Artzy-Randrup, et al. (2020). Their work uses simulations to provide a proof of concept of the idea behind ET. A population in a steady state is put under environmental stress (primary therapy). As a result, some cell lineages have a selective advantage depending on the degree of resistance in that lineage. Resistance is modelled as a continued trait. This means that there is no single wild-type and resistant variant, but rather it is a distribution of fitness values in which the higher values are relatively more resistant. This way of modelling resistance in cellular populations is an alternative to considering fully resistant/sensitive phenotypes, but both are considered equally valid.

A second therapy or ”strike” is given some time after the onset of primary therapy, with the expectation that the population will go extinct, owing to the vulnerability of small populations due to the stochastic risk of extinction and Allee effects (Gatenby, Zhang, and Brown 2019). The timing and severity of the second strike are essential determinants of extinction dynamics. The original model (Gatenby, Artzy-Randrup, et al. 2020) provides estimates of a window of opportunity in which a second strike is the most effective, and how Allee effects increase the efficacy of the treatment.

What the existing model lacks is a quantitative calculation of the extinction probabilities of population under this treatment. While it sets the stage for further work on extinction therapy, there is much to be done. Most of the questions regarding the timing of the second strike, the time till extinction, the effect of environmental and demographic factors, and most importantly the conditions under which ET is a feasible alternative to other therapies, remain unanswered.

In this thesis, we achieve three main objectives – 1) developing the first analytical model to describe extinction therapy. We obtain expressions for extinction probabilities

and optimal timing for the second strike. We derive from concepts in evolutionary rescue to build a simple, tractable mathematical model with few assumptions. 2) We run stochastic simulations to understand the system in a more general setting. We understand the effect that demographic stochasticity plays when we compare the analytical results with simulation results. Ultimately, we draw inferences about the behaviour of a system under extinction therapy with many different parameters. 3) In anticipation of experiments to be conducted on extinction therapy, we perform predictive modelling to aid design and analyses in the laboratory.

Chapter 2

Theory

A tumour cell population under treatment is analogous to a population of a species faced with a severe environmental change. Extinction in both scenarios can be avoided by adaptation to the new environment via the evolution of resistant phenotypes (either by genetic adaptation or plasticity) (Alexander et al. 2014). This phenomenon of the prevention of extinction due to adaptive evolutionary changes is termed evolutionary rescue. It is different from other forms of rescue like demographic rescue which occurs when the population is rescued due to population dispersal and immigration of fitter phenotypes. In evolutionary rescue, resistance must emerge at a timescale similar to that of population decay under environmental stress, which is why it is often described as a "race against extinction" (Carlson, C. J. Cunningham, and Westley 2014). While the conventional goal of evolutionary rescue is to maximise the probability of rescue, the same mathematical formulations can be applied with the aim of extinction.

In this thesis, we use the extensive theory of evolutionary rescue (ER) to find the probability of extinction in cancer populations undergoing extinction therapy. Mathematical models of ER analytically study systems that experience external stress, determining how and under what conditions is it possible for the population to undergo evolutionary changes in order to rescue the population (Bell 2013). There are many factors that affect this process, broadly divided into: genetic factors, demographic factors and external factors (Carlson, C. J. Cunningham, and Westley 2014). Depending on the system one is working with, different factors end up playing important roles.

For our case, a mathematical model of evolutionary rescue in an isolated, asexual population under two different environmental stresses, we must have three essential components corresponding to the interactions between the three most important determinants of population extinction (Alexander et al. 2014; G. Martin et al. 2013; Bell 2017). These interactions are summarised in Fig. 2.1A. Small population size, low

genetic variance and a high degree of environmental stress are factors that might lead to species extinction (Bell 2013).

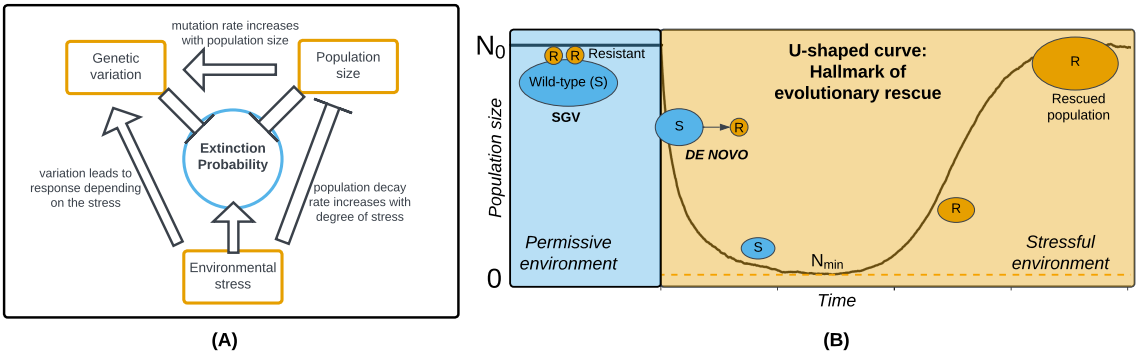

Figure 2.1: **(a)** Components of a typical mathematical formulation of evolutionary rescue. The three most important determinants of extinction risk and how they affect each other must be specified to quantify the probability of rescue. **(b)** A schematic of evolutionary rescue, adapted from Alexander et al. (2014). Standing genetic variation (SGV) and *de-novo* mutations are shown. The rescued population is composed of resistant cells, while the sensitive (wild-type) population goes extinct. This is an illustration of the typical U-shaped trajectory seen in evolutionary rescue.

Small populations are at a higher risk of stochastic extinction. Furthermore, the total mutation rate is lower because the number of individuals is small, so the generation of resistant lineages is slow. Even if resistant lineages exist, they might die out due to demographic stochasticity. The risk is even greater if one considers Allee effects (reduced growth rate at low population sizes) (Dennis et al. 2016). Therefore, changes in population size with environmental stress must be specified.

Prior to the onset of stress, it is assumed that the resistant variants are rare or non-existent. If they are abundant, it is almost certain that the population will survive (Orr and Unckless 2014). Given these conditions, one can assume density-independent growth (in discrete or continuous time) for the resistant mutants in the beginning and use branching processes to model it (Orr and Unckless 2008).

A population is rescued from extinction when one or more resistant lineages fix in the population. The generation of resistant mutants is controlled by the mutation rate, which may or may not depend on the degree of stress. The mutants that ultimately establish themselves in the population and lead to evolutionary rescue are called rescue mutants. The probability of establishment of a resistant mutant can be approximated by Haldane's $2s$ (Haldane 1927), where s is the selective advantage of the mutant over the wild-type, which depends on the degree of stress. For a more general, continuous-time estimate, we use diffusion approximation to get the establishment probabilities of resistant lineages staring with one cell (Lambert 2006; Alexander et al. 2014),

$$\pi_e = 1 - \exp\left(\frac{-2r}{\sigma}\right) \tag{2.1}$$

where r is the growth rate of the resistant lineage and σ is its reproductive variance. This expression is derived using a stochastic population growth equation from (Lambert 2006). We use Eq 2.1 in the following sections to calculate extinction probabilities.

An important distinction is that the probability of rescue by pre-existing mutants (from before the onset of stress, called standing genetic variation) is different from that of new mutants (via *de-novo* mutations) (Orr and Unckless 2008). This is because the pre-existing mutants (SGV) typically get more time than the *de-novo* mutants (DN) to grow their lineage. There are several expressions by different authors (G. Martin et al. 2013; Uecker, Sarah P. Otto, and Hermisson 2014; Orr and Unckless 2008) for the probability of evolutionary rescue by both these classes of mutants, but considering the common conceptual basis, they all reduce to the following,

$$P_{\text{SGV}} = 1 - \exp(-N_0 \pi_e f_0) \tag{2.2}$$

$$P_{\text{DN}} = 1 - \exp\left(-\mu \pi_e \int_0^{t_{\text{ext}}} N_t dt\right) \tag{2.3}$$

where N_0 is the population size at the onset of stress $(t = 0)$, μ is the per capita, per unit time mutation rate after the onset of stress, f_0 is the frequency of resistant variants at $t = 0$ and t_{ext} is the time at which the population would go extinct if it is not rescued. We usually try to get the integral in Eq 2.3 with respect to population size, because t_{ext} is not easy to determine analytically. The above equations are derived on the basis of one main assumption – the generation of rescue mutants in the population is a Poisson process (Alexander et al. 2014). Once the rate of the Poisson process is calculated, it is easy to obtain the probability of extinction which will be equal to the probability that no rescue mutants are generated.

Additionally, we make assumptions about the density independence of rescue lineages and that the probability of establishment of a resistant mutant is constant throughout. There is extensive literature on what happens when these assumptions do not hold. Analytical results can be derived in all such situations using stochastic methods and taking continuous time approximations (Uecker, Sarah P. Otto, and Hermisson 2014; G. Martin et al. 2013; S. P. Otto and Whitlock 1997; Kuosmanen n.d.).

To summarise, with evolutionary rescue models, one can derive the probability of extinction of a population under stress, given its initial state. In Section 4.2, we use

these results extensively in the context of extinction therapy. The principles of evolutionary rescue provide a foundation for building the theoretical formulation of extinction therapy. There are, of course, significant differences to account for – firstly, most evolutionary rescue models consider either one abrupt change in the environment or a continuous, gradual change (Bell 2017). However, extinction therapy is based on using two or more subsequent strikes, all of which are abrupt changes in the environment. Second, most existing models consider a single resistant variant (an exception is G. Martin et al. (2013)), while the existing model for extinction therapy works with a continuum of resistance effects. Both these constructions are equally valid in literature, so we choose to develop the simpler case in which we have discrete genetic variation. Third, the assumption of density-independent interactions is often made in evolutionary rescue models, but this neglects Allee effects which may become important in the case of extinction therapy (Brown, J. J. Cunningham, and Gatenby 2017). We check if Allee effects are needed for extinction therapy to work and include appropriate density dependence. We therefore theoretically understand extinction therapy as the prevention of evolutionary rescue.

Chapter 3

Methods and Model Building

3.1 Reproducing and Analysing the Existing Model

As ideated in the paper by Gatenby, Zhang, and Brown (2019), extinction therapy is a broad concept in the sense that it relies on fundamental properties of an ecological system moving towards extinction. One of their key arguments is that a two-strike model exploits the stochasticity of small and vulnerable populations to drive the cancer population to extinction (Section 1.4). Their model uses stochastic simulations of the system under different treatment schedules, using the Gillespie algorithm (Section 3.3.1) to obtain their results. The major results from the model define a window of time in which the second strike is the most effective. This window is shown to be close to the minimum population reached before relapse. These results, although promising, are limited by the assumptions and constraints of their specific model, which are particularly complex, involving many assumptions and illustrating a limited number of cases.

The main assumptions and limitations of their model are:

- The nature of resistance: An arbitrary measure of resistance $x \in [0, 1]$ is used to generate heterogeneity in the cancer cell population. The proliferation rates depend on the value of x, which is chosen at random in the beginning of the simulation. This means that there are no fully resistant or fully sensitive cells. There is a continuum of resistance levels within the population and the response to therapy on any individual cell depends on its value of x.
- Nature of the second strike: The second strike takes the form of an instant reduction in the cell population, with the efficacy of 20-40%. The primary treatment is continued throughout.

- Allee effects: The reduction in growth rates at small population sizes plays an important role in the model. The window of opportunity for the second strike is shown to increase in size if the Allee threshold increases. It is not shown, however, if the model works without any Allee effects.
- Particular functional forms: Response to therapy and proliferation rate of an individual cell has a specific dependence on its level of resistance. This may be restrictive and the results might be sensitive to changes in these functional forms. The model parameters and descriptions are given in Table 3.1 Moreover, we are left with a number of parameters used to specify the functional form of these variables, but have no physical interpretation. This makes the model hard to translate to a laboratory setting.

Parameter/Variable	Symbol	Default values/formulae
Initial population size	N_0	10000
Carrying capacity	K	10000
Allee threshold	A	15
Resistance level of a cell	$x \in [0,1]$	None
Parameter to control Allee effects	a	0
Concentration of drug (treatment 1)	C	0.4
Minimum intrinsic birth rate	$\lambda_{\min}$	0.2
Maximum intrinsic birth rate	$\lambda_{\max}$	1
Intrinsic death rate	μ	0.1
Parameter to control cost of resistance	s	0.25
Parameters to control drug-induced death rate	(α, β)	$(0.8, 0.2)$
Birth rate of a cell	$b(t)$	$\lambda_{\max} - x^s(\lambda_{\max} - \lambda_{\min})$
Death rate of a cell	$d(t)$	$\mu + D$
Drug-induced death rate	D	$\max[0, \beta(C-x)^\alpha]$

Table 3.1: List of parameters and variables, their descriptions and default values in the model given by Gatenby, Artzy-Randrup, et al. (2020)

Understanding the generality and applicability of the existing model was the first step towards building our own model. However, despite repeated attempts over the last year, we were unable to obtain the simulation code from the authors of the papers. Therefore, following the model description in Gatenby, Artzy-Randrup, et al. (2020), we attempted to reproduce, verify and analyse their results. A few additions were made in order to understand our results better. The subsequent algorithm was employed for these set of simulations:

1. Set initial population size N_0 and specify simulation parameters including carrying capacity, Allee threshold, function parameters for treatment-induced death and for

cell proliferation. Set seed.

2. Pick x values for all cells from a uniform distribution between 0 and 1. Grow the population till carrying capacity to obtain a steady-state distribution of x values in the population. Use this population as a staring point for therapy. Label all cells with lineage numbers (1,2,...).
3. Start treatment 1. Birth events result in two daughter cells that replace the mother cell, each with an x value sampled from a Gaussian distribution around the x value of the mother cell with a standard deviation of 0.02. Note the label of the parent cell to keep track of cell lineages. Death rates are affected by the treatment, depending on the level of resistance of an individual cell. After each demographic event, update the population size, cell labels and x values.
4. For the first run, do not apply the second strike. For subsequent runs, apply the second strike at a specified time point, which can be determined by the first run (which does not have a second treatment). Note the distribution of x before and after the second strike.

This algorithm allows us to track individual cell lineages to observe the impact of treatment on tumour heterogeneity. Additionally , we can also analyse the distributions of growth rates and x values at key time points throughout the simulation, allowing us to derive insights into the effect of the second strike on the population under this model.

The major limitation of this model is that it is very complex, without explaining the need for such complexity. There are many inter-connected parameters which do not have a physical interpretation and are therefore hard to translate to the laboratory. The parameter space for this model becomes very large and impractical to explore meaningfully. A secondary limitation is that these simulations are extremely time-consuming, even for small populations of sizes up to 10^3 cells. We were therefore unable to simulate larger populations in this framework. For the same reason, we (and the original authors) were unable to obtain actual extinction/relapse probabilities, which is an important component of a stochastic system like this one. Moreover, the lack of an analytical model leaves us with no other way to analyse these probabilities. In an attempt to simplify this model, while keeping the core concept of extinction therapy intact, we formulate our own framework including an analytical model and simulations, which we describe in the following sections.

3.2 Analytical Insights

The theory of evolutionary rescue (Section 2) from conservation biology provides us with a strong theoretical basis for an analysis of extinction therapy. We start with an

Description	Symbol
Initial tumour population size	N_0
Probability of fixation of a single R_i lineage	π_i
Probability of mutating to the R_i variant	p_i
Mutation rate	μ
Growth rate of S cells	g_S
Growth rate of R_i cells	g_{R_i}

Table 3.2: List of parameters for the analytical model. Note that the probability of mutating to a certain cell type is conditional on the mutation event occurring. Wherever relevant, $i = 1, 2$.

original permissive environment, which in our case will be the absence of treatment. The primary therapy (presence of the first drug) constitutes the first kind of stressful environment. Further strikes lead to different environments, possibly with varying levels of stress. The way we model the second (and further) strikes depends on the type of treatment administered and their method of action. Within this framework, we have multiple cases which lead to qualitatively different results. These preliminary models are described in Appendix A. Our final analytical model builds on the preliminary models.

Taking insights from cases 1-4 in Appendix A, we understand that a bottleneck therapy is not viable due to the nature of its effect on cells. An instantaneous elimination of a fixed fraction of the population is an unrealistic treatment method. Something similar can be achieved in case of a surgical removal of tumour, but to do so at very small population sizes (when the second strike is expected to be the most effective) is not possible. The trivial results of case 1 also demonstrates that the model is too limited to describe this system. From case 2, we understand that overlapping treatments 1 and 2 results in the emergence of a cell type that is resistant to both the therapies, which almost always leads to rescue. The nature of the results indicates that the treatment method is unsuitable for extinction therapy. Finally for this case, we consider two resistant cell types (R_1 and R_2) and one sensitive or wild-type cell type (S). However, this time we work with two non-overlapping treatments, each inducing an environment (E_1 and E_2). Cell type $R_{1/2}$ is resistant in $E_{1/2}$ but sensitive in $E_{2/1}$, while S cells respond to therapy in both the environments. Model results are given in Section 4.

With this model, we can calculate extinction probabilities for different parameter ranges. Moreover, it accounts for pre-existing mutants and the growth of R_1 before the second strike. We make a few assumptions here – 1) The second strike threshold $N_\tau \geq N_{\min}$,

which is the minimum population reached without a second strike, 2) the fraction of R_2 cells remains the same at time 0 and at N_τ, 3) The growth of R_1 cells before the second treatment is logistic, and 4) The population decays logistically during the second treatment. However, this framework allows for further analytical development in addition to being easy to simulate.

The limitations of this analytical framework arise due to the importance of accounting for stochastic dynamics at small population sizes, which is something that this model does not achieve. This is where the simulations add value. They account for the stochasticity of the system and allow us to make predictions that the analytical model cannot. We therefore simulate this same framework to gain further insights.

3.3 Stochastic simulations of our model

To test the hypothesis of Extinction therapy on the simple system described in Section 4.2.1, we consider two strikes. Each strike induces a different environment (E_1 and E_2) and possibly a different level of stress on the population. The total population at time t is $N(t) = N_S(t), +N_{R_1}(t) + N_{R_2}(t)$. We use a Gillespie-like algorithm to simulate this system. The code developed for its implementation is designed to be highly versatile and easy to modify. It can be extended to analyse more complex systems with more number of treatments and corresponding cell types.

3.3.1 Gillespie-like Implementation

The stochastic simulation algorithm (SSA), commonly referred to as the Gillespie algorithm (Gillespie 1977), is a Monte Carlo method initially devised in 1977 for modelling the temporal evolution of chemical reactions in a well-mixed system. It has since become an important tool in computational chemistry and systems biology. By simulating specific reactions or events given their rate of occurrence, the Gillespie algorithm mimics the time evolution of a complex system.

We implement a version of the Gillespie algorithm for simulating the population dynamics of our system in the context of Extinction therapy. The idea is that given a set of rates corresponding to birth, death and mutation events, we can track the population sizes of each subpopulation. These rates are specified for all the cell types and can change with time or in response to the environment. The basic steps of our algorithm are as follows:

1. Initialize the system by specifying an initial population for all cell types ($N_S(0), N_{R_1}(0), N_{R_2}(0)$) and setting the time to zero. Define all possible demographic events and

their corresponding rates.

2. Calculate the rate of any one event occurring, which is the sum of all individual rates $\omega_i(t)$. From this we can obtain the time interval after which the next event will take place. To do this, generate a sample from an exponential random variable with the rate parameter equal to the total rate ($\sum_i \omega_i(t)$). Alternatively, one can generate a random number z_1 from the uniform distribution between 0 and 1 and use the following formula to determine the time interval for the next event: $t_{\text{int}} = -\ln(z_1)/\sum_i \omega_i(t)$.
3. Calculate event probabilities for the next event by dividing the individual event rates by the total rate: $p_i(t) = \omega_i(t)/\sum_i \omega_i(t)$. Use these probabilities to select the next event by generating a random number z_2 between in 0 and 1. The chosen event k would be the largest j such that $(z_2 - \sum_{i=0}^{j} p_i(t)) > 0$.
4. Implement the chosen event by updating the population of the corresponding cell types. Increment time $t = t + t_{\text{int}}$. For example, if the chosen event is the birth of an R_2 cell, then $R_2(t_{\text{int}}) = R_2(t) + 1$.
5. Repeat steps 2-4 till a stopping condition is reached.

Note that the effects of carrying capacity and treatment are included while specifying birth and death rates (Section 3.3.2). All simulation parameters used to compute the individual event rates are listed in Table 3.3.

Simulations are stopped under one of three conditions: if the population goes extinct, if it exceeds the maximum simulation time, or if it exceeds some maximum population size. Similarly, the outcome of one run of a simulation can be one of three possibilities: extinction ($N(t) = 0$), progression ($N(t) \geq N(0)$) or persistence ($N(t) < N(0)$). Note that the stopping conditions do not have an equivalent simulation outcome (Table 3.1(b)).

3.3.2 Determining Demographic Event Rates

Following our variant of the Gillespie algorithm, one must define all possible demographic events at the beginning of the simulation and define rates corresponding to those events at each time step. Note that an individual event includes the type of event and the type of cell. For example, a mutation event $S \to R_1$ is one individual event and has a rate specified for it. Similarly, birth of an S cell is a different event than the birth of an R_1 cell.

We derive the birth and death rates for all cell types from the deterministic Logistical model for population growth.

Description	Parameter	Default value
Initial population sizes	$(N_S(0), N_{R_1}(0), N_{R_2}(0))$	$(10^4, 1, 1)$
Carrying capacity	K	10002
Per capita base birth rate of S cells	α_S	1
Cost of resistance	(c_{R_1}, c_{R_2})	$(0.5, 0.5)$
Per capita base birth rates of R_1 cells	$\alpha_{R_1} = \alpha_S - c_{R_1}$	0.5
Per capita base birth rates of R_2 cells	$\alpha_{R_2} = \alpha_S - c_{R_2}$	0.5
Per capita base death rates	$(\beta_S, \beta_{R_1}, \beta_{R_2})$	$(0.1, 0.1, 0.1)$
Per capita death rate due to treatment	$(D_S^{E_i}, D_{R_1}^{E_i}, D_{R_2}^{E_i}), i = 1, 2$	$(0, 1, 1)$
Mutation rate	μ	10^{-5}
Strike-2 threshold	N_τ	500

Table 3.3: Model parameters with default values adopted from Gatenby, Artzy-Randrup, et al. (2020). Per capita birth/death rates are specified for each cell type. Per capita death rates due to treatment are the same for all the resistant cell types, but depend on the environment (induced by the ongoing treatment). As default values, we keep the drug-induced death rate same in both the environments. The second treatment begins at the population threshold of N_τ.

$$\frac{dN(t)}{dt} = g(t)N(t); \qquad g(t) = r\left(1 - \frac{N(t)}{K}\right) - D^{E_i}, \tag{3.1}$$

where $g(t)$ is the per capita growth rate of the entire population at time t and r is the per capita base growth rate. The total treatment-induced death rate in environment i is D^{E_i} For a subpopulation within the entire population, the per capita time-dependent growth rate will be,

$$g_{\{S,R_1,R_2\}}(t) = r_{\{S,R_1,R_2\}}\left(1 - \frac{N(t)}{K}\right) - D^{E_1}_{\{S,R_1,R_2\}} \tag{3.2}$$

$$g_{\{S,R_1,R_2\}}(t) = (\alpha_{\{S,R_1,R_2\}} - \beta_{\{S,R_1,R_2\}})\left(1 - \frac{N(t)}{K}\right) - D^{E_1}_{\{S,R_1,R_2\}} \tag{3.3}$$

From Eq 3.3, we can separate the positive and negative parts to obtain the total birth and death rates,

$$b_{\{S,R_1,R_2\}}(t) = \alpha_{\{S,R_1,R_2\}} - (\alpha_{\{S,R_1,R_2\}} - \beta_{\{S,R_1,R_2\}})\frac{N(t)}{K} \tag{3.4}$$

$$d_{\{S,R_1,R_2\}}(t) = \beta_{\{S,R_1,R_2\}} + D^{E_1}_{\{S,R_1,R_2\}} \tag{3.5}$$

The compound birth and death rates for each cell type account for the effects of environmental factors like carrying capacity and treatment-induced death. We use Eqs 3.4 and 3.5 for our simulations.

For mutation events, we consider a constant mutation rate μ for all cell types. Once the source population is chosen (say, S cells), the target cell is chosen at random from the remaining cell types (R_1 and R_2). However, these mutation probabilities from one cell type to the other can be modified. This can be important in cases where back mutation rates or mutation rates between resistant cell types are much lower.

3.3.3 Switching to the Second Treatment

The treatment induced death rate depends on the cell type and the environment, which in turn changes with time. For this study, we consider two environments E_1 and E_2, each corresponding to the two strikes or treatments. Now we ask the question: how does the environment change with time? In other words, what is the condition under which we switch treatments and apply the second strike? This is a very important question because the hypothesis of Extinction therapy relies on the timing of the second strike. However, it is an indirect relation. Extinction therapy aims to exploit the stochasticity of a small population at the time of the second strike. This means that the switch between treatments depends on the population size. Therefore, we define the parameter N_τ – the population threshold at which we switch to the second treatment.

It is necessary to understand the behaviour of this threshold N_τ in order to evaluate the efficacy and limitations of Extinction therapy. Specifically, its relation with the minimum population reached in the absence of the second strike ($N_{\min}$) is important to explore. As reasoned by Gatenby, Artzy-Randrup, et al. (2020), the hypothesis is that the optimal N_τ (the N_τ at which probability of extinction is the highest) would be close to $N_{\min}$. To test this, we run a set of simulations with the same parameter values but with different seeds. For a single seed, we run multiple simulations, each with a different N_τ, at different distances from $N_{\min}$. This ensures that any differences due to stochasticity are eliminated. With 100 seeds for a single set of parameters, we calculate extinction probabilities for N_τ at various distances from the $N_{\min}$ (Fig 3.1(a)). The algorithm is as follows:

1. Select a set of parameter values, excluding N_τ. All other parameters need to be specified and are kept constant throughout.
2. Set seed for this set of simulations, thus eliminating differences due to stochasticity.
3. Run the first simulation without any second strike. Equivalently, set N_τ equal to zero for this run. From the results of this run, save the values of $N_{\min}$ and its

corresponding time point ($t_{\min}$).

4. Run the remaining set of simulations with $N_\tau = mN_{\min}$ where the factor m goes from 1 to a specified value (set to 20 for our simulations). Record the outcomes of the all the runs (according to Table 3.1(b)).
5. Repeat steps 2-4 for the desired number of seeds. Each set of simulations with a different seed is independent and will have different values of $N_{\min}$. This means that the absolute N_τ values will be different for each independent realisation. Therefore, in order to calculate extinction probabilities, we keep the factors m constant for all seeds.
6. Calculate extinction probabilities using outcomes for each seed, corresponding to the values of m. We thus obtain extinction probabilities for various values of $N_\tau/N_{\min}$.

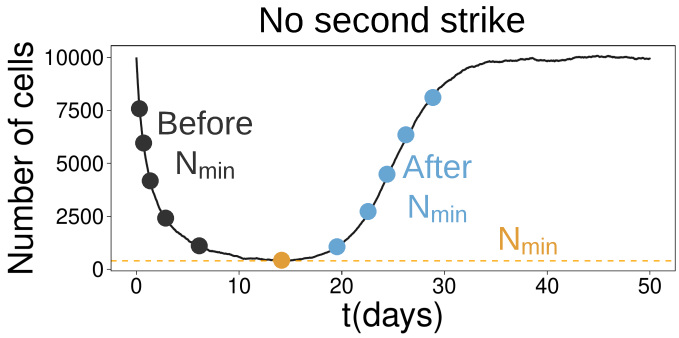

(a)

Stopping condition	Outcome condition	Outcome
$N(t) = 0$	$N(t) = 0$	Extinction
$t \geq T$	$N(t) \geq N(0)$	Progression
	$N(t) < N(0)$	Persistence
$N(t) \geq \min(0.99K, N_{\max})$	$N(t) \geq N(0)$	Progression
	$N(t) < N(0)$	Persistence

(b)

Figure 3.1: **(a)** Illustration of the set of simulations to find the relation between $N_{\min}$ and the optimal N_τ. Shown is the first run without any second treatment, which defines $N_{\min}$ as the minimum population size before relapse. The factor $m = N_\tau / N_{\min}$ goes from 1 to 20 on both sides of $t_{\min}$: before reaching $N_{\min}$ and after crossing it. Subsequent runs take these marked points as N_τ and record the outcome in each case. **(b)** Stopping conditions and corresponding outcomes of a simulation. In the second condition, T is the maximum simulation time defined at the beginning. Generally, if we see a significant number of outcomes with persistence (more than 10%, for example), it means that T is not large enough. In the third condition, the threshold $0.99K$ is arbitrary. The outcome remains the same as long as the threshold is close to K. We take the minimum of two quantities for cases where K is much larger than the initial population size, and a threshold of $N_{\max}$ is enough to declare the outcome. Typically, K and $N_{\max}$ are greater than $N(0)$, so persistence in the last condition is not observed. However, the condition is mentioned for generality.

Chapter 4

Results

In this section, we will review our results and briefly discuss the inferences and ramifications. Starting with the reproduction of the existing model, we get some new insights into the dynamics of cell lineages during extinction therapy, as described in Gatenby, Artzy-Randrup, et al. (2020). We derive analytical results from our model and compare those findings to simulation results. We discuss the significance of these approaches as well as their advantages and disadvantages. We then turn to the simulation results and their in-depth examination. Lastly, we model errors in genetic barcoding experiments, a technique expected to be used extensively in *in-vitro* experiments of extinction therapy.

4.1 Results and observations from the reproduced model

We simulated the model described in the paper by Gatenby, Artzy-Randrup, et al. (2020), attempting to reproduce their results in order to understand the population dynamics better. However, we were compelled to do so only based on the details provided in the paper, without their simulation code. We were able to obtain qualitatively similar results but with different parameter values. Our results are shown in Figure 4.1. Original simulation parameters are given in Table 3.1.

We make three major observations – first, that Allee effect is essential for the second strike to work in this model. This is something that the authors of the original paper have not explored. Although the results depend heavily on the specific combination of parameters, in our exploration of the parameter space, there were no cases in which the second strike led to the expected result in the absence of Allee effects. Therefore, we demonstrate that Allee effects affect the outcome of the therapy in this model. Second, there is a significant shift in the mean resistance values (mean x) in the

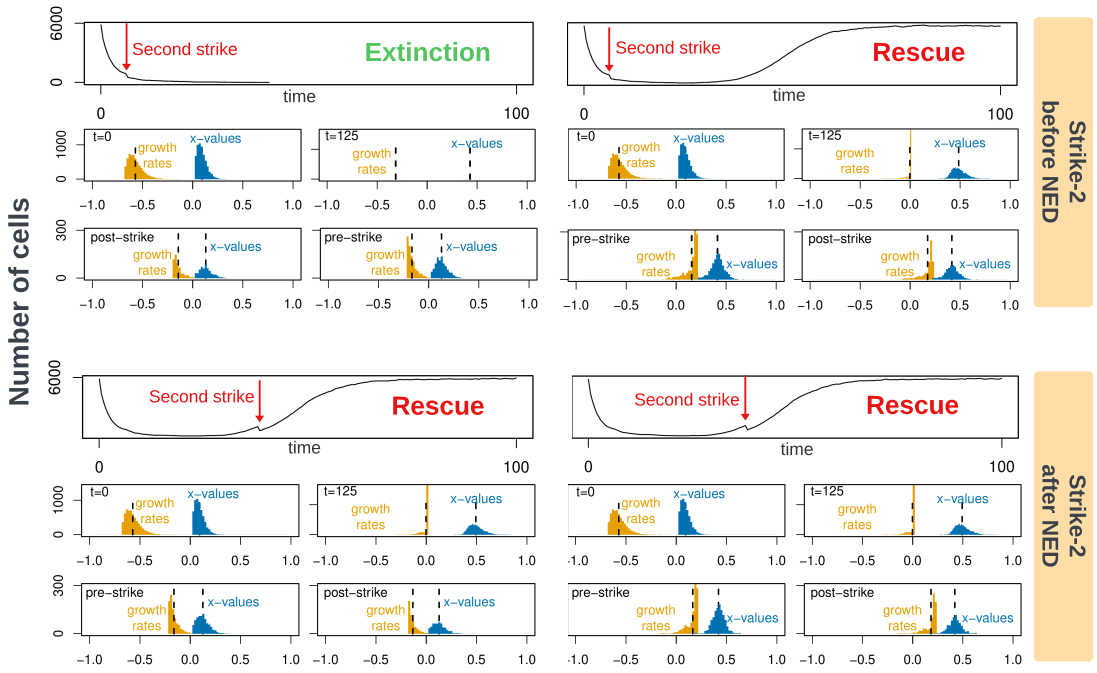

Figure 4.1: Demonstration of the model presented in Gatenby, Artzy-Randrup, et al. (2020). Each histogram shows the distribution of growth rates and resistance values of the population at the given time point. Pre(post)-strike time points show the distributions right before(after) the strike is implemented. The black dashed lines indicate the respective means. We use different parameter values than the original model to get qualitatively same results. This was done because the model and values, as described in the paper, did not produce the same results. Possible explanations are mentioned in the text. The parameter values are chosen in order to illustrate the contrast between different treatment combinations. We observe that the only instance in which extinction therapy seems to work is the case with Allee effect in which the second strike is applied before (or as soon as) the population reached the NED (No Evidence of Disease) threshold. Additionally, we notice the shift in mean x in relapsed populations due to selection of highly resistant cells. Lastly, we also see a shift in the distributions between cases where the strike is applied before and after NED. Growth parameters: $\lambda_{\min} = 0.2, \lambda_{max} = 1, s = 0.25, \mu = 0.1$. Population level parameters: $A = 15, a = 0, K = 6000, N_0 = 6000$. Treatment related parameters: $C = 0.4, \beta = 1.5, \alpha = 0.8$

population in cases when the population relapses. This, again depends on the specific treatment related parameters used for this simulation. Third, we notice that the growth rate distributions are skewed in different directions when the second strike is applied before and after the NED. This demonstrates the exact mechanism of relapse and the importance of the timing of the second strike. If the population is struck while the growth rates are skewed to the left and the mean x value is low, the population has little chance of being rescued. If the strike is after the growth rates are already skewed to the right, or equivalently the mean x value is high, then the tumour will likely relapse.

For these results, we used different parameter values than those given in the original paper because with the given description, we were unable to reproduce their results. The treatment related parameters were altered to give a stronger treatment. Specifically, β was increased from 0.2 to 1.5. This was necessary because the population did not show decrease in size, as expected from the results in the original paper. Additionally, the method to obtain the initial steady-state distribution was not detailed in order to be reproduced. We used a two-sample Kolmogorov-Smirnov test to determine convergence to steady state. The distribution obtained with this algorithm, although visually similar to the one in the original paper, did not produce the same results. From a preliminary analysis, it seemed that the tail of the x distribution needed to be quite thin in order to lead to extinction. A Beta(3,30) distribution was therefore used to approximate this form. This allowed us to analyse the results with a range of initial x distributions.

Following an analysis of various starting tumour cell populations (both from steady-state analyses and beta distributions), we observed that the outcome of the treatment is very sensitive to the distribution of x values (resistance values) in the beginning. If the number of cells with $x > C$ is high, then the evolution of resistance is fast and the population does not reduce in size enough for the second strike to work. Further, the initial distribution of x values in the population is determined by the cost of resistance parameter s. This parameter is not a direct additive cost to the intrinsic proliferation rates, but an exponent which indirectly determines the cost of resistance (formula given in Section 1.4). The higher the value of s, the higher the cost of resistance.

Practically, in an actual tumour it is hard to find the distribution of costs of resistance. This becomes an issue because a parameter which plays an important role in determining therapy outcomes cannot be easily predicted or monitored. To classify a cancer as suitable for extinction therapy, we would need a proxy to indirectly determine the effect of the costs of resistance in a given population. To that end, we added a function to track the lineages of each cell in the initial tumour population in our simulations. The motivation for this was to perform predictive modelling for future experiments in extinction therapy. Proposed experiments include barcoding the initial cell population

(see Section 4.4 in which we mathematically model experimental errors in barcoding experiments) to track their dynamics and categorise populations that are likely to respond well to extinction therapy. Focusing on differentiating cases with different cost of resistance distributions, an analysis of cell lineages in our simulations shows that the statistics of growth rates across cell lineages can help inform the ultimate outcome of therapy (shown in Figure 4.2).

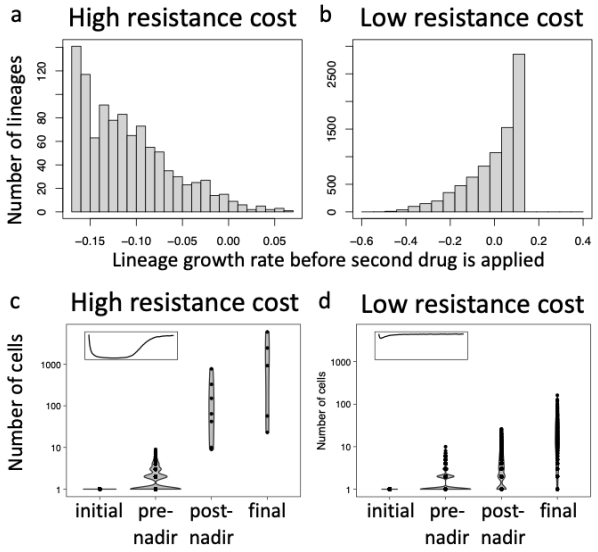

Figure 4.2: A comparison between lineage dynamics of systems with high and low resistance costs. **(a)** In the high resistance case, the lineage growth distribution before second strike is skewed to the left, indicating that there are few lineages with a positive growth rate (resistant lineages) while in **(b)**, the low resistance case, the distribution is skewed to the right indicating that there are many resistant lineages. **(c,d)** Experimentally, one can classify a system suitable for extinction therapy by observing the number of lineages in a relapsed population. A small number of lineages with the same total population size means that the bottleneck due to treatment was very small, in which case extinction therapy is likely to lead to extinction.

To demonstrate the effect of varying s, we compared two sets of model parameters with different outcomes post-therapy due to different trade-offs between resistance level and proliferation rate. Case 1 leads to a successful treatment outcome if the second treatment is given before the population reaches the nadir (minima, or equivalently, before NED), while case 2 does not. These parameter values are chosen such that they represent two extreme and contrasting behaviours of the initial resistance trait value distribution. The primary therapy is constant throughout the treatment schedule.

A high fitness cost of resistance leads to a left-skewed lineage growth rate distribution (Fig. 4.2a), such that highly-resistant (positive growth rate) lineages are extremely rare by the time the second strike is applied. In this scenario, the primary treatment reduces the tumour to a very small size and subsequent application of the second strike is

expected to lead to extinction. In contrast, a weaker cost of resistance leads to a right-skewed lineage growth rate distribution (Fig. 4.2b) and treatment failure. In this case, the scales are tipped the other way, and there are enough high-resistance lineages that the population size does not decrease much.

When the second strike fails, our simulations predict that the nadir population size can be inferred from the final lineage frequency distribution. Each lineage originates from a single cell at $t = 0$. A narrow bottleneck (low nadir, as expected in the high resistance cost scenario) shifts the distribution towards a small number of lineages with a lot of cells (Fig. 4.2c), whereas a wide bottleneck preserves a multitude of smaller lineages (Fig. 4.2d). In an experiment, this method can help determine the feasibility of extinction therapy for a given population. To summarise, a tumour relapse which consists of only a few lineages of large sizes, indicates a system with high cost of resistance, and therefore relatively higher sensitivity to extinction therapy. On the contrary, a relapsed population that consists of many small lineages indicates a low cost of resistance system, which is expected to be less responsive to extinction therapy.

With this limited analysis of the existing model, we realise the need for a simpler model with fewer degrees of freedom. We require a way to calculate extinction probabilities and to meaningfully explore parameter spaces in the context of reasonably large cancer populations.

4.2 Analytical results

4.2.1 The model

We consider two collaterally sensitive drugs, each corresponding to a resistant cell type (R_1 and R_2). A simplifying observation we make here is that ultimately, any case of evolutionary rescue will be due to R_2 cells if we switch to the second treatment at any point of time. If R_1 rescues the population, then at some point, due to treatment 2 (second strike) it will go extinct and further rescue by the R_2 variant will be subject to the environment at that time. In this framework, we can also take into account standing genetic variation. We are therefore left with two conditions to consider – rescue due to pre-existing R_2 cell that persist till the onset of the second treatment and new rescue variants that are generated during the second treatment. Here, we make an assumption that any R_2 lineages generated during the first treatment go extinct due to a negative growth rate. All model parameters are listed in Table 3.2.

Assuming that the initial number of R_2 cells is given by $N_{R_2}(0)$, we can estimate the

expected number of these cells remaining at time τ.

$$\mathbb{E}[N_{R_2}(\tau)] = \frac{N_{R_2}(0)N_\tau}{N_0} \tag{4.1}$$

We obtain this expression by assuming that the R_2 cells are equivalent to sensitive cells in the first environment. We know that the distribution of pre-existing rescue variants is Poisson with a rate equal to $\lambda_1 = \pi_2\mathbb{E}[N_{R_2}(\tau)]$, with which we can calculate the probability that all these pre-existing variants go extinct in E_2. Here, π_2 is the probability of fixation of an R_2 lineage starting from one cell (see Section 2). Note that the number of R_2 cells at time τ are treated as standing genetic variance for treatment 2.

For *de-novo* mutations in E_2, we have the rate of generation of R_2 cells

$$\lambda_2 = \pi_2\mu p_2 \int_\tau^{t_{\text{ext}}} N_t dt \tag{4.2}$$

This rate is conditional to the population going extinct at time t_{ext}. With the assumption of Logistic growth, we can make the substitution $N_t = (dN_t/dt)(1/g(t)(1 - N_t/K))$:

$$\lambda_2 = \pi_2\mu p_2 \int_{N_\tau}^{0} \frac{dN_t}{g(t)(1 - N_t/K)} \tag{4.3}$$

where $g(t)$ is the growth rate of the population as a function of t. The growth rate is not constant in this case because as the sensitive cell population declines, the R_1 population increases. The growth rate of the entire population at any point of time would be dependent on the fraction of R_1 cells in the population.

$$g(t) = \frac{N_{R_1}(t)}{N_t} g_{R_1} + \frac{N_S(t)}{N_t} g_S \tag{4.4}$$

where $g_{R_1/S}$ is the growth rate of R_1/S cells in the first environment. We do not need include the growth rate of R_2 cells because before the onset of the second treatment, their population is negligible compared to R_1 and S cells. To find the fraction of R_1 cells in the population at time τ, we look at the growth equations for both the cell types.

$$\dot{N_{R_1}}(t) = g_{R_1} N_{R_1}(t)(1 - N_t/K) \quad \text{and} \quad \dot{N_S}(t) = g_S N_S(t)(1 - N_t/K) \tag{4.5}$$

After simplifying and eliminating the density dependence term, we integrate twice from 0 to τ,

$$\frac{\dot{N_{R_1}}(t)}{g_{R_1}N_{R_1}(t)} = \frac{\dot{N_S}(t)}{g_S N_S(t)} \tag{4.6}$$

$$\implies \int_0^\tau \frac{1}{g_{R_1}} \frac{d\ln N_{R_1}(t)}{dt} dt = \int_0^\tau \frac{1}{g_S} \frac{d\ln N_S(t)}{dt} dt \tag{4.7}$$

$$\implies \frac{N_{R_1}(\tau)}{N_{R_1}(0)} = \left(\frac{N_S(\tau)}{N_S(0)}\right)^{g_{R_1}/g_S} \tag{4.8}$$

$$\implies N_\tau = N_{R_1}(\tau) + N_S(\tau) = N_{R_1}(\tau) + N_S(0)\left(\frac{N_{R_1}(\tau)}{N_{R_1}(0)}\right)^{g_S/g_{R_1}} \tag{4.9}$$

We get the expected number of R_1 cells at time τ as a function of the total population at time τ. Since Eq 4.9 is an implicit equation, we solve it numerically to get the growth rate at population size N_τ. Using this result, we can now compute the extinction probabilities.

$$-\ln P_E = \lambda_1 + \lambda_2 \tag{4.10}$$

$$= \frac{N_{R_2}(0)}{N_0} N_\tau \pi_2 + \pi_2 \mu p_2 I \tag{4.11}$$

$$\tag{4.12}$$

Here, we have replaced the integral term in Eq 4.3 with I. Since explicit calculation of the integral I is complicated and unnecessary, we obtain an expression to find the rate of change of the probability of extinction (P_E) with N_τ.

$$-\frac{d\ln P_E}{dN_\tau} = \frac{N_{R_2}(0)}{N_0}\pi_2 + \pi_2 \mu p_2 \frac{dI}{dN_\tau} \tag{4.13}$$

$$= \frac{N_{R_2}(0)\pi_2}{N_0} - \frac{\pi_2 \mu p_2}{g(\tau)(1 - N_\tau/K)} \tag{4.14}$$

Given a boundary condition, we can calculate P_E for different values of N_τ and find the optimal N_τ in this scenario. It is important to note that Eq 4.13 is only valid for the range of N_τ values that are practically possible. The lower limit to the range of N_τ values is $N_{\min}$, the minima reached by the population without any second treatment. Any value below $N_{\min}$ cannot be a threshold to switch to the second treatment because the population will never reach that threshold. To determine this range, we need to calculate $N_{\min}$. With the assumption that the R_2 population is negligible before the second strike, we know that the total population before reaching $N_{\min}$ has a negative

growth rate $g(t)$. After $N_{\min}$, the population will have a positive growth rate. Therefore, we need to find the population at which the growth rate of the population (including S and R_1 cells) becomes zero. Taking the frequency of R_1 cells in the population at the minima to be q, we solve:

$$g(t) = 0 \implies qg_{R_1} + (1-q)g_S = 0 \tag{4.15}$$

$$\implies q(g_{R_1} - g_S) = -g_S \tag{4.16}$$

$$\implies q = \frac{g_S}{g_S - g_{R_1}} \tag{4.17}$$

At this frequency, using Eq 4.9, we can get the value of $N_{\min}$:

$$N_{\min} = qN_{\min} + N_S(0)\left(\frac{qN_{\min}}{N_{R_1}(0)}\right)^{g_S/g_{R_1}} \tag{4.18}$$

Dividing by $N_{\min}$ throughout and simplifying, we get,

$$N_{\min} = \left(\frac{(1-q)}{N_S(0)}\left[\frac{N_{R_1}(0)}{q}\right]^{g_S/g_{R_1}}\right)^{g_{R_1}/(g_S - g_{R_1})} \tag{4.19}$$

$$= \left(\frac{g_{R_1}}{N_S(0)(g_{R_1} - g_S)}\left[\frac{N_{R_1}(0)(g_S - g_{R_1})}{g_S}\right]^{g_S/g_{R_1}}\right)^{g_{R_1}/(g_S - g_{R_1})} \tag{4.20}$$

We note that the value P_E at $N_\tau = 0$, according to the model, is 1. This means, trivially, that the population is sure to go extinct if the second strike is applied when the population goes extinct. This becomes a convenient boundary condition to solve the differential equation for P_E.

Another crucial observation relates to the extrema of the P_E function against N_τ. In Eq 4.13, if we equate the differential to zero, we get,

$$\frac{N_{R_2}(0)}{N_0} g(\tau)\left(1 - \frac{N^*_\tau}{K}\right) + \mu p_2 = 0 \tag{4.21}$$

$$N^*_\tau = K\left(1 - \frac{\mu p_2 N_0}{N_{R_2}(0)|g(\tau)|}\right) \tag{4.22}$$

Assuming $g(\tau)$ to be negative, which is a valid assumption to make before the population reaches $N_{\min}$, we obtain the expression of N_τ. Substituting the default parameter values from Table 3.3, we find that in most cases, the value of N^*_τ is close to K. To

determine the nature of the extrema, we differentiate Eq 4.13 again with respect to N_τ to see that the resulting expression is negative, implying that $-\ln P_E$ reaches a maximum at this point and therefore the probability of extinction is minimum at N_τ^*. With this information, we conclude that P_E decreases as N_τ increases, with a maximum at $N_\tau = 0$ and a minimum at N_τ^*. However, since the relevant values of N_τ are greater than $N_{\min}$, for our purposes, the maximum P_E is expected to be at $N_{\min}$.

Building on the results in this section, we analyse and compare them with the simulation results. Note that the simulations have the same framework as the analytical model, using the same parameter values wherever applicable. To begin, we calculated the probability of extinction for different values of N_τ by numerically solving the differential equation given in Eq 4.13, using the boundary condition that $P_E{=}\ 1$ at $N_\tau{=}\ 0$. The results are shown in Figure 4.3. We study the trend of P_E with N_τ, over a range of initial population sizes (N_0) and treatment values (D). The treatment parameter is equal to the death rate induced by the treatment, and is kept the same for both treatments. Major observations and results from the analytical model are described in the following sections.

4.2.2 Significance of $N_{\min}$ in simulations and analytical results

As seen in Figure 4.3, the behaviour of P_E for $N_\tau{<}N_{\min}$ is very different from the expected trend from analytical calculations. This is because these thresholds are never reached by the population, so the outcomes are the same as would be without a second strike (only primary treatment throughout). For this reason, in the analytical model, the expression for P_E holds only for values of $N_\tau{>}N_{\min}$. The starting condition of $P_E{=}\ 1$ for $N_\tau{=}\ 0$ is interpreted as the probability of extinction if the second strike is applied when the population has already gone extinct. Because this is not possible due to evolutionary rescue of the population by R_1 cells, the simulations show a low probability of extinction in these cases.

Ideally, if the population is large enough, these probabilities will be zero. However, in a few cases with low population sizes (Figure 4.3A,B,C), we see non-zero P_E below the $N_{\min}$ threshold. This is because of the stochastic nature of population dynamics. Every time we run an independent simulation, the value of $N_{\min}$ is different, more so in smaller populations. This leads to a gradual decrease in P_E in this range. As we increase the population size to 1 million cells, we see a sharp decrease in P_E for N_τ below $N_{\min}$, as expected.

On the other side of $N_{\min}$, we see a decrease in P_E, which provides evidence for the hypothesis that the maximum P_E is reached at $N_{\min}$. To verify the relative positions of

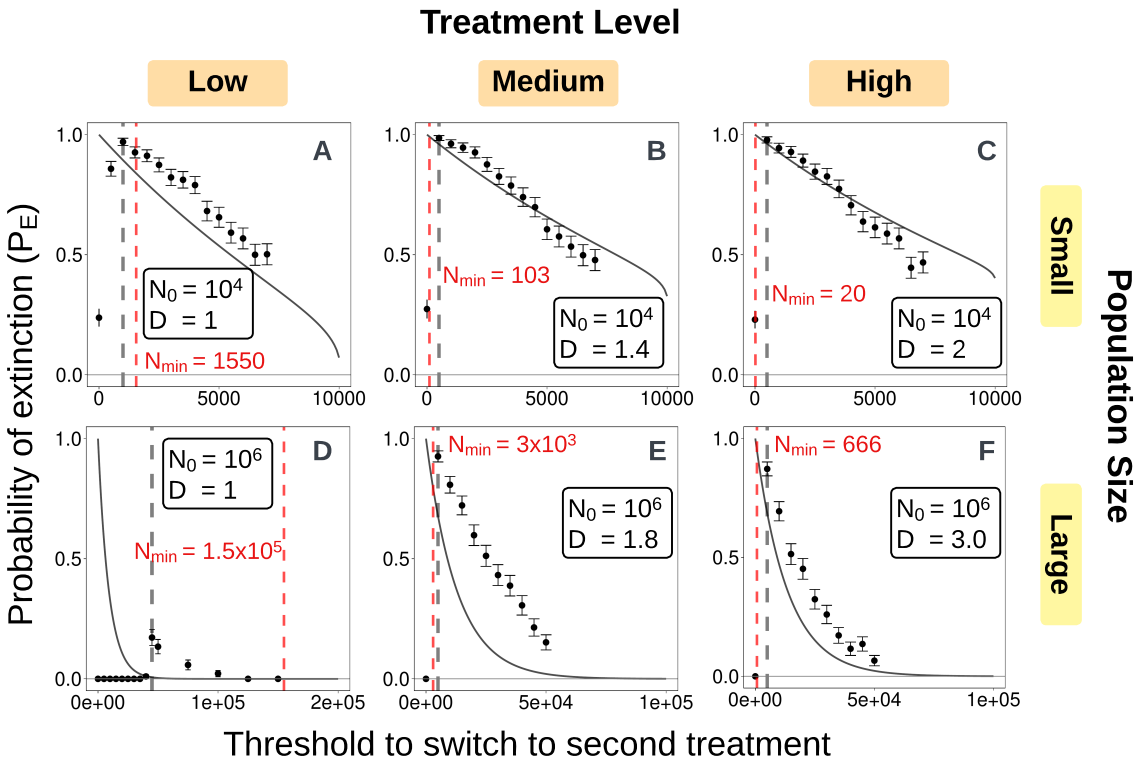

Figure 4.3: Analytical predictions compared to simulation results for low, medium and high treatment levels for both small and large populations. Each plot shows the trend of P_E with varying N_τ. The initial population sizes (N_0) and treatment parameters (D) are mentioned for each case. Black dots are simulation results and the grey solid line indicates analytical predictions. Error bars show 95% binomial proportion confidence intervals. The analytical expectation of the minimum population reached in absence of second treatment ($N_{\min}$) is shown with a dashed red line, with values mentioned alongside. Grey dotted lines indicate the second strike thresholds at which the extinction probability is maximum, as observed in the simulations. From our simulations, this optimal N_τ is shown to be a good estimate of the empirical $N_{\min}$ (see Section 4.3.2). All deviations from expected P_E values are reasoned in the text with possible explanations.

$N_{\min}$'s in the simulations, we run a separate set of simulations that compare $N_{\min}$ with other values of N_τ to verify that the optimal N_τ is $N_{\min}$ (Section 4.3.2). Therefore, we use the optimal N_τ as an estimate of the empirical $N_{\min}$ in Figure 4.3 and Figure 4.4.

4.2.3 The analytical model provides reliable estimates of the optimal second strike threshold in most cases

Our analytical model provides a way to calculate the optimal N_τ (or $N_{\min}$) using only the growth rates and initial conditions of the population (refer to Eq 4.20). We compare this analytical estimate to the observed optimal N_τ's in the simulations in Figure 4.4A,B. We see that the analytically calculated $N_{\min}$ describes the observed behaviour well for most of the treatment levels, with the exception of very low treatment levels ($D = 1$).

Note that only absolute values of treatment-induced death rates (D) in the population are shown in the figures, but the overall effect of treatment is relative to the intrinsic growth rates. Since the intrinsic growth rate of treatment-sensitive (S) is taken to be 0.9 by default (Table 3.2), $D = 1$ is a low treatment level. This is also why lower values of D are not shown. For higher levels of treatment, we see good agreement in the expected and observed values of optimal N_τ. Therefore, our analytical model can be used to easily calculate the optimal second strike threshold which leads to the maximum extinction probability.

4.2.4 Minimum population size ($N_{\min}$) is higher than expected for low treatment levels

For larger populations, we observe that the $N_{\min}$ does not correspond to the maximum P_E when the treatment level is low (Figure 4.3A,D and Figure 4.4A,B). A hypothesis to explain this behaviours is that the R_1 cell population does not grow as fast as expected in these cases. Since the contribution of the R_1 population is what ultimately allows for relapse, the value of the minima depends on how fast it is able to grow large enough such that the total growth rate becomes positive.

In our analytical model, we assume logistic growth for R_1 subpopulation from the beginning of the primary therapy. This assumption may not be valid for cases in which the sensitive cell population decays slowly so that the R_1 lineages take more time than expected to escape stochastic extinction. Our analytical framework does not account for this extra time because we do not include the stochasticity of the R_1 subpopulation while it is small. This results in a higher $N_{\min}$ than observed in the simulations.

4.2.5 Extinction probability is higher than expected for large population sizes

In large populations, we see that the observed extinction probability is higher than what is estimated using the analytical model. While the broad trend is estimated well, the analytical model systematically underestimates the extinction probabilities for all treatment levels starting with $D = 1$. A possible hypothesis to explain this could be the error in calculation of R_2 cells at the second strike threshold (N_τ). Since our model assumes that the proportion of R_2 cells remains constant throughout the primary therapy, it ignores stochastic death due to the small abundance of R_2 cells. The starting population of the resistant cells is very small, and is subject to stochastic extinction. This effect results in a lower number of R_2 cells at the beginning of the second treatment, and consequently, a higher rate of extinction.

However, this may be just one of the reasons for the observed deviations from expected trends. This does not explain why the deviation decreases as the treatment level is increased (Figure 4.3E,F). We hypothesize that in cases in which the error in estimation of R_2 cells at the beginning of the second strike does not play a big role, we see less deviation from expected behaviour. These may be the cases in which our approximation holds or cases in which the observed extinction probability is low due to other factors. For instance, the same deviations are not observed for small population sizes. The observed trend of extinction probabilities in that case is well estimated by the analytical predictions. A possible reason for that could be that even though the same factors affect growth dynamics in all cases, in smaller population sizes, the dynamics are much faster than in larger populations ($N_{\min}$ is reached earlier, for example) so these effects do not manifest such that we can observe them.

4.2.6 Conditions for feasibility of analytical predictions

We observe that the analytical predictions match with the simulation outcomes only under certain conditions. There are many factors at play, some of which are not accounted for by our analytical model. The major limitation of our model is that we do not include the effects of stochasticity on small populations. Our assumption of Logistic growth and decay also does not hold completely in some cases. Most of the deviations from the predicted behaviour can be explained if we include these effects. Even so, our model manages to predict $N_{\min}$ values within acceptable error bounds (barring cases with very low treatment levels). Moreover, the overall trend of extinction probability varying with N_τ is also estimated well for small N_τ values and higher treatment levels.

As a result, we can carve out a region of feasibility for our model. Generally, we can

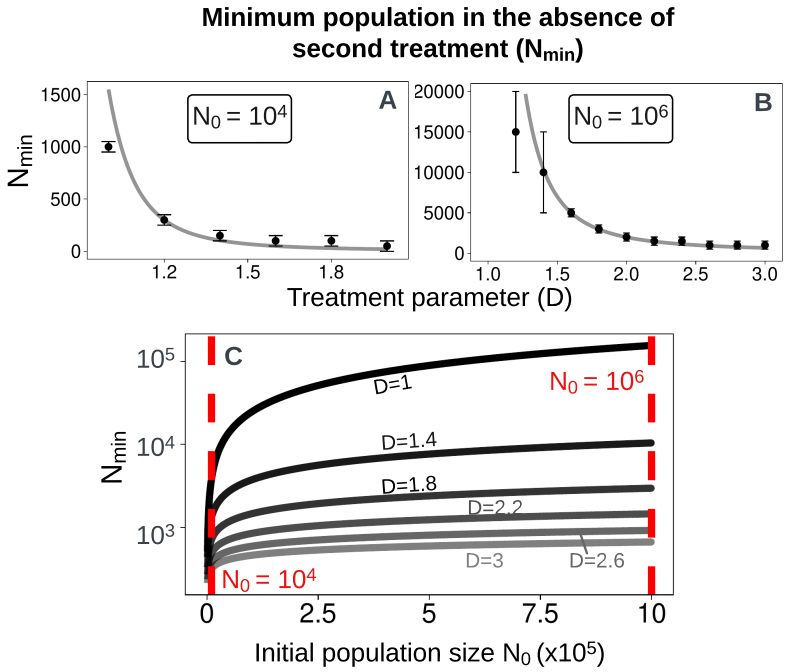

Figure 4.4: Behaviour of $N_{\min}$ with N_0 and D. **(A,B)** Comparison of observed (black dots) and expected (grey line) $N_{\min}$ for a range of treatment values. This range of D is different for small and large population sizes. In both the cases, values below 0.9 are not considered because they result in a net positive growth rate. Values of D above 2 are not considered for the smaller population size because it leads to extinction without the need of a second strike ($N_{\min}$ close to 0). These cases are not of interest to us because our aim is to prevent relapse in cancer populations using a second strike. The error bars in the observed $N_{\min}$ indicate least count errors. The observed values of $N_{\min}$ are estimated by the optimal N_τ in each case, which is the the threshold at which the extinction probability is the maximum. **(C)** Change in $N_{\min}$ with increase in initial population size. As the treatment parameter is increased, $N_{\min}$ decreases for all starting population sizes. The red dashed lines indicate populations at which variation of P_E with N_τ is shown in Fig 4.3. The lower region of this plot is the region of feasibility for $N_{\min}$ predictions using the analytical model.

expect the analytical estimate for $N_{\min}$ or optimal N_τ to be accurate except in cases with very low treatment levels. In Figure 4.4C, the lower region of the plot corresponds to the parameter values for which our analytical model can provide reliable estimates of the optimal N_τ.

The drawbacks of the analytical model prevent reliable exploration of certain regions of the parameter space. Especially, the exact change in extinction probability as the second strike threshold changes is not accurately predicted by the model. We therefore simulate this model for wide ranges of various parameters, in order to obtain a complete understanding of the system in context of extinction therapy. Additionally, simulations allow us to verify our hypothesis that P_E peaks at $N_{\min}$, while accounting

for the stochasticity in the system.

4.3 Stochastic simulation results

Using the methods outlined in Section 3.3, we simulate our model for two initial population sizes – 10^4 and 10^6. To begin with, we take all relevant parameter values from the existing model to aid comparison. Building on that, we explore more of the parameter spaces for both the population sizes and compare them. The difference between the two systems shows the importance of a stochastic growth component needed in our model. The main results are described in this section.

4.3.1 Second treatment increases the probability of extinction

As a proof of concept, we demonstrate that switching to the second treatment (the second strike) works. We see an expected increase in the probability of extinction when a second strike is applied. To compare with the original model in Gatenby, Artzy-Randrup, et al. (2020), we also simulate the system with model parameters directly taken from their paper, but with relevant changes to make it suitable for our framework. For example, the original model has a treatment-induced death rate which is a function of the resistance values. The maximum value of this function is taken as the treatment level for our simulation. Figure 4.5 demonstrates the effect of extinction therapy.

We note here that since the staring conditions of both the resistant cell types are the same, $N_\tau = 0$ is equivalent to $N_\tau = N_0$ in terms of extinction probability. Switching from treatment 1 to 2 is equivalent to switching from treatment 2 to 1. After this proof of concept, we show the effect of change in (primarily) N_τ and other parameters.

4.3.2 The optimal threshold for second strike is equal to the minimum population

One of the results of the paper by Gatenby, Artzy-Randrup, et al. (2020) is that the optimal threshold for the second strike is close to the minimum population reached without a second strike. To verify this, we ran simulations on the same system with different values of the second strike threshold (See Section 3.3). Due to the stochastic nature of the simulations, $N_{\min}$ is different for each independent simulation. We therefore analyse the extinction probability P_E as a function of the factor $N_\tau/N_{\min}$. The hallmark U-shaped curve of a rescued population means that this factor will give two instances of each N_τ – one before the minima and one after. Figure 4.6 shows both these classes of thresholds relative to $N_{\min}$.

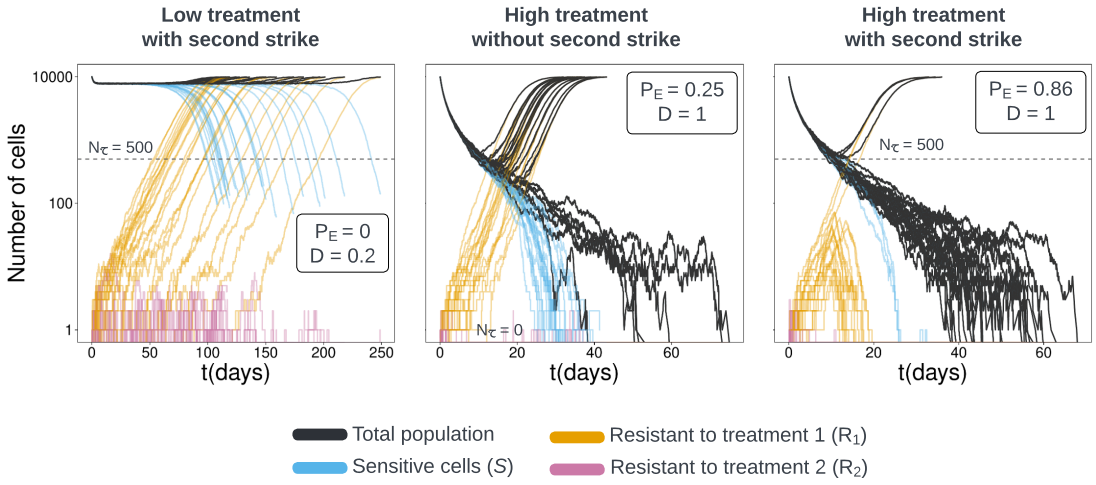

Figure 4.5: Simulations showing the effect of a second strike. The case with low treatment takes parameters from the original model. In the case with high treatment, a second strike significantly the extinction probability (P_E). The N_τ threshold is the same as in the original model. Plots show 25 out of 500 total runs. Persistence probabilities in all cases is equal to zero. Other parameter values are the default ones specified in Table 3.3. Note that the treatment levels in both the environments (under both the treatments) are the same, i.e. $D_{E_1} = D_{E_2} = D$.

We observe that the extinction probability is consistently the highest for $N_\tau{=}N_{\min}$, both for small and large populations, thus verifying the hypothesis presented by previous models. Moreover, we notice a significant difference in extinction probabilities for same N_τ's before and after the minimum population is reached. The extinction probability is higher if the switch to the second treatment is made after the minimum population has been crossed. This is because the growth rate of the relapsed is less than the initial population (before the minima). Since the relapsed population is dominated by resistant variants (R_2), it suffers from a cost of resistance which makes it easier to push the population to extinction. This result is robust to changes in the treatment level during the second treatment. In Figure 4.6D, the treatment-induced death for the second treatment is lower than the first treatment. This results in a lower P_E overall, but the extinction probabilities for N_τ's after the minima are still much higher than the ones before the minima.

An inference of this result is that it is safer to initiate the second strike after the population has reached its minima, not before as suggested by previous work. We get a "window of opportunity" which is skewed to the right.

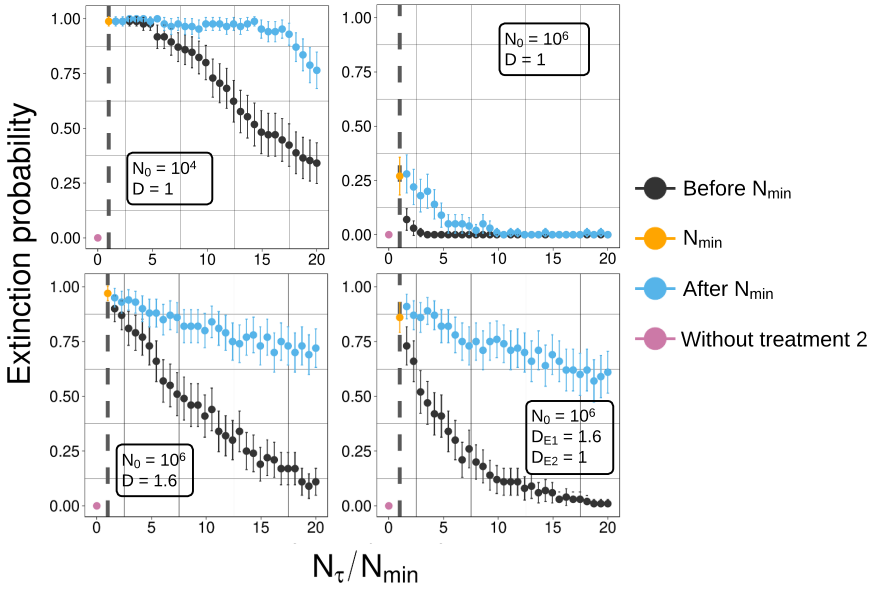

Figure 4.6: Results of simulations to determine the relationship between the second strike threshold (N_τ) and the minimum population reached without a second treatment ($N_{\min}$). Extinction probability is shown for many values of $N_\tau/N_{\min}$ and for two initial population sizes. The maximum P_E is consistently at $N_\tau{=}N_{\min}$ (indicated with a grey dashed line), implying that the optimal N_τ is equal to $N_{\min}$. Additionally, N_τ's after the minima are better than N_τ's before the minima has been reached. Note that for panels A,B and C, the treatment-induced death (or treatment level) is the same in both the environments (under both the treatments), i.e. $D_{E_1} = D_{E_2} = D$. For panel D, we consider two different values of treatment levels in the two environments, but the general trend remains the same.

4.3.3 Optimal threshold for second strike increases with carrying capacity

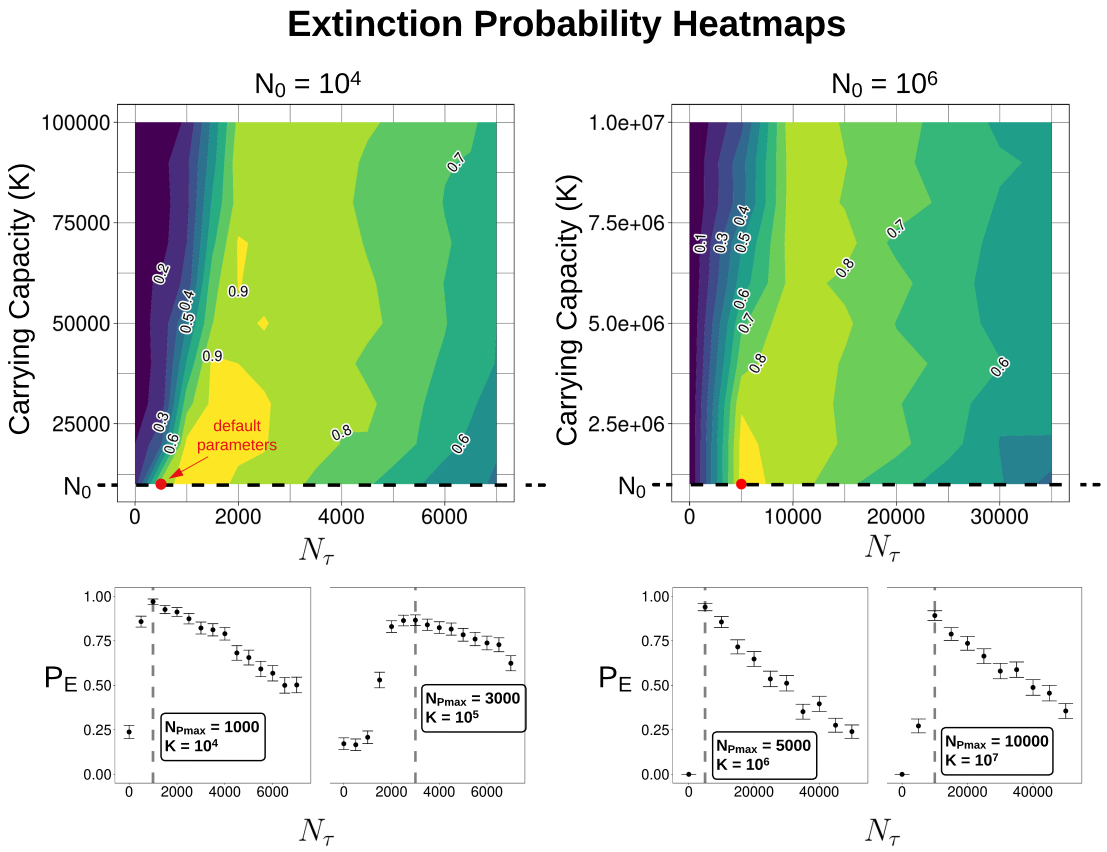

Figure 4.7: Heatmaps and scatter plots for extinction probabilities over a range of N_τ and K values. We see an increase in the value of the optimal N_τ as the carrying capacity is increased. The scatter plots at the bottom show this increase for two specific K values in each case. The default simulation parameters are marked with red dots on the heatmaps. Grey dashed lines indicate the empirical optimal N_τ in each case, as observed in the stochastic simulations.

Focusing on N_τ as our main parameter of interest, we simulate how our model responds to changes in N_τ with other parameters. The carrying capacity (K) in our model's default parameters is equal to the initial population N_0. When we increase K, we observe an increase in the optimal N_τ, defined as the N_τ at which we get the maximum extinction probability. In Figure 4.7, we see this increase while comparing two values each of K for two different population sizes. For both the populations (10^4 and 10^6 cells), we see a marked increase in the peak of the extinction probability vs N_τ plot. However, also note the decrease in the value of P_E (extinction probability) at the peak as K increases.

Since we know that the optimal N_τ is at the minimum population reached in the ab-

sence of a second strike ($N_{\min}$), we can infer that the $N_{\min}$ for a system with higher carrying capacity is higher. We can visually verify this with Figure 4.8, where we see that the mean $N_{\min}$ is higher for the system with higher carrying capacity. Moreover, we see that the initial decay rate is different in both the cases (higher for lower K). This is because in a population with a lower K, there is an extra constraint on population growth since it is closer to its carrying capacity. This difference leads to a lower $N_{\min}$ in lower K systems.

In our analytical model, the expression for $N_{\min}$ (Eq 4.20) does not depend on K, because we assume logistic growth for the R_1 population (resistant to the first treatment). While in the simulations, we can infer that the initial stages of R_1 population growth is not logistic due to stochastic effects. This leads to a delay in the time needed for the R_1 population to reach a certain size. If the growth were logistic, the change in growth rate in in the R_1 clone would balance the difference in decay rates experienced by the sensitive population. However, due to stochastic growth at small population sizes, these two effects do not balance out, giving a lower $N_{\min}$ than expected, which depends on the carrying capacity.

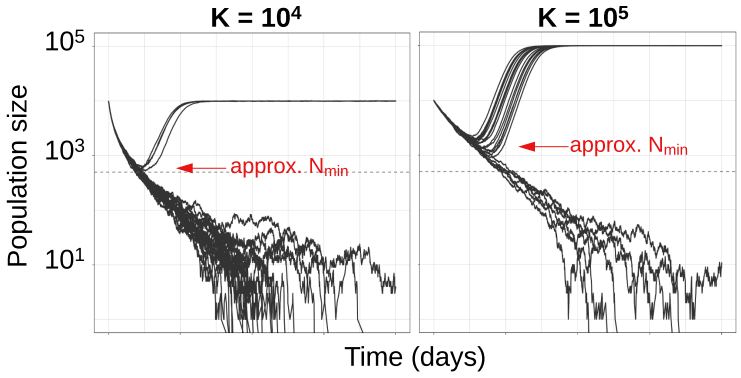

Figure 4.8: Plots showing the difference between population minima in systems with different carrying capacities. The approximate $N_{\min}$ is seen to increase with K. The initial decay rate of the population decreases with an increase in K.

4.3.4 Extinction probability increases with cost of resistance for all second strike thresholds

The default parameters in our model simulations (Table 3.3) include an intrinsic birth rate of 1. Relative to that, we vary the cost of resistance (from 0 to 1) and observe the change in extinction probabilities (P_E) for different second strike thresholds (N_τ). As expected, we see an increase in the extinction probabilities for all N_τ's with $P_E = 1$

for a cost of resistance equal to the intrinsic birth rate. In Figure 4.9, we can see this behaviour for two initial population sizes.

We expect this behaviour because for a low cost of resistance, the resistant variants have a greater advantage over the sensitive cells, so the rescue probability increases. It is interesting to note that the region of high P_E narrows down as we go towards low costs of resistance. This gives us an optimal N_τ for almost all costs of resistance.

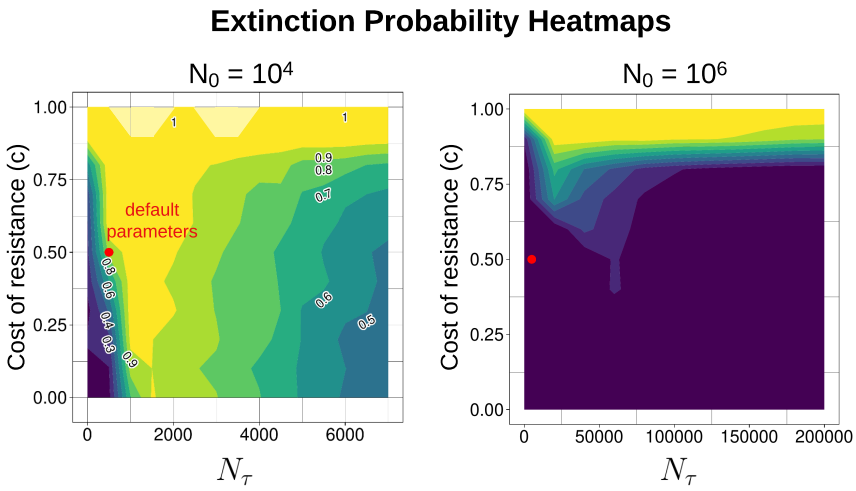

Figure 4.9: Heatmaps showing extinction probability as we vary N_τ and the cost of resistance. We see that the region of high P_E increase as the cost of resistance increases. The default parameters are marked as red dots. Note that the costs for both resistant variants R_1 and R_2 are the same ($c_{R_1} = c_{R_2} = c$).

4.3.5 Extinction probability is higher with a lower treatment level at a given second strike threshold

Perhaps the most important behaviour to understand in extinction therapy is the effect of treatment level. We simulate a range of treatment levels for two different population sizes and several N_τ values. In Figure 4.10A, the first thing we notice is the sharp drop in extinction probability below threshold near 1. This is because the intrinsic growth rate of the sensitive (S) cells is 0.9, so the minimum level of the treatment needed in order to result in a net negative growth rate must be greater than 0.9. Below this limit, P_E will be zero. We see a small region below the line indicating a treatment level of 0.9 because of the heterogeneity in the population. A similar limit, even though not visible in the heatmap (Figure 4.10B) is present in the system with a higher initial population. Above the limit of intrinsic growth rate, we see that with increasing treatment levels, the probability of extinction goes down. This trend is more clearly visible in the system

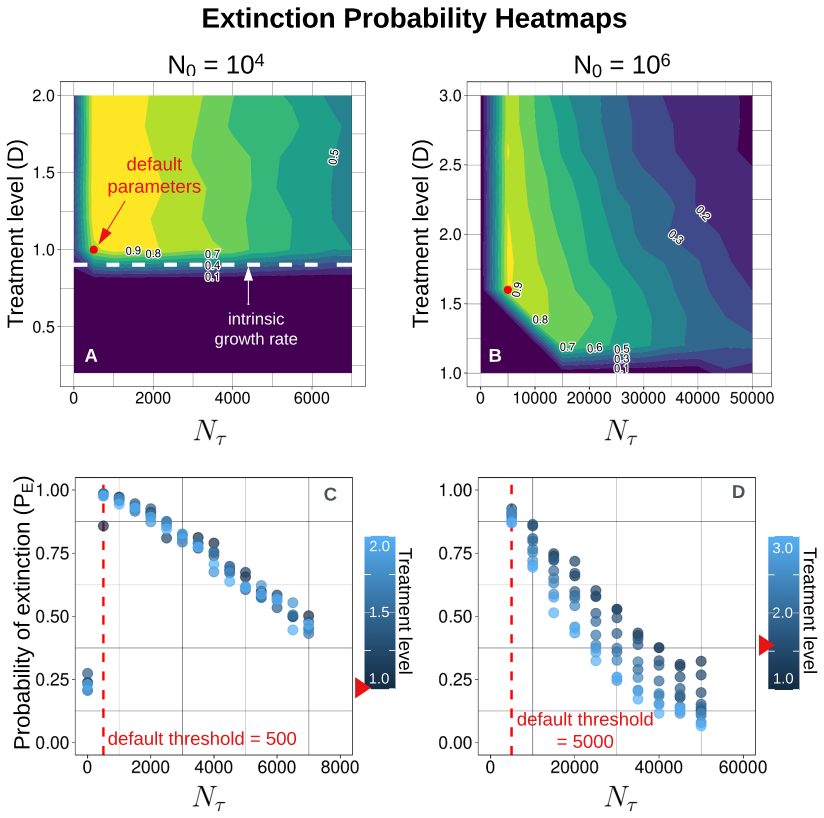

Figure 4.10: Heatmaps of extinction probability as we vary N_τ and the treatment level (treatment-induced death). The intrinsic growth rate of the sensitive population, marked as a white dashed line, is the limit below which $P_E = 0$ due to a net positive growth rate. Above this limit, for a constant N_τ, a higher treatment level has lower P_E than a lower treatment level. Panels C,D show this behaviour in the form of scatter plots. Default parameter values are marked on the heatmaps (red dots) and the scatter plots (red dashed lines and red pointers on the colourbar). Note that the treatment levels for both the treatments are kept the same ($D_{E_1} = D_{E_2} = D$)

with initial population of 10^6 cells (Figure 4.10D). In the system with a smaller starting population, the trend is masked due to higher level of stochasticity, and possibly due to the fact that we only increase the treatment level till $D = 2$.

Lower P_E with a higher treatment level is an unintuitive result. We hypothesize that this is due to a shift in the $N_{\min}$ with change in treatment levels. As predicted by our analytical model, a higher treatment level will result in a lower $N_{\min}$ due to a faster decay of the sensitive population. Because the resistant population (R_1 cells) experience no change in growth rate, the effect of treatment is reflected in the minimum population. While the $N_{\min}$ decreases with treatment level, we keep N_τ's constant in our analysis. Therefore, at the same second strike threshold, we have different $N_{\min}$'s due to

increase in treatment levels. Consequently, a given threshold N_τ will be further away from the $N_{\min}$ for a higher treatment level. With our observation that the maximum P_E is when $N_\tau{=}N_{\min}$, an N_τ that is further away will from the $N_{\min}$ (the optimal N_τ) will have a lower extinction probability.

Notably, the hypothesis explained above does not directly imply that we get the observed trend. Another crucial factor is the faster population decay due to a higher treatment level. Even though a fixed N_τ will be further from the optimal N_τ, it will have the advantage of a greater effect on population size. It is not intuitive as to which factor will dominate in such a situation, but from our simulations, we see that the disadvantage of a suboptimal N_τ is greater than the advantage of a higher decay rate due to treatment.

4.4 Error modelling for Barcode dynamics

For comparing model predictions to data, individual lineages must be tracked experimentally. Genetic barcoding is a method to track cell lineages using unique, short sequences of DNA that are taken up by each cell of the starting population (Bhang et al. 2015). These genetic markers are stably inherited. Each cell gives rise to clones which can be uniquely identified via sequencing techniques following amplification. The relative abundance of barcodes tagging a particular type of cell (say, cells with high growth rates) can be used to infer the absolute number and frequency of cells of that type. This technique is expected and proposed to be a great tool for experiments on extinction therapy because it can help us identify lineages with high growth rates and subsequently classify systems according to their predicted response to extinction therapy. Overall, we can better understand the population dynamics in a system undergoing extinction therapy.

We need to model experimental errors in order to compare data to simulation predictions (Thielecke et al. 2017). There are primarily three sources of noise when analysing cellular barcoding data – stochastic growth effects, DNA amplification and sequencing. We model them separately at first and then combine them to get a total error estimate. DNA amplification via PCR – the first step to quantify barcode abundance in the DNA sample, can lead to errors in measurement due to DNA damage, errors made by polymerase, etc. The amplification process is exponential so the measurements with noise are assumed to follow a log-normal distribution (Noble and Recker 2012).

$$R_{\text{amp}} = 2^X R_{\text{sim}} \tag{4.23}$$

where $X \sim \mathcal{N}(0,\, \sigma^2)$ and R_{sim} is the simulated number of reads. Next, we assume that

sequencing introduces Poisson sampling errors in the number of reads corresponding to a given barcode (Levy et al. 2015). Given the frequency F of the chosen barcodes in the sample, the observed number number of reads R, for that cell lineage is given by a Poisson distribution with mean $R_{\text{amp}}F$ where R_{amp} is the total number of reads post amplification.

$$\mathbb{P}(R = r|F = f) = \frac{(R_{\text{amp}}f)^r \exp(-R_{\text{amp}}f)}{r!} \tag{4.24}$$

Consequently, the abundance of reads corresponding to the barcode follows a mixed Poisson-lognormal distribution. It gives the conditional probability of observing r reads of the barcode given the actual frequency of the corresponding cells in the population is f. Lastly, we consider stochastic growth effects which lead to differences across replicates. Quantifying the stochastic dynamics of multiple clonal lineages, all competing with each other, is extremely complex, so for our purposes, to estimate the variation in growth curves, we can run our simulation several times and obtain approximate distribution parameters like mean and variance.

A complete description of the model is given below where N_r is the number of reads of the desired barcode.

$$X \sim \mathcal{N}(0,\, \sigma^2) \tag{4.25}$$

$$R_{\text{amp}} = 2^X R_{\text{sim}} \tag{4.26}$$

$$r_{|F} \sim Poisson(R_{\text{amp}}F) \quad (\text{where } F = N_r/R_{\text{sim}}) \tag{4.27}$$

$$\mathbb{E}[N_r] = \gamma, \quad \text{Var}[N_r] = \delta \tag{4.28}$$

4.4.1 Deriving the expression for error in number of reads

The conditional probability distribution of the number of reads corresponding to a particular barcode given its true frequency is assumed to be a Poisson distribution (Eq (4.24)) with rate equal to $R_{\text{amp}}F$. Since the rate parameter itself is a lognormal random variable, we have a Mixed Poisson variable R (given F), for which the expected value is given by,

$$\mathbb{E}[R|F] = \mathbb{E}[R_{\text{amp}}F|F = f] \tag{4.29}$$

$$= \mathbb{E}[R_{\text{sim}}2^X F|F = f] \tag{4.30}$$

$$= \mathbb{E}[2^X N_r|N_r = n_r] \tag{4.31}$$

$$= n_r\mathbb{E}[2^X] \tag{4.32}$$

Since $X \sim \mathcal{N}(0,\, \sigma^2)$, the expected value of the lognormal value 2^X is given as follows,

$$v = \mathbb{E}[2^X] = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} 2^x \exp\left(\frac{-x^2}{2\sigma^2}\right) dx \tag{4.33}$$

$$= \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(\frac{-(x-\sigma^2 \ln 2)^2}{2\sigma^2}\right) \exp\left(\frac{\sigma^2 \ln^2 2}{2}\right) dx \tag{4.34}$$

$$= \exp\left(\frac{\sigma^2 \ln^2 2}{2}\right) \tag{4.35}$$

Consequently, $\mathsf{Var}[2^X] = \exp(2\sigma^2 \ln^2 2) - \exp(\sigma^2 \ln^2 2) = v^4 - v^2$. Using these quantities, we can compute the variance of the Mixed Poisson-lognormal variable, given by,

$$\mathsf{Var}[R|F] = \mathbb{E}[R_{\mathsf{amp}} F | F = f] + \mathsf{Var}[R_{\mathsf{amp}} F | F = f] \tag{4.36}$$

$$= n_r v + n_r^2 (v^4 - v^2) \tag{4.37}$$

The unconditional variance of the number of resistant reads R can be calculated using the conditional moments.

$$\mathsf{Var}[R] = \mathbb{E}[\mathsf{Var}[R|F]] + \mathsf{Var}[\mathbb{E}[R|F]] \tag{4.38}$$

$$= \mathbb{E}[N_r v + N_r^2 (v^4 - v^2)] + \mathsf{Var}[N_r v] \tag{4.39}$$

$$= v\mathbb{E}[N_r] + (v^4 - v^2) E[N_r^2] + v^2 \mathsf{Var}[N_r] \tag{4.40}$$

$$= v\gamma - v^2\gamma^2 + v^4(\gamma^2 + \delta) \tag{4.41}$$

where v is a function of σ.

$$v = \exp\left(\frac{\sigma^2}{2} \ln^2 2\right) \tag{4.42}$$

If the value of σ is small, v will be close to 1, which means that the dominant terms in $\mathsf{Var}(R)$ will be δ and γ, the parameters of the growth distribution. This serves as a check of the accuracy of the expression because, in the limit of small experimental error, the total error reduces to noise across replicates.

Chapter 5

Discussion

Extinction therapy is a novel evolutionary therapy which aims to push a cancer to extinction by exploiting the stochasticity of small and vulnerable populations. This is done via the use of multiple "strikes" or treatments at appropriate times. A tumour that responds well to a primary therapy is in a way, primed for a second strike because it is small and susceptible to stochastic effects. When applying extinction therapy on a cancer ecosystem, we aim to "kick it while it's down". The origins of this idea come from Anthropocene extinctions and background extinctions in large-scale ecosystems (Gatenby, Zhang, and Brown 2019). Extinction therapy translates these eco-evolutionary phenomena into the context of cancer treatments.

Previous models on extinction therapy provide a working model and proof of concept for the treatment (Gatenby, Artzy-Randrup, et al. 2020). We build on that with a quantitative, analytical model of extinction therapy with which we can estimate quantities of importance like extinction probabilities (P_E) and optimal timing for the second strike (N_τ). Our simulation results show us how a population under extinction therapy changes when crucial system parameters are altered. Moreover, these results point to the gaps left by the analytical model due to its limitations, and at the same time verify the main conclusions made by it. The combination of analytical and computational analyses, both derived from the principles of evolutionary rescue, arms us with powerful tools to explore extinction therapy in a wide range of parameters with a solid basis in eco-evolutionary theory. Lastly, we model experimental errors and predictions for upcoming experiments in extinction therapy using mathematical modelling and simulations tracking cellular lineages.

Analytical results: Our analytical model, building on our previous preliminary models (Appendix A), is simple, tractable and easy to simulate. We attempt to make as few assumptions as possible and find the minimal set of parameters needed to accurately

model extinction therapy. Using the theory of evolutionary rescue, we find explicit expressions for calculating extinction probabilities. On analysing the behaviour of extinction probability as a function of the second strike threshold (N_τ), we find that the maximum probability of extinction is expected to occur when the second strike threshold is equal to the minimum population reached in the absence of a second strike ($N_{\min}$). This result verifies the hypothesis made by the model proposed by Gatenby, Artzy-Randrup, et al. (2020).

Based on the analytical model, we constructed a simulation model with the same parameters. Comparing analytical predictions to simulation results highlights the limitations of the analytical results. Major deviations from the analytical model are seen in the form of higher extinction probabilities than expected for larger populations. We suggest that deviations from predictions occur because the mathematical model does not account for stochastic population growth at small sizes and makes assumptions of Logistic growth which do not necessarily hold in all cases. We propose plausible explanations to support this hypothesis.

Apart from these limitations, the analytical model is able to accurately predict the optimal second strike threshold – the population size at which if the second strike is administered, the extinction probability is maximised. We also show that this optimal threshold is expected to be equal to the minimum population size reached in the absence of a second strike (a result in agreement with the simulations). We note that the optimal second strike threshold is one of the most important quantities to determine. The ability to analytically predict this value for a large range of parameter values can allow us to design treatment schedules for extinction therapy. Following these results, we define a region of feasibility for accurate prediction of the optimal second strike threshold.

Stochastic simulation results: To explore a wider range of parameter values, including those outside the region of feasibility of the analytical predictions, we ran a large number of simulations which allowed us to empirically calculate extinction probabilities and optimal second strike thresholds (N_τ). First, we show that a second strike increases extinction probability and that reaching a low enough population is necessary for the second strike to be effective. A small population size is key because the population must be in a regime where it is susceptible to stochastic effects. To characterise what "small" means in the context of extinction therapy, we compare different second strike thresholds relative to $N_{\min}$, the minimum population size reached without a second strike. In our cases of interest, the minimum population size is non-zero because those are the cases in which relapse is imminent but can be avoided if the second strike is implemented. As a result, we conclude that $N_{\min}$ is the optimal second

strike threshold, which agrees with our analytical predictions. So when we say that the population must be small enough for extinction therapy to be effective, it means that the population size must be close enough to $N_{\min}$.

The range of values around $N_{\min}$ that result in an extinction probability greater than a given tolerance level (say, 90%) can easily be calculated using both the analytical and simulation models. An important distinction here is between second strike thresholds implemented before and after the minimum population size is reached. In the characteristic U-shaped trajectory of a population undergoing evolutionary rescue, a given population threshold N_τ is met twice, once before the minima and once after. Previous models do not distinguish between these two and hypothesize that the "before-minima" thresholds are better in terms of treatment efficacy. On the contrary, we show that the "after-minima" thresholds result in a much higher extinction rate. We hypothesize that this is due to the difference in fitness of resistant and sensitive clones. Since the population is dominated by resistant clones after the minima has been reached, treatment-induced death adds to the cost of resistance resulting in an overall higher death rate than would be observed before the minima. Consequently, the so-called "window of opportunity" extends further to the right of the minima than it does to the left. In fact, on the right side of the minima, the probability of extinction can be as much as double the value for the same N_τ on the left side. In practice, this would imply that an error in the estimation of $N_{\min}$ is much costlier to the left than to the right (after the minima), so it is better to wait and risk missing the optimal N_τ than to make the same error before the minima has reached.

Next, we explore the parameter space of our model, with a focus on the second strike threshold (N_τ). We find that optimal N_τ increases with carrying capacity (K). This is because we see an increase in the minimum population $N_{\min}$ as K increases. However, the analytical calculation for $N_{\min}$ does not include this dependence on carrying capacity possibly because it comes about due to stochastic growth effects in the initial stages of growth for the resistant clone R_1, which leads to a slower overall growth than expected by the analytical model. Meanwhile, the sensitive population continues to decay at the expected rate, resulting in a lower $N_{\min}$ for lower carrying capacities. For a given second strike threshold, however, the peak extinction probability will tend to decrease as the $N_{\min}$ goes higher and further away from the threshold. Considering the effect of carrying capacity on extinction therapy is very important. It plays a role in *in-vitro* experiments (size of the petri dish, for example) and in actual tumours in patients. There are processes like angiogenesis that can alter the carrying capacity of a tumour. Ultimately, in order to implement extinction therapy, one must take into account the effect of carrying capacity and its implications for the given system.

We also assess the effect of fitness of the resistant clones on the N_τ threshold. As expected, we see a higher extinction probability for less fit clones (high cost of resistance). Moreover, we see an increase in the size of the "window of opportunity" as the cost of resistance increases. This implies that if the resistance clones are not too fit compared to the sensitive clones, then it is highly likely that the population goes extinct regardless of the second strike. However, these are the cases in which the first strike is enough to kill the entire population so there is no relapse. We are interested in cases where evolutionary rescue is imminent. Therefore, we care about fit resistant clones (low cost of resistance) and infer that even in the extreme cases (no cost of resistance), there exists a window of opportunity, albeit narrow, making extinction therapy a viable option.

Finally, we investigate the effect of treatment-induced death rate on extinction probabilities. Surprisingly, we find that for a given second strike threshold, higher treatment levels result in a lower extinction probability. We hypothesize that this trend for high treatment levels is a result of two opposing effects. A higher death rate pushes the population towards extinction, but also pushes the $N_{\min}$ (or optimal N_τ) towards lower values. Since the N_τ is fixed, a lower $N_{\min}$ means that the threshold is further away from the optimum. The disadvantage of suboptimal timing ends up dominating over other factors. This result emphasizes the importance of timing in extinction therapy – A stronger treatment with bad timing is worse than a weaker treatment given at the right time.

Predictive modelling for experiments in extinction therapy: Given the current nascent stage of the development of extinction therapy, experimental verification and testing of the concept is necessary in the near future. These experiments are expected to use techniques like genetic barcoding to track cellular lineages to test the efficacy of extinction therapy. To aid the design and analysis of future experiments, we performed predictive mathematical modelling with two aims – 1) to help characterize systems based on their growth distributions in order to determine if they are suitable for extinction therapy, and 2) to model errors in barcoding experiments to determine the accuracy of lineage tracking experiments. For aim 1, we used stochastic simulations to show that systems with high and low costs of resistance respond differently to treatments and illustrated how such differences in intrinsic growth parameters can be experimentally determined. We find that growth rate distributions and lineage counts at certain time points are useful tools for characterizing systems undergoing extinction therapy. For aim 2, we derived an expression for error in the number of reads of a given barcode, taking into account the errors introduced by PCR amplification and sequencing. Using this expression with experimental data can provide an estimate of the experimental error. Predictive modelling of extinction therapy with experiments in

mind helps us understand the limitations of theoretical work and tells us what to expect when these experiments do get conducted.

Limitations of this work: Our results are based on a few assumptions and consequently subject to limitations. First, our analytical model does not account for stochastic growth effects which are a crucial component of population dynamics at low population sizes. Even from the stochastic simulation results, it is clear that the analytical predictions deviate from empirical outcomes in cases where stochastic dynamics are expected to play an important role. Second, we use a very simple and specific model formulation to obtain our results. As model specifications change, the results might vary. While the qualitative insights would still be valid, the significance of individual parameters depends highly on the model formulation. For example, the effect of carrying capacity seen in our results may be dependent on the way density dependence manifests in the population. Third, for the sake of simplicity we limit our analysis to a few cases with a limited number of parameters while keeping everything else constant. For example, treatment-induced death rates are kept equal for both the treatments in our model. Fourth, we only consider small populations relative to a typical tumour burden. Therefore, our work may be valid for tumour metastases but when the population size is in the order of a billion cells, the results may vary. With this work, we aim to produce the simplest working model of extinction therapy in order to understand the fundamental dynamics that take place in such systems. However, several aspects of our work need to be expanded upon to gain a more realistic understanding of the system.

When should extinction therapy be used? Extinction therapy involves eradicating at least a large portion of the initial population in order to take it to a stage where it is susceptible to further strikes. However, such a strategy may not always be the best way to go. For instance, in systems with very fit resistant phenotypes, extinction therapy might just give way to competitive release. More importantly, we are faced with practical considerations like the availability of a large and diverse group of drugs/treatments that the tumour responds to, and the feasibility of techniques to monitor tumour burden at a chosen time point (Reed et al. 2020). The most added benefit of ET would show up in cases where the rate of relapse is high following a good response to treatment.

Another situation in which extinction therapy may be useful is when the second treatment is not as effective as the first one (also suggested by Gatenby, Artzy-Randrup, et al. (2020)). The first treatment will drastically reduce the population size but may result in a relapse. However, even a weaker treatment is effective as a second strike because of two reasons – 1) small population sizes are subject to stochastic extinctions and even a weak treatment can push the population into the extinction vortex, and 2) if the population is dominated by treatment-resistant cell types at the beginning

of the second strike, a weaker treatment is sufficient to cause population decay because the population's intrinsic growth rate is already limited by the cost of resistance. Ultimately, extinction therapy is about the strategic use of multiple treatments to obtain the best possible outcome (Gatenby, Artzy-Randrup, et al. 2020).

Extinction therapy in the clinic: The closest practical realisation of the concept of extinction therapy is seen in the treatment of paediatric Acute Lymphoblastic Leukaemia (ALL), which has multiple "phases" of chemotherapy of varying intensities (Gatenby, Zhang, and Brown 2019). However, there is little exploration of the optimal time difference between two phases or the optimal intensity of each phase. Nonetheless, the success of the ALL treatment regime motivates further investigation into the causal eco-evolutionary factors behind it. In this case, there are already considerations of extinction therapy in the context of diseases like LARC (locally advanced rectal adenocarcinoma, Felder, Fleming, and Gatenby (2021)), mPC (metastatic prostate cancer, Gatenby, Zhang, and Brown (2019)) and paediatric sarcomas (Reed et al. 2020). As the influence of evolutionary therapies increases, we see more and more clinical and pre-clinical trials being conducted to test this new paradigm.

Conclusion and future directions: Extinction therapy is still a theoretical model at this stage, lacking any experimental evidence, with perhaps a little empirical support. A robust, complete theory of extinction therapy is needed for a comprehensive understanding of this phenomenon and its applicability in the context of cancer. In this thesis, we set the stage for further theoretical and experimental development of extinction therapy. Our model emphasizes the importance of the timing of the second strike and analyses its effects on other model parameters. We predict the behaviour of systems undergoing extinction therapy within the scope of our analytical and stochastic simulations model. However, it is important to integrate stochastic growth effects and plasticity into analytical models and explore the implications of continuous resistance traits. Overcoming the limitations presented in our model could be the first step to a more general understanding of extinction therapy.

Bibliography

Aktipis, C. Athena et al. (Nov. 2011). “Overlooking Evolution: A Systematic Analysis of Cancer Relapse and Therapeutic Resistance Research”. en. In: *PLOS ONE* 6.11, e26100. ISSN: 1932-6203.

Alexander, Helen K. et al. (2014). “Evolutionary rescue: linking theory for conservation and medicine”. en. In: *Evolutionary Applications* 7.10, pp. 1161–1179. ISSN: 1752-4571.

Bell, Graham (Jan. 2013). “Evolutionary rescue and the limits of adaptation”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 368.1610, p. 20120080.

— (2017). “Evolutionary Rescue”. In: *Annual Review of Ecology, Evolution, and Systematics* 48.1, pp. 605–627.

Bhang, Hyo-eun C et al. (May 2015). “Studying clonal dynamics in response to cancer therapy using high-complexity barcoding”. en. In: *Nature Medicine* 21.5, pp. 440–448. ISSN: 1078-8956, 1546-170X.

Brady-Nicholls, Renee et al. (Sept. 2021). “Predicting patient-specific response to adaptive therapy in metastatic castration-resistant prostate cancer using prostate-specific antigen dynamics”. en. In: *Neoplasia* 23.9, pp. 851–858. ISSN: 1476-5586.

Brown, Joel S., Jessica J. Cunningham, and Robert A. Gatenby (Mar. 2017). “Aggregation Effects and Population-Based Dynamics as a Source of Therapy Resistance in Cancer”. In: *IEEE Transactions on Biomedical Engineering* 64.3, pp. 512–518. ISSN: 1558-2531.

Carlson, Stephanie M., Curry J. Cunningham, and Peter A. H. Westley (Sept. 2014). “Evolutionary rescue in a changing world”. en. In: *Trends in Ecology & Evolution* 29.9, pp. 521–530. ISSN: 0169-5347.

Chakrabarti, Shaon and Franziska Michor (July 2017). “Pharmacokinetics and Drug Interactions Determine Optimum Combination Strategies in Computational Models of Cancer Evolution”. en. In: *Cancer Research* 77.14, pp. 3908–3921. ISSN: 0008-5472, 1538-7445.

Cunningham, Jessica J. et al. (Dec. 2018). “Optimal control to develop therapeutic strategies for metastatic castrate resistant prostate cancer”. en. In: *Journal of Theoretical Biology* 459, pp. 67–78. ISSN: 0022-5193.

Dennis, Brian et al. (Sept. 2016). “Allee effects and resilience in stochastic populations”. en. In: *Theoretical Ecology* 9.3, pp. 323–335. ISSN: 1874-1746.

Enriquez-Navas, Pedro M., Yoonseok Kam, et al. (Feb. 2016). “Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer”. In: *Science Translational Medicine* 8.327. Publisher: American Association for the Advancement of Science. ISSN: 19466242.

Enriquez-Navas, Pedro M., Jonathan W. Wojtkowiak, and Robert A. Gatenby (Nov. 2015). “Application of Evolutionary Principles to Cancer Therapy”. In: *Cancer Research* 75.22, pp. 4675–4680. ISSN: 0008-5472. eprint: https://aacrjournals.org/cancerres/article-pdf/75/22/4675/2937393/4675.pdf.

Felder, Seth I., Jason B. Fleming, and Robert A. Gatenby (2021). “Treatment-induced evolutionary dynamics in nonmetastatic locally advanced rectal adenocarcinoma”. en. In: *Advances in Cancer Research*. Vol. 151. Elsevier, pp. 39–67. ISBN: 978-0-12-824078-6.

Gallaher, Jill et al. (Aug. 2022). *The sum and the parts: dynamics of multiple and individual metastases during adaptive therapy*. en. Pages: 2022.08.04.502852 Section: New Results.

Gallaher, Jill A., Joel S. Brown, and Alexander R. A. Anderson (Feb. 2019). “The impact of proliferation-migration tradeoffs on phenotypic evolution in cancer”. en. In: *Scientific Reports* 9.1. Number: 1 Publisher: Nature Publishing Group, p. 2425. ISSN: 2045-2322.

Gatenby, Robert A., Yael Artzy-Randrup, et al. (Feb. 2020). “Eradicating metastatic cancer and the eco-evolutionary dynamics of Anthropocene extinctions”. In: *Cancer Research* 80.3, pp. 613–623. ISSN: 15387445.

Gatenby, Robert A. and Joel S. Brown (Nov. 2020). “Integrating evolutionary dynamics into cancer therapy”. en. In: *Nature Reviews Clinical Oncology* 17.11. Number: 11 Publisher: Nature Publishing Group, pp. 675–686. ISSN: 1759-4782.

Gatenby, Robert A., Ariosto S. Silva, et al. (June 2009). “Adaptive therapy”. In: *Cancer Research* 69.11, pp. 4894–4903. ISSN: 00085472.

Gatenby, Robert A., Jingsong Zhang, and Joel S. Brown (2019). “First strike-second strike strategies in metastatic cancer: Lessons from the evolutionary dynamics of extinction”. In: *Cancer Research* 79.13, pp. 3174–3177. ISSN: 15387445.

Gillespie, Daniel T. (1977). “Exact stochastic simulation of coupled chemical reactions”. In: *The Journal of Physical Chemistry* 81.25, pp. 2340–2361. eprint: https://doi.org/10.1021/j100540a008.

Greaves, Mel and Carlo C. Maley (Jan. 2012). “Clonal evolution in cancer”. en. In: *Nature* 481.7381. Number: 7381 Publisher: Nature Publishing Group, pp. 306–313. ISSN: 1476-4687.

Haldane, J. B. S. (1927). “A Mathematical Theory of Natural and Artificial Selection, Part V: Selection and Mutation”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 23.7, pp. 838–844.

Iwasa, Yoh, Martin A Nowak, and Franziska Michor (Apr. 2006). “Evolution of Resistance During Clonal Expansion”. In: *Genetics* 172.4, pp. 2557–2566. ISSN: 1943-2631.

Kavran, Andrew J. et al. (Mar. 2022). “Intermittent treatment of BRAFV600E melanoma cells delays resistance by adaptive resensitization to drug rechallenge”. In: *Proceedings of the National Academy of Sciences* 119.12. Publisher: Proceedings of the National Academy of Sciences, e2113535119.

Korolev, Kirill S., Joao B. Xavier, and Jeff Gore (May 2014). “Turning ecology and evolution against cancer”. en. In: *Nature Reviews Cancer* 14.5, pp. 371–380. ISSN: 1474-1768.

Kuosmanen, Teemu (n.d.). “Drug Resistance as an Evolutionary Rescue: Insights to Treatment Optimization in Cancer”. en. In: (), p. 47.

Lambert, Amaury (June 2006). “Probability of fixation under weak selection: A branching process unifying approach”. en. In: *Theoretical Population Biology* 69.4, pp. 419–441. ISSN: 00405809.

Levy, Sasha F. et al. (Mar. 2015). “Quantitative evolutionary dynamics using high-resolution lineage tracking”. en. In: *Nature* 519.7542, pp. 181–186. ISSN: 1476-4687.

Martin, Guillaume et al. (Jan. 2013). “The probability of evolutionary rescue: towards a quantitative comparison between theory and evolution experiments”. en. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 368.1610, p. 20120088. ISSN: 0962-8436, 1471-2970.

Martin, R. B. et al. (July 1992). “Optimal control of tumor size used to maximize survival time when cells are resistant to chemotherapy”. en. In: *Mathematical Biosciences* 110.2, pp. 201–219. ISSN: 0025-5564.

Meille, Christophe et al. (Aug. 2016). “Revisiting Dosing Regimen Using Pharmacokinetic/Pharmacodynamic Mathematical Modeling: Densification and Intensification of Combination Cancer Therapy”. en. In: *Clinical Pharmacokinetics* 55.8, pp. 1015–1025. ISSN: 0312-5963, 1179-1926.

Monro, Helen C. and Eamonn A. Gaffney (2009). “Modelling chemotherapy resistance in palliation and failed cure”. In: *Journal of Theoretical Biology* 257.2, pp. 292–302. ISSN: 0022-5193.

Noble, Robert and Mario Recker (June 2012). “A Statistically Rigorous Method for Determining Antigenic Switching Networks”. en. In: *PLOS ONE* 7.6, e39335. ISSN: 1932-6203.

Orr, H. Allen and Robert L. Unckless (Aug. 2008). “Population Extinction and the Genetics of Adaptation”. en. In: *The American Naturalist* 172.2, pp. 160–169. ISSN: 0003-0147, 1537-5323.

— (Aug. 2014). “The Population Genetics of Evolutionary Rescue”. en. In: *PLOS Genetics* 10.8, e1004551. ISSN: 1553-7404.

Otto, S. P. and M. C. Whitlock (June 1997). “The Probability of Fixation in Populations of Changing Size”. In: *Genetics* 146.2, pp. 723–733. ISSN: 0016-6731.

Pressley, Mariyah et al. (2021). “Evolutionary Dynamics of Treatment-Induced Resistance in Cancer Informs Understanding of Rapid Evolution in Natural Systems”. In: *Frontiers in Ecology and Evolution* 9. ISSN: 2296-701X.

Reed, Damon R. et al. (2020). “An evolutionary framework for treating pediatric sarcomas”. In: *Cancer* 126.11. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cncr.32777, pp. 2577–2587. ISSN: 1097-0142.

Strobl, Maximilian et al. (Mar. 2020). *Turnover modulates the need for a cost of resistance in adaptive therapy*. en. Pages: 2020.01.22.914366 Section: New Results.

Thielecke, Lars et al. (Mar. 2017). “Limitations and challenges of genetic barcode quantification”. en. In: *Scientific Reports* 7.1, p. 43249. ISSN: 2045-2322.

Uecker, Hildegard, Sarah P. Otto, and Joachim Hermisson (Jan. 2014). “Evolutionary Rescue in Structured Populations”. en. In: *The American Naturalist* 183.1, E17–E35. ISSN: 0003-0147, 1537-5323.

Viossat, Yannick and Robert Noble (June 2021). “A theoretical analysis of tumour containment”. In: *Nature Ecology and Evolution* 5.6, pp. 826–835. ISSN: 2397334X.

Walther, Viola et al. (May 2015). “Can oncology recapitulate paleontology? Lessons from species extinctions”. en. In: *Nature Reviews Clinical Oncology* 12.5, pp. 273–285. ISSN: 1759-4782.

West, Jeffrey, Fred Adler, et al. (Mar. 2023). “A survey of open questions in adaptive therapy: Bridging mathematics and clinical translation”. In: *eLife* 12. Ed. by Richard M White. Publisher: eLife Sciences Publications, Ltd, e84263. ISSN: 2050-084X.

West, Jeffrey, Yongqian Ma, and Paul K. Newton (Oct. 2018). “Capitalizing on competition: An evolutionary model of competitive release in metastatic castration resistant prostate cancer treatment”. en. In: *Journal of Theoretical Biology* 455, pp. 249–260. ISSN: 0022-5193.

You, Li et al. (Dec. 2017). “Spatial vs. non-spatial eco-evolutionary dynamics in a tumor growth model”. en. In: *Journal of Theoretical Biology* 435, pp. 78–97. ISSN: 0022-5193.

Zhang, Jingsong et al. (Nov. 2017). “Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer”. en. In: *Nature Communications* 8.1, p. 1816. ISSN: 2041-1723.

Appendices

Appendix A

Analytical insights from preliminary models

While attempting to build a simple model for extinction therapy, we analysed some preliminary models and understood the components necessary for our main model. Here are 4 such preliminary models.

A.1 Case 1: Bottleneck treatment with exponential decay

To begin with, we try to simplify the type of treatment described in Gatenby, Artzy-Randrup, et al. (2020) and model it using evolutionary rescue theory. Here, we consider only one resistant variant, since there is only one drug-based treatment. The second strike is administered in the form of a bottleneck on the population, after which a certain fraction q of cells are killed. The second strike does not distinguish between resistant (R) and sensitive (S) cells and kills cells at random. The rationale for this could be that the second strike would typically be chosen such that it has a different mechanism of action than the primary drug, which means that both types of cells will be sensitive to this strike.

For simplicity, we first consider exponential decay (with rate $g < 0$) of the population due to the primary treatment. Additionally, we assume that the rescue mutants have a small growth rate so that the total number of rescue mutants can be estimated by the number of mutations from $S \to R$. Following the theory developed in Section 2, the instantaneous rate of generation of resistant mutants is $N_t \mu \pi$ where μ is the per capita mutation rate ($S \to R$) and π is the fixation probability of a single resistant mutant (see

Section 2). The rescue variants generated till time τ are Poisson distributed with the rate,

$$\lambda = \mu\pi \int_0^\tau N_t dt \tag{A.1}$$

$$= \mu\pi \frac{N_0 - N_\tau}{|g|} \tag{A.2}$$

We apply the second treatment at time τ, with the assumption that $N_\tau > N_{\min}$. This means that N_τ is a decreasing function of τ. At time τ, a fraction $q = k/N_\tau$ of all cells are killed at random. The population can be rescued either if any rescue mutant survives the bottleneck or if resistant mutants generated after the second strike escape stochastic extinction. This means that if the population is to go extinct, two conditions must be met: no rescue mutants should survive after the bottleneck *and* any resistant mutants generated after the bottleneck should not fix in the population.

For condition 1, we derive the probability that the second strike kills all rescue mutants in the population till time τ. Assuming that the number of cells k that die during the second strike is much larger than the rate λ, we obtain that the probability of no rescue mutants surviving the bottleneck is $\exp[-\lambda(1-q)]$ (derivation in Section A.1.1). For condition 2, analogous to Eq A.2, rate of generation of rescue mutants after the second strike is $\lambda' = N_\tau(1-q)\mu\pi/|g|$.

Therefore, the total extinction probability for this case is,

$$P_E = \exp[-\lambda(1-q)] \exp[-\lambda'] \tag{A.3}$$

To find the optimal time for the second strike, we find the value of τ for which $\lambda(1-q)+\lambda'$ is minimum, which corresponds to the maximum P_E.

$$\underset{\tau}{\operatorname{argmin}} \quad \lambda(1-q) + \lambda' \tag{A.4}$$

$$= \underset{\tau}{\operatorname{argmin}} \quad \mu\pi(1-q)\frac{N_0 - N_\tau}{|g|} + \mu\pi(1-q)\frac{N_\tau}{|g|} \tag{A.5}$$

$$= \underset{\tau}{\operatorname{argmin}} \quad \mu\pi(1-q)\frac{N_0}{|g|} \tag{A.6}$$

From Eq A.6, we see that the time of the second strike does not matter in this case. This result implies that if a therapy eliminates a fixed fraction of the population, the probability of extinction does not depend on the population size at which the elimination occurs. This can be explained by looking at the two terms in Eq A.3. The first

part, corresponding to the probability of killing all rescue mutants in the second strike, decreases with N_τ because the fraction of rescue mutants in the population keeps on increasing while the total population decays. It becomes harder and harder to eliminate all rescue mutants as the threshold population decreases and with it, the number of cells that die in the second strike decreases. However, as a result the time after the second strike till extinction also decreases with N_τ, because N_τ is lower for higher values for τ. This does not allow many rescue mutants to escape stochastic death after the second strike. So, the second term of the expression is higher at lower N_τ's, as opposed to the first term, which is lower at lower N_τ's. These two terms perfectly balance each other and P_E in this case turns out to be constant throughout,

$$P_E = \exp\left[-\mu\pi(1-q)\frac{N_0}{|g|}\right]. \tag{A.7}$$

A.1.1 Derivation of probability of rescue mutants after a bottleneck

Given that a fraction q of the total population is killed in the bottleneck, we derirve the probability that there are no rescue mutants left after the bottleneck. Assume that $k = qN_\tau$ cells are eliminated in the bottleneck and X is a random variable denoting the number of rescue mutants before the bottleneck, then the probability that the number of rescue mutants after the bottleneck (X') is equal to zero is,

$$\mathbb{P}(X'=0) = \sum_{i=0}^{k} \frac{\mathbb{P}(X=i)\binom{N_\tau - i}{k-i}}{\binom{N_\tau}{k}} \tag{A.8}$$

$$= \frac{k!(N_\tau-k)!}{N_\tau!}\sum_{1=0}^{k}\frac{\mathbb{P}(X=i)(N_\tau-i)!}{(k-i)!(N_\tau-k)!} \tag{A.9}$$

$$= \frac{k!}{N_\tau!}\sum_{i=0}^{k}\frac{\exp(-\lambda)\lambda^i}{i!}\frac{(N_\tau-i)!}{(k-i)!} \tag{A.10}$$

where λ is the Poisson rate of generation of resistant mutants. Now, assuming that $i << k, N_\tau$ because the probability of an i close to k would be small, meaning that the product of the terms inside summation will be small, we get:

$$\mathbb{P}(X'=0) = \frac{k!}{N_\tau!}\sum_{i=0}^{k}\frac{\exp(-\lambda)\lambda^i}{i!}\frac{N_\tau!}{k!}\left(\frac{k}{N_\tau}\right)^i \tag{A.11}$$

$$= \sum_{i=0}^{k}\frac{\exp(-\lambda)\lambda^i q^i}{i!} \tag{A.12}$$

Assuming that $k >> \lambda$ which, again is reasonable, given that λ is the rate of generation of resistant mutants, we have:

$$\mathbb{P}(X' = 0) = \sum_{i=0}^{\infty} \exp(-\lambda(1-q)) \frac{(\lambda q)^i \exp(-\lambda q)}{i!} \tag{A.13}$$

$$= \exp(-\lambda(1-q)) \tag{A.14}$$

A.2 Case 2: Bottleneck treatment with density dependence

In this case, we add density dependence in the framework of Case 1. The carrying capacity of the system is K and the rest of the parameters remain the same. We have the distribution of rescue mutants described by a Poisson distribution with rate given by Eq A.1. Given the Logistic growth equation with growth rate $g < 0$, we substitute the integral term,

$$\lambda = \mu\pi \int_{N_0}^{N_\tau} \frac{dN_t}{g(1 - N_t/K)} \tag{A.15}$$

$$= \frac{\mu\pi K}{g} \int_{N_0}^{N_\tau} \frac{dN_t}{K(1 - N_t/K)} \tag{A.16}$$

$$= \frac{\mu\pi K}{|g|} \left[\ln\left(1 - \frac{N_\tau}{K}\right) - \ln\left(1 - \frac{N_0}{K}\right) \right] \tag{A.17}$$

We apply the second treatment at time τ, eliminating a fraction q of all cells. Again, the probability that all rescue mutants die in this strike is $\exp[-\lambda(1-q)]$ (assuming that the growth rate of rescue mutants is positive but small). After the bottleneck, rescue mutants are generated with rate λ',

$$\lambda' = \frac{\mu\pi K}{g} \int_{N_\tau(1-q)}^{0} \frac{dN_t}{K(1 - N_t/K)} \tag{A.18}$$

$$= \frac{\mu\pi K}{|g|} \left[-\ln\left(1 - \frac{N_\tau(1-q)}{K}\right) \right] \tag{A.19}$$

Therefore, the total extinction probability for this case is,

$$P_E = \exp[-\lambda(1-q)] \exp[-\lambda'] \tag{A.20}$$

To find the optimal time for the second strike, we find the value of τ for which $\lambda(1-q)+\lambda'$ is minimum, which corresponds to the maximum P_E.

$$-\ln P_E = \lambda(1-q) + \lambda' \tag{A.21}$$

$$= \frac{\mu\pi K(1-q)}{|g|}\left[\ln\left(1-\frac{N_\tau}{K}\right) - \ln\left(1-\frac{N_0}{K}\right)\right] + \frac{\mu\pi K}{|g|}\left[-\ln\left(1-\frac{N_\tau(1-q)}{K}\right)\right] \tag{A.22}$$

$$= -\frac{\mu\pi K(1-q)}{|g|}\ln\left(1-\frac{N_0}{K}\right) + \frac{\mu\pi K}{|g|}\left[(1-q)\ln\left(1-\frac{N_\tau}{K}\right) - \ln\left(1-\frac{N_\tau(1-q)}{K}\right)\right] \tag{A.23}$$

Since the value of N_τ can range from 0 to N_0, we see that the maximum extinction probability is obtained at $N_\tau = N_0$. Therefore, we get minimum P_E at $N_\tau = 0$ and it keeps increasing for higher values of N_τ.

$$\min\ P_E = \exp\left[\frac{\mu\pi K(1-q)}{|g|}\ln\left(1-\frac{N_0}{K}\right)\right] \tag{A.24}$$

$$\max\ P_E = \exp\left[\frac{\mu\pi K}{|g|}\ln\left(1-\frac{(1-q)N_0}{K}\right)\right] \tag{A.25}$$

In this case, the second treatment increases the probability of extinction as N_τ increases. This is different from case 1 where due to exponential decay, the probability of killing rescue mutants in the second strike does not depend on the population size at the time of the second strike because the two terms in the expression for P_E cancel out to give a constant value throughout. That balance is skewed here because the population experiences a faster decay due to the extra constraint of a carrying capacity. At the same population size, a population under density dependence will have fewer rescue mutants than a population with exponential decay.

Note that the biggest drawback both in cases 1 and 2 is the assumption that rescue mutants have a small growth rate. We make this assumption so that we can estimate the total population of rescue mutants with the number of rescue mutant lineages. The expression for the probability distribution and expected number of rescue mutants at a given time τ is hard to compute analytically, given the stochastic nature of population growth. For the same reason, we do not take into account standing genetic variation. Heuristically, if we do consider these two effects, then the probability of extinction will decrease overall.

A.3 Case 3: Two overlapping drug treatments with exponential decay

A bottleneck treatment that instantly eliminates a fixed proportion of cells is not a realistic system. In practice, one would tend to combine two or more drugs that have complementary mechanisms of action. When we consider two drug treatments, we have to model two resistant variants and a single wild-type or sensitive variant. Each resistant variant corresponds to a different environment. The environment in turn depends on the type of treatment in both the strikes. In this case, we take both the treatments to be drugs that alter the growth rate of tumour cells. We thus have two environments with different effects on the cells. Environment 1 (E_1) is when only the first drug (drug 1) is administered. Environment 2 (E_2) is when the second drug (drug 2) is given in addition to the first drug. Consequently, we have two types of resistant mutants – clone R_1 resistant to drug 1 and clone R_2 resistant to both the drugs. Note that saying a cell is "resistant" to a drug is equivalent to saying that it has positive growth rate in the corresponding environment. Table A.1 lists all the relevant parameters for this case.

Parameter	Description
$d_i < 0$	Decay rate of the sensitive cells in E_i $(i = 1, 2)$
μ_i	Rate of mutation to resistant variants in E_i $(i = 0, 1, 2)$
π_j	Probability of establishment of a single resistant cell in clone R_j $(j = 1, 2)$
p_j	Probability of mutating to the j^{th} resistant variant $(j = 1, 2)$

Table A.1: List of parameters for case 3

The instantaneous rate of generation of rescue mutants in E_1 is given by $N_t\mu_1\overline{\pi}dt$, where N_t is the total population of cells at time t and $\overline{\pi} = p_1\pi_1 + p_2\pi_2$ is the expected probability of establishment of a single resistant cell, which consists of both the resistant cell type fixation probabilities because both R_1 and R_2 cells are resistant in E_1. The rate of generation of rescue mutants in E_1 is,

$$\lambda_1 = \int_0^\tau N_t\mu_1\overline{\pi}dt = \mu_1\overline{\pi}\int_0^\tau N_t dt \tag{A.26}$$

$$= \frac{\mu_1\overline{\pi}}{|d_1|}(N_0 - N_\tau) \tag{A.27}$$

Similarly, for environment E_2, we have,

$$\lambda_2 = \frac{\mu_2\pi_2 N_\tau}{|d_2|} \tag{A.28}$$

We take the expected probability of establishment for λ_1 but not for λ_2 because clone R_2 is resistant to both the drugs and can lead to the generation of rescue mutants in both the environments, but R_1 is resistant only in E_1 and does not play a role in E_2.

Therefore, the probability of extinction is given by,

$$P_E = \exp[-(\lambda_1 + \lambda_2)] \tag{A.29}$$

Our aim is to maximise the probability of extinction with respect to the time of the second strike (τ).

$$\begin{aligned}
&\underset{\tau}{\operatorname{argmax}}\ P_E && \text{(A.30)}\\
=\ &\underset{\tau}{\operatorname{argmin}}\ \lambda_1 + \lambda_2 && \text{(A.31)}\\
=\ &\underset{\tau}{\operatorname{argmin}}\ \frac{\mu_1\overline{\pi}}{|d_1|}(N_0 - N_\tau) + \frac{\mu_2\pi_2 N_\tau}{|d_2|} && \text{(A.32)}\\
=\ &\underset{\tau}{\operatorname{argmin}}\ N_0\frac{\mu_1\overline{\pi}}{|d_1|} + N_\tau\left(\frac{\mu_2\pi_2}{|d_2|} - \frac{\mu_1\overline{\pi}}{|d_1|}\right) && \text{(A.33)}\\
=\ &\underset{\tau}{\operatorname{argmin}}\ N_\tau\left(\frac{\mu_2\pi_2}{|d_2|} - \frac{\mu_1\overline{\pi}}{|d_1|}\right) && \text{(A.34)}
\end{aligned}$$

Hence, the optimal time for the second strike in this case, depends on the relative values of parameters in both the environments. Let $U = \mu_2/\mu_1$ and $D = |d_2/d_1|$ be the relative change in mutation rates and wild-type decay rates in E_1 and E_2. If τ_E is the time of extinction of the population, then

$$\frac{\pi_2}{\pi_1} > \frac{p_1 D}{U - p_2 D} \implies \underset{\tau}{\operatorname{argmax}}\ P_E = \tau_E \quad \text{and,} \tag{A.35}$$

$$\frac{\pi_2}{\pi_1} < \frac{p_1 D}{U - p_2 D} \implies \underset{\tau}{\operatorname{argmax}}\ P_E = 0 \tag{A.36}$$

We observe that in this case, the second drug should either not be given or should be given simultaneously with the first drug to get a combined effect. The decision depends on the ratio of the probabilities of fixation of the two resistant variants. If $U = D = 1$, then the comparison is just between the fixation probabilities. If the second resistant variant has a higher probability of fixation, then it is better to not administer the second strike at all.

One must note that this simple condition resulting in only two outcomes – a complete absence of drug 2 versus simultaneous administration of the two drugs – is a result of the type of treatments considered. If there is an environment with both the drugs present (i.e. a second strike is given), it is likely that a cell type resistant to both the treatments will emerge and fix in the population. If there is no second strike, then the

R_1 cells will fix in the population, if they escape stochastic extinction. Therefore, the optimal N_τ only depends on the difference in growth rates between the two resistant cell types and one of the treatments is rendered obsolete in any case.

A.4 Case 4: Two overlapping drug treatments with density dependence

Next, if we add density dependence to case 3, from Eq A.32 we get,

$$\lambda_1 + \lambda_2 \tag{A.37}$$

$$= \frac{\mu_1 \overline{\pi} K}{|d_1|}\left[\ln\left(1-\frac{N_\tau}{K}\right) - \ln\left(1-\frac{N_0}{K}\right)\right] + \frac{\mu_2 \pi_2 K}{|d_2|}\left[-\ln\left(1-\frac{N_\tau}{K}\right)\right] \tag{A.38}$$

$$= -\frac{\mu_1 \overline{\pi} K}{|d_1|}\ln\left(1-\frac{N_0}{K}\right) + \ln\left(1-\frac{N_\tau}{K}\right)\left[\frac{\mu_1 \overline{\pi} K}{|d_1|} - \frac{\mu_2 \pi_2 K}{|d_2|}\right] \tag{A.39}$$

Since N_τ takes values between 0 and N_0, we again have the same result, with minimum/maximum P_E at $N_\tau = 0$ depending on the parameter values. The conditions for the same are described in Eqs A.35.