Dynamics of infection, mutation, and eradication, in HIV and other evolving populations

Citation
Rosenbloom, Daniel Scholes. 2013. Dynamics of infection, mutation, and eradication, in HIV and other evolving populations. Doctoral dissertation, Harvard University.

Permanent link
http://nrs.harvard.edu/urn-3:HUL.InstRepos:11157925

Terms of Use

Share Your Story

Dynamics of infection, mutation, and eradication, in HIV and other evolving populations

Thesis advisor: Martin A. Nowak   Daniel Scholes Rosenbloom

Dynamics of infection, mutation, and eradication, in HIV and other evolving populations

Abstract

This work uses mathematical models of evolutionary dynamics to address clinical questions about HIV treatment, public health questions about vaccination, and theoretical questions about evolution of high mutation rates.

Chapters 2 and 3 explore HIV treatment. Despite the high inhibition of viral replication achieved by anti-HIV drugs, many patients fail treatment, often with emergence of drug resistance. The observed relationship between adherence and likelihood of resistance differs across drug class. Chapter 2 presents a model that explains these observations by considering drug properties, fitness of susceptible and resistant strains, and adherence. Poor adherence to boosted protease inhibitors, which have sharp dose-response curves and short half-lives, is shown to cause failure via growth of susceptible strains, not resistant ones.

Current HIV treatment cannot eradicate the infection due to a reservoir of latent virus in resting memory $CD4^+$ T cells. Chapter 3 models infection dynamics during treatment interruption, in order to assess investigational therapies that reduce the size of the latent reservoir. Calculations suggest that reducing the reservoir by 1,000- to 20,000-fold will let half of patients interrupt treatment for one year without rebound, but that rebound may occur suddenly after years of success.

Chapter 4 considers vaccination against seasonal infections. One's decision to

vaccinate affords indirect protection to non-vaccinators. Reliance on this protection may prevent establishment of herd immunity. Here epidemiology is integrated into a model of adaptive learning, in which individuals use anecdotes from peers to estimate benefits of vaccination. Herd immunity is found to establish in network-structured populations, but to break down if vaccination costs exceed a critical threshold. This result suggests parallels to historical "vaccine scares" following periods of high vaccination coverage.

The final two chapters examine evolution of mutation rates under frequency-dependent competition. Cyclical "rock-paper-scissors" competition is found to exert upward selective pressure on mutation rates. Competition resulting in a stable equilibrium trait distribution exerts downward selective pressure on mutation rates. Recombination lowers the evolutionarily stable mutation rate but may permit stable coexistence of rates above and below this level. Biological scenarios are discussed which may meet theoretical requirements for the adaptive evolution of high mutation rates.

Contents

Author List

The author for Chapter 1 is Daniel S. Rosenbloom, incorporating material derived from abstracts written by authors of the subsequent five chapters.

The authors for Chapter 2 are Daniel S. Rosenbloom, Alison L. Hill, S. Alireza Rabi, Robert F. Siliciano, and Martin A. Nowak.

The authors for Chapter 3 are Alison L. Hill, Daniel S. Rosenbloom, Feng Fu, Martin A. Nowak, and Robert F. Siliciano.

The authors for Chapter 4 are Feng Fu, Daniel S. Rosenbloom, Long Wang, and Martin A. Nowak.

The authors for Chapter 5 are Daniel S. Rosenbloom and Benjamin Allen.

The authors for Chapter 6 are Benjamin Allen and Daniel S. Rosenbloom.

Acknowledgments

I would like to thank all who have made this work possible. Alison Hill, Ali Rabi, Feng Fu, and Ben Allen have been simply outstanding collaborators. We have plotted, calculated, celebrated, traveled, and (not too often) commisserated. I have learned so much while working with them and cannot imagine what graduate school would have been like without them.

I especially would like to thank my advisor, Martin Nowak, who introduced me to mathematical biology first through his writing, and then personally. His guidance has been inspiring to me, and his sense of humor is unmatched.

My co-advisor, Bob Siliciano, has welcomed me warmly as a visitor in his lab, and he has been generous with his attention and wisdom. Working with him and his team has been most rewarding.

Feng Fu's Ph.D. advisor, Long Wang, was especially generous with his time, hospitality, and advice as we worked on Chapter 4.

I am grateful to have had John Wakeley and Michael Desai as members of my thesis committee. They have been insightful and helpful teachers. In particular, Dr. Wakeley's population genetics seminars have been a highlight of my graduate education.

I am indebted to those who have read and commented on parts of this dissertation. In particular, I thank Tibor Antal, Ivana Božić, Joel Gallant, Yoh Iwasa, Chris Marx, J-B Michel, Pleuni Pennings, Maame Sampah, Lin Shen, Olivier Tenaillon, Mary Wahl, and Scott Wylie for giving generously of their time

and effort. I also thank Emily Chakwin for watching and critiquing presentations of my research.

For their advice and support in navigating academic administration, I thank May Huang and Michael Wojcik of the Program for Evolutionary Dynamics, Chris Preheim and Jason Green of the Department of Organismic and Evolutionary Biology, and Bonita Grant of Johns Hopkins University. Their patience is invaluable.

My time at the Program for Evolutionary Dynamics has been a Party Every Day. In addition to those I have already mentioned, I thank Anna Dreber Almenberg, Michael Manapat, Stefany Moreno Gámez, Jordi van Gestel, Oliver Hauser, Moshe Hoffman, Jan Humplik, Christina Pawlowitsch, Thomas Pfeiffer, Dave Rand, Corina Tarniță, and Matthijs van Veelen for making this so. I am particularly thankful that Erez Lieberman Aiden, Erick Matsen, and Franziska Michor introduced me to the group during my undergraduate days.

My students in Math 153 and Math 243 have made teaching a delight. I am especially honored to have had the opportunity to advise Nathan Georgette, Lee Seligman, and Ellen Rice on their undergraduate theses.

Generous friends have welcomed me into their homes, both in Cambridge and Baltimore, enabling me to travel during the latter part of my studies. To Oliver Hauser, Alison Hill, Greg Laird, Ali Rabi, Dan Ross, Vardit Haimi Samuels and Richard Samuels, and Stephanie Neely and Max Gladstone: my home is your home. I especially thank Sarah Winawer-Wetzel and Catherine Deneke, who have been gracious hosts more times than I can count.

Finally, I would like to thank my loving family. My parents' encouragement has helped me pursue the work that I love. Their exhortations to find time to rest are always appreciated, if often unheeded. My sister (and fellow biologist) Marissa has been ever supportive of me, and her enthusiasm is infectious. Susan, David, and Lexi have been the most wonderful in-laws I could ever wish for, and they too are a source of great strength to me. To Jessie: You have been with me through joys and challenges, and I cannot thank you enough.

1

Introduction

Mathematical models of evolving populations can be used to investigate biological mechanisms, explaining and predicting trends in these populations. In this thesis, I explore a range of biological scenarios with models of evolutionary dynamics. My work addresses clinical questions of HIV treatment and evolution of resistance, public health questions regarding vaccination and herd immunity, and theoretical questions regarding the evolution of mutation rates.

In Chapters 2 and 3, I investigate two questions about current and investigational HIV treatment efforts. Since the early days of HIV treatment,

dynamic modeling approaches have contributed to our understanding of viral growth and evolution under various treatment regimes [33, 55, 239]. Highly active antiretroviral therapy (HAART) inhibits HIV replication and can maintain a low viral load indefinitely, but it cannot eradicate the infection due to a reservoir of latent virus present in resting memory $CD4^+$ T-cells. The treatment battle against HIV is therefore a long one that may involve multiple rounds of treatment modification following the development of resistance. Two recent trends in HIV treatment research have emerged: development of highly potent drugs with relatively few adverse effects that make it easier for patients to adhere to their HAART regimen, and investigation of treatments that may deplete the latent reservoir. I investigate two mathematical models that may be of service in the design and testing of treatments that arise from these trends.

Adherence is typically measured as the fraction of a patient's prescribed doses that he or she takes, though measures such as the frequency of extended treatment interruptions provide complementary information [233, 234]. In the context of the first HIV research trend, it is important to understand the relationship between adherence and treatment outcomes. Despite the high inhibition of viral replication achieved by current anti-HIV drugs, many patients fail treatment — suffering rebounding viral loads, often with emergence of drug-resistant virus. The relationship between adherence and likelihood of resistance differs dramatically across drug class, with boosted protease inhibitors generally exhibiting least vulnerability to resistance, even at sub-par adherence levels [15, 185]. In Chapter 2, I present a mathematical model that explains these

observations and makes novel predictions. This work extends the literature on viral dynamics by incorporating empirical measurements of antiretroviral pharmacokinetics and pharmacodynamics into a model considering adherence, mutation, and fitness differences between susceptible and resistant strains. Drug resistance may emerge either from mutations pre-existing in the latent reservoir or from mutations arising *de novo* from reverse transcription during the viral lifecycle. The model shows that antiviral activity falls quickly for drugs with sharp dose-response curves and short half-lives, such as boosted protease inhibitors, limiting the time between doses when resistance can be selected. Poor adherence to such drugs therefore causes treatment failure via growth of susceptible virus. This chapter also proposes the hypothesis that particular single-pill combination therapies may prevent resistance at any patient adherence level, even though the same drugs administered as separate pills could be vulnerable to resistance. This hypothesis follows from a calculation that the mutant selection window for a particular drug combination is small if the drug concentrations are forced to vary in lockstep — as they would be in a coformulated pill.

The number of potential combination regimens grows rapidly with each investigational compound, yet only a small fraction of these regimens can feasibly be tested in randomized control trials. Moreover, a patient's optimal regimen depends on personal factors such as viral load history, CD4 count, pre-existing resistant variants, and adverse effects of and adherence to treatment, meaning that results of large trials may not generalize to all patients. As standards of care change over time, these aspects of the overall patient population may change as

3

well, potentially destabilizing correlations observed in past clinical studies. Against an evolving adversary, investment in drugs is like investment in the stock market: Past performance does not guarantee future results. Since the modeling approach proposed in this chapter incorporates patient characteristics and underlying — presumably stable — mechanisms of viral growth, it has the potential to predict patient outcomes in a broader range of circumstances than previous approaches.

In the context of the second HIV research trend, great strides are being made to develop reliable therapies for eradicating cells latently infected with proviral HIV. While risky and costly stem cell transplantation remains the only verified method for curing HIV by complete eradication of the latent reservoir [139], both early treatment initiation [29, 243, 296] and use of drugs to induce targeted activation of latently infected cells [8–10, 27, 197, 249, 285, 327, 328] may reduce the size of the latent reservoir. It is not yet known how reservoir reduction, short of complete eradication, will benefit patients. In Chapter 3, I address the specific question of how much reduction of the reservoir is required to allow a patient to maintain low viral loads during extended interruption of HAART. By triangulating measurements of latent reservoir size, activation rate of latently infected cells, and growth rate of virus issuing from actively infected cells, and then by incorporating these measurements in a model of infection dynamics, I offer an initial answer to this question, setting specific goals for investigational therapies that have yet seen limited clinical testing. The model suggests that reduction of 1,000- to 20,000-fold is required to let a majority of patients

interrupt HAART for one year without rebound and that viral rebound may yet occur after multiple years of successful interruption. Investigational drugs thus far have not been shown to reduce the reservoir even tenfold. Readers who are daunted by this finding may wish to recall that AZT, the first antiretroviral drug approved for HIV treatment, inhibits replication by drug-susceptible virus roughly 20-fold as measured *in vitro*, while the first protease inhibitor, saquinavir, approved nine years later, inhibits replication over one million-fold by the same measurement [279], rendering virus undetectable by clinical assays. As mathematical models designed in the early HAART era served to frame debates about therapeutic goals and virological consequences of antiretroviral therapy [122, 239], so too may models of latency reduction serve investigations of this emerging therapy. As an important point of contrast with earlier viral dynamics literature, the work presented in this chapter uses a stochastic model, appropriate for understanding small populations of actively infected cells, the time until first activation, and heterogeneity among patients receiving equally effective therapy. As data from longer-term clinical trials of latency-reducing drugs becomes available, the distribution of times until viral rebound generated by the model may be tested against it, leading to a clearer understanding of the viral dynamics of activation and small infections.

The work in Chapter 4 takes place in an epidemiological setting. I examine the dynamics of the spread of an infection through a community of individuals who make the recurring seasonal decision whether to vaccinate against the infection. Successful vaccination campaigns have been known to seed their own demise, for

achievement of herd immunity through vaccination indirectly protects even those who cannot or who choose not to vaccinate, allowing them to "free-ride" on vaccination efforts of others [24]. As memories of the virulent consequences of a pathogen wane, deniers of vaccination efficacy and safety may gain prominence and credence [208]. In the context of a flu-like infection, against which a new vaccination is required each season, individuals have the opportunity to update their vaccination decision annually and may condition their decision upon information about vaccination costs and benefits that they gathered in previous seasons. In this chapter, I investigated the consequences of a simple rule by which members of a population update their vaccination decisions. Under this rule, an individual compares his or her outcome from last season (sick or healthy) with that of a peer; the individual then switches his or her strategy (vaccinate or not) to that of the peer with a probability depending on whether the peer fared better last season. This rule provides a simple description of learning, by which individuals seek to improve their own outcomes by selectively imitating others. Similar learning rules have been investigated theoretically and experimentally in the context of human cultural evolution [312], have been shown to be optimal under certain conditions of bounded rationality [267], and have been used to describe learned foraging behavior adequately in at least one nonhuman species [245].

The learning rule under investigation gives rise to dueling contagions: after the seasonal epidemic spreads through the population, so too may the strategy of vaccination spread via imitation learning, limiting the following season's

epidemic. Paradoxically, as individuals become more adept at imitating successful strategies, the equilibrium level of vaccination falls below the rational individual optimum. Herd immunity is not established in this setting. In network-structured populations where individuals may transmit the vaccination strategy or infection only to close contacts, the picture is guardedly optimistic: vaccination is widespread over a range of low vaccination costs, but coverage plummets after cost exceeds a critical threshold. This result suggests parallels to historical scenarios in which vaccination coverage provided herd immunity for some time, but then rapidly dropped. The model is therefore capable of representing both the establishment and fragility of herd immunity.

Chapters 5 and 6 explore theoretical issues in the evolution of mutation rates. The adaptive potential of a population is both enabled and constrained by random mutation, and numerous studies have investigated the causes and consequences of mutation rate evolution [14, 295]. In the context of infections, both the mutation rate and intra-host population size of a pathogen govern the rate at which it may escape host defenses, a phenomenon that has been studied in great detail in the context of HIV epitope evolution [94, 220]. In mammals, the adaptive $CD4^+$ T cell response is a form of "domesticated evolution" that exploits strong selective pressures for antigenic specificity in a large population ($\approx 10^{12}$ cells in humans) to combat this shifting pathogen threat. In bacterial species confronting phage parasites, hypermutation may play a similar protective role [227], albeit with the attendant risks of increased mutational load. Theoretical arguments suggest that host-pathogen coevolution may drive

increased mutation rates in both species [150, 201]. Stepping beyond host-pathogen competition, I argue in these two chapters that numerous forms of frequency-dependent competition may mediate both upward and downward selective pressures on mutation rates. Chapter 5 uses numerical simulation to investigate how a basic form of cyclical competition among three or more traits within a population — such as "rock-paper-scissors" competition — favors lineages with high mutation rates that are able to respond quickly to periodic shifts in trait frequencies. Despite these never-ending shifts, the mutation rate in this system generally converges to a unique high evolutionarily stable rate. It is widely known that recombination weakens selection for high mutation by decoupling the genetic determinants of mutation rate from the adaptive consequences of mutation [295]. In agreement with this principle, recombination lowers the evolutionarily stable mutation rate in the model presented. This chapter adds a wrinkle to this principle by showing that recombination also allows for stable coexistence between mutation rates above and below the evolutionarily stable rate. Even considering strong mutational load and ignoring the costs of faithful replication, evolution favors positive mutation rates in the model if the selective advantage of prevailing in competition exceeds the ratio of recombining to non-recombining offspring. In conclusion, I hypothesize that local mutation rates may be relatively high on genes influencing cyclical competition, and that global mutation rates in asexual species may be higher in populations subject to strong cyclical competition.

Chapter 6 abstracts from the particular form of cyclical competition studied in

the previous chapter by analyzing a general mathematical framework of mutation rate evolution under frequency-dependent selection. Using deterministic population dynamics, the evolution of mutation rates can be studied under any form of frequency-dependent competition using an adaptive dynamics approach, whereby a resident mutation rate may be invaded by a small population with a slightly different mutation rate. This framework adequately reflects the evolutionary regime where (1) the mutation rate evolves gradually and slowly relative to the dynamics of trait competition and (2) the population as a whole is expected to produce at least one mutant offspring per generation. Two general principles can be proven in this framework. First, if frequency-dependent competition dynamics lead to a stable mutation-selection equilibrium distribution of traits, then lower mutation rates outcompete higher ones. Second, if competition leads to a particular type of cyclical trait dynamics (a heteroclinic cycle), of which the rock-paper-scissors dynamic is a special case, then the zero mutation rate can be invaded by a range of higher rates. These theoretical results suggest that the previous chapter's conclusions regarding mutation rates on genes influencing cyclical competition or in asexual species under cyclical competition are robust to the quantitative details of competition. The standard-bearer for cyclical competition observed in nature has been the side-blotched lizard *Uta stansburiana*, in which male morphotypes of three different colors engage in a rock-paper-scissors mate competition [289, 290]. Cyclical competition should be common where individuals face fitness tradeoffs in multiple interacting characters [288], suggesting that frequency dependence-modulated mutation

9

rate evolution may be of broad evolutionary importance.

2

Antretroviral dynamics determines HIV evolution and predicts treatment outcomes

2.1 Introduction

The prognosis of HIV infection has dramatically improved since the introduction of highly active antiretroviral therapy (HAART), which, when successful, can bring viral loads below the detection limit, improve immune function and prevent progression to AIDS [229]. Although a complete

understanding of how virologic, pharmacologic and host factors interact to determine therapeutic outcome is still lacking, it is clear that a major obstacle to successful treatment is suboptimal drug adherence. Non-adherence can lead to virologic failure and the emergence of drug resistance [100, 117, 185, 315].

Because of their high antiviral activity, protease inhibitors are crucial in HIV-1 treatment and are used in three of the five recommended initial regimens and many salvage regimens [225]. Clinical trials have shown that for many drug combinations involving protease inhibitors, treatment failure occurs without resistance mutations in the viral gene encoding protease [12, 118, 247, 300], though mutations conferring resistance to other drugs in the regimen are often found. It is generally believed that combination therapy works because it is unlikely that multiple mutations conferring resistance to all drugs in the combination will appear in the same viral genome. Thus, failure without protease inhibitor resistance is puzzling, because it seems to contradict this fundamental explanation for the success of HAART. It is commonly believed that protease inhibitors have a higher 'barrier to resistance' than other drugs, meaning that clinically significant protease inhibitor resistance requires the accumulation of multiple mutations in the protease gene [58]. Protease inhibitor resistance also typically occurs at a narrower range of adherence levels than resistance to other drug classes [100, 155]. Although these concepts are suggestive, no theory has been developed to explain why patients fail protease inhibitor-based regimens without protease inhibitor resistance.

A resistance mutation may exist before treatment in the latent or active viral

populations or may arise during treatment [210]. Drug resistance develops clinically if the mutant strain is selected for over the wild-type strain. Selection depends on the fitness costs and benefits of the mutation, as well as on drug levels, which vary with the dosing interval, the drug half-life and the patient's adherence. Here we use a modeling approach to integrate these factors, enabling us to determine when a resistance mutation will be selected and to predict the outcome of therapy with different drugs. Our results explain the unique adherence-resistance relationship for protease inhibitors and show why patients fail protease inhibitor-based therapy without protease inhibitor resistance.

2.2 RESULTS

2.2.1 DEFINING THE MUTANT SELECTION WINDOW

Antiretroviral drugs reduce viral fitness in a dose-dependent manner (Fig. 2.2.1a). Viral fitness can be summarized as a single parameter, the basic reproductive ratio R_0, which encompasses all phases of the viral life-cycle [215] (Supplementary Methods). The Hill dose-response curve describes the relationship between drug concentration and R_0:

$$R_0 = \frac{R_{00}}{1 + \left(\frac{D}{IC_{50}}\right)^m} \tag{2.1}$$

Here D is drug concentration, IC_{50} is the concentration at which 50% inhibition occurs, and m is a parameter determining steepness of the curve

Figure 2.2.1: Drug concentrations determine the relative fitness of the wild-type virus and a resistant mutant. (a) The fitness of the wild-type virus (R_o, blue line) decreases with increasing drug concentration (here shown normalized by wild-type IC_{50}), following equation (1). A drug-resistant strain (R'_o, red line) is less fit than the wild type at low concentrations but more fit at higher concentrations, owing to an increased IC_{50} or a reduced slope. The MSW is the range of concentrations where a resistant mutant, if present, will grow faster than the wild type and still has $R'_o > 1$. The WGW is the range of low concentrations where the wild type has $R_o > 1$, leading to treatment failure without the need for resistance. For drug concentrations in the overlapping range of these windows, virologic failure can occur even without resistance but will be hastened by the appearance of a faster-growing mutant. (b) As drug concentrations decay after the last dose is taken, the viral fitness passes through four different selection ranges. Depending on the drug, dose level and mutation, not all of these ranges may exist. The time spent in each selection window is also determined by the drug half-life. WT, wild type.

[48, 279]. The numerator R_{00} is baseline fitness in the absence of treatment.

A drug-resistant mutant is any viral variant that is less inhibited than the wild type for some drug concentration, described by the altered dose-response curve that determines viral fitness R_0':

$$R_0' = \frac{R_{00}(1-s)}{1 + \left(\frac{D}{\rho IC_{50}}\right)^{m(1+\sigma)}} \qquad (2.2)$$

Mutations have a fitness cost, meaning that the drug-free fitness of the mutant virus is reduced by a fraction s ($0 < s < 1$). In the presence of the drug, the mutation confers a benefit, multiplying the IC_{50} by a factor ρ (the fold change in IC_{50}, $\rho > 1$). Many mutations also reduce the slope (m) of the dose-response curve by a fraction $\sigma < 0$ (ref. [265]).

Virologic failure occurs when treatment fails to prevent the growth of virus to high levels. A viral strain grows when $R_0 > 1$. The strain with highest R_0 outcompetes others [215]. The range of drug concentrations where a resistant mutant can cause virologic failure is called the mutant selection window (MSW) [79, 80]. Above the MSW, even replication of the mutant is suppressed ($R_0'(D) < 1$), although toxicity may prevent these drug concentrations from being achieved clinically. We here define the wild-type growth window (WGW), where drug concentrations are so low that wild-type virus is not adequately suppressed and failure can occur even without resistance ($R_0(D) > 1$).

2.2.2 THE MSW EXPLAINS THERAPY OUTCOME PATTERNS

To predict how well each drug suppresses growth of resistant and susceptible strains, we computed the time during a treatment interruption that a patient spends in the MSW and WGW. During treatment interruption, both R_o and R'_o increase. Up to four selection ranges can be identified (Fig. 2.2.1b). Using pharmacokinetic and pharmacodynamic data $[265, 279]$ (Supplementary Table 7.1.1), we determined the time spent in these ranges for 66 drug-mutation pairs (Fig. 2.2.2a) on the basis of their specific dose-response curves (Fig. 2.2.2b-e). For each pair, we determined how soon after the most recent dose the mutant or wild-type virus starts to grow. This quantity is shorter than the expected time until virologic failure, which requires plasma HIV RNA to reach detectable levels and may also depend on the time until mutant virus appears. We examined here only single-point mutations that are fully characterized by their effect on the dose-response curve (Eq. (2.2), Supplementary Tables 7.1.2, 7.1.3). For this reason, we caution that our results may be over-optimistic, as virus with multiple resistance mutations often appears during infection. Use of our results for clinical recommendations is therefore premature. Below, we discuss extending the model to multiple mutations.

Successful treatment must both minimize the time spent in the MSW and delay entry into the WGW. These two goals are in tension, as shortening the time spent in the MSW (for example, by decreasing drug half-life) can also hasten
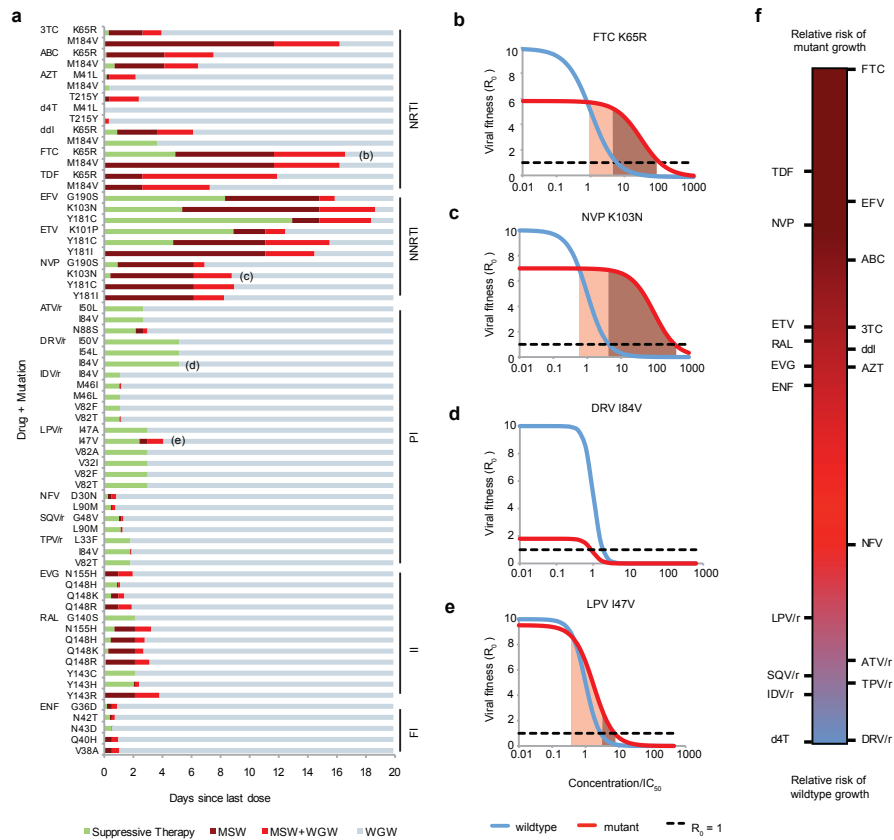
Figure 2.2.2: Selection windows can be calculated for particular drug-mutation pairs. (a) The distance to the right along each horizontal bar is the time since the last dose, and the color corresponds to the selection window during that time interval (described in Fig. 2.2.1b). (b-e) Examples of dose-response curves (showing drug concentration normalized by wild-type IC_{50}) for drug-mutation combinations indicated in a. Shading indicates the MSW. If the cost of a mutation is too high or its benefit (ρ or σ) too low, it is possible that the MSW does not exist. (f) Rank of each drug for relative risk of wild-type versus mutant virus growth, independent of the overall risk of therapy failure. For each drug, we show a 'synthetic', worst-case, single-nucleotide mutation (Supplementary Methods and Supplementary Fig. 7.2.12). PI, protease inhibitors; FI, fusion inhibitors; II, integrase inhibitors; ABC, abacavir; FTC, emtricitabine; ATV, atazanavir; TPV, tipranavir; EVG, elvitegravir; ENF, enfuvirtide. Protease inhibitors are often boosted (co-formulated) with ritonovir (/r), which interferes with breakdown in the liver and increases half-life.

17

entry into the WGW (Fig. 2.2.1b). Results from our model (Fig. 2.2.2a) suggest that non-nucleoside reverse-transcriptase inhibitors (NNRTIs) are protected against failure via wild-type virus due to their long half-lives but are vulnerable to mutation due to the time spent in the MSW. Protease inhibitors are at the opposite end of the spectrum, with little time spent in the MSW but rapid entry into the WGW. This behavior is caused by high slope parameters (extreme sensitivity to changes in concentration) and short half-lives. These results explain the unique trade-off presented by protease inhibitor therapy: greater protection against the evolution of resistance but vulnerability to wild-type-based virologic failure after short treatment interruptions. This feature is depicted schematically by plotting the drugs along a single axis, which measures the relative risk of mutant growth versus wild-type growth, independent of the overall risk of virologic failure (Fig. 2.2.2f and Supplementary Methods).

2.2.3 Simulation of clinical outcomes

Whereas the MSW and WGW concepts describe instantaneous growth of mutant and wild-type virus for a given drug concentration, virologic failure depends on sustained growth and, therefore, drug concentrations over time. To explain clinical observations across drug classes and adherence levels, we developed a stochastic model of viral evolution (Fig. 2.2.3 and Methods). Our model builds on the large body of previous work modeling HIV therapy [215, 259, 293, 318, 325] by integrating new data on class-specific drug properties [279] and realistic costs and benefits of mutations [265]. We also

modified past approaches by allowing drug concentrations, and hence R_o, to fluctuate, rather than taking time-averages.

We first simulated 48-week trials of single agents in a cohort of patients. The results are presented in two ways: as outcome versus patient adherence at the trial endpoint (Fig. 2.2.4a) and as outcome versus time for a distribution of patient adherence levels (Fig. 2.2.4b,c).

Consistent with a previous meta-analysis of combination therapy clinical trials [15], our model predicts that the level of adherence necessary for mutant virologic failure differs by drug class (Fig. 2.2.5). Specifically, for the NNRTIs efavirenz (EFV) and etravirine (ETV), the risk of mutant virologic failure is greatest at low adherence levels; for unboosted protease inhibitors, the risk peaks at a higher adherence level and remains substantial up to 100% adherence; for boosted protease inhibitors (paired with ritonavir to increase half-life), resistance occurs infrequently and at intermediate adherence levels. Researchers have previously argued that drug half-life and fitness costs of mutations are key factors explaining these general trends [100, 155]. By incorporating these factors as parameters, our model formalizes this argument.

In examining simulations of each drug individually (Supplementary Figs. 7.2.1–7.2.7), we found four qualitative patterns of outcome, which correspond closely-but not exactly-to drug class (Fig. 2.2.4).

For most nucleoside reverse-transcriptase inhibitors (NRTIs), the integrase inhibitors, the fusion inhibitor, and the NNRTI nevirapine (NVP), even perfect

Figure 2.2.3: Schematic of algorithm for simulating viral dynamics in a patient undergoing treatment. (a) A single simulated patient takes a particular drug (or drug combination) with a designated adherence level, starting with an initial viral load (VL). Over a 48-week clinical trial, drug levels fluctuate and viral load levels are simulated according to a viral dynamics model. (b) Drug levels fluctuate according to patient's dosing pattern and pharmacokinetics (dose size, half-life, bioavailability); gaps show missed doses (figure shows single drug). (c) Wild-type viral fitness (R_o) fluctuates in response to drug concentration depending on the dose-response curve. (d) Fitness of drug-resistant strain (R'_o) depends on an altered dose-response curve; at high drug concentrations, mutant fitness exceeds that of the wild type. (e) Wild-type viral load depends on viral dynamics equations, which account for active replication, exit from the latent reservoir and competition between strains. (f) A mutant virus may appear (red star) but be below the threshold for detection (dotted red line) before eventually leading to virologic failure.
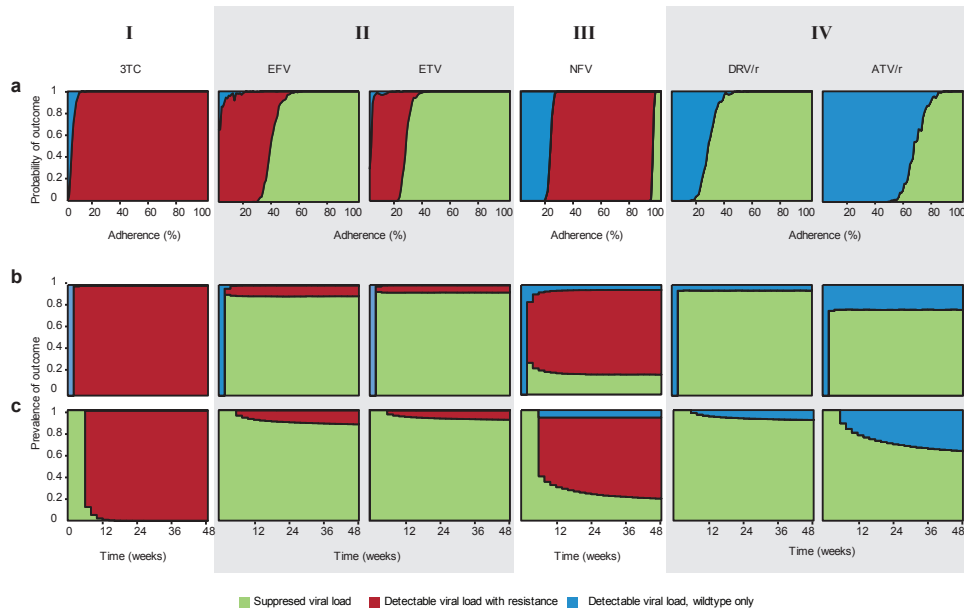
Figure 2.2.4: Outcomes for simulated patients in a clinical trial. (a-c) The height of the area shaded indicates probability of the corresponding outcome at a given adherence level (a) or time point (b,c). (a) Adherence is defined as the fraction of scheduled doses taken. These are maintenance trials (see Methods). (b,c) Measurements are taken every 2 weeks for simulated patients with a distribution of adherence levels (Supplementary Methods and Supplementary Fig. 7.2.13b). (b) Suppression trials (see Methods). (c) Maintenance trials. (1) 3TC therapy (pattern includes AZT, ABC, d4T, ENF, EVG, FTC, NVP, RAL, TDF). (2) EFV and ETV therapy. (3) NFV therapy (pattern includes ddI). (4) DRV/r and ATV/r therapy (pattern includes ATV, TPV/r; variation on this pattern described in the Results includes LPV/r, SQV, SQV/r IDV, IDV/r).

adherence led to mutant virologic failure in all simulated patients. As adherence declined, some wild-type virologic failure occurred. Virologic failure and resistance occurred soon after the trials started. These results are consistent with the notion that monotherapy often leads to rapid evolution of resistance.

For most protease inhibitors and the NNRTIs EFV and ETV, however, perfect adherence resulted in treatment success in simulations. Control of viral replication has been observed in a substantial fraction of patients in protease inhibitor monotherapy trials [241], but similar trials with EFV and ETV have not been carried out. In simulations, declining adherence affected performance of these two drug classes differently.

For the NNRTIs EFV and ETV, there was a large range of low-to-intermediate adherence for which mutant virologic failure was likely. Below this range, wild-type virologic failure became increasingly likely, whereas above this range the simulated patients succeeded. The size of this range is explained by the low fitness costs of drug-resistant mutations and long half-lives of NNRTIs, which allowed the patient to remain within the MSW for a substantial duration (suggested in ref. [17]).

The protease inhibitor nelfinavir (NFV) and the NRTI didanosine (ddI) showed a large range of intermediate adherence leading to mutant virologic failure. Near-perfect adherence was required for treatment success. Under most clinical settings (adherence < 95%), our model predicts that these drugs perform similarly to monotherapy with other NRTIs, typically leading to mutant virologic failure.
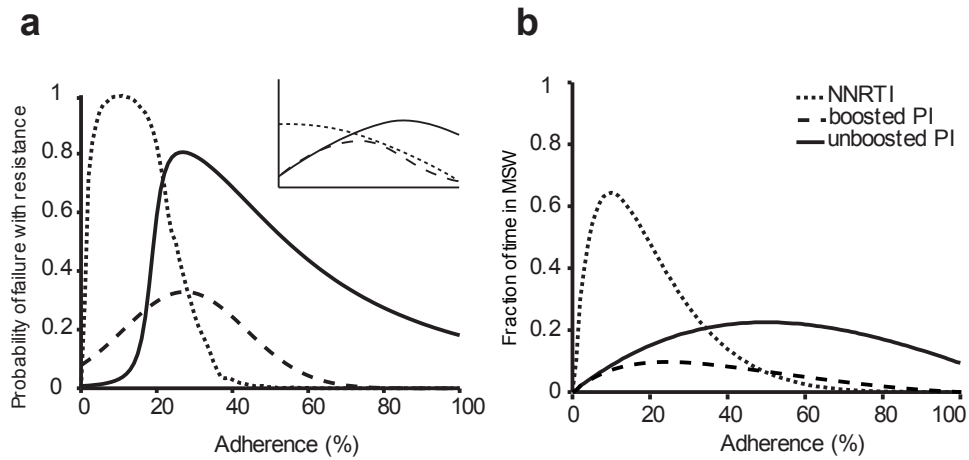
Figure 2.2.5: Our calculated adherence-resistance relations are in agreement with those observed in clinical trials. (a) Adherence versus simulated probability of resistance in a 48-week suppression trial for a protease inhibitor, a boosted protease inhibitor and an NNRTI. The inset shows a qualitative summary of results from a meta-analysis of clinical trials [15], which agrees with our simulations. (b) Adherence versus fraction of time spent in the MSW for the same drugs. Adherence-resistance trends demonstrate that time in MSW is a good proxy for the risk of mutant-based virologic failure. For both plots, curves were generated by averaging over all boosted protease inhibitors, all unboosted protease inhibitors, and the NNRTIs EFV and ETV. Protease inhibitor curves in a were fitted to skewed-T distributions to smooth step-like behavior. NVP, which was excluded from this figure, shows a different pattern from the other two NNRTIs; specifically, mutant virologic failure can occur even for perfect adherence (Supplementary Figs. 7.2.1, 7.2.2).

For many protease inhibitors, a decline from perfect adherence led abruptly from success to wild-type virologic failure, with little or no intermediate range for mutant virologic failure. This result explains the outcomes of clinical studies, which have shown that virologic failure in many boosted protease inhibitor-based regimens (including monotherapy) does not require the evolution of resistance [12, 118, 300]. Variations on this pattern exist for some protease inhibitors: simulations of lopinavir (LPV/r), saquinavir (SQV, SQV/r), and indinavir (IDV, IDV/r) showed mutant virologic failure at low and moderate adherence levels, mainly for trials where initial viral load was high. Still, like all the protease inhibitors simulated except NFV, as adherence declined from the successful range, the first failing outcome observed was wild-type virologic failure (Supplementary Figs. 7.2.1, 7.2.2).

We also examined the sensitivity of our results to changes in the baseline viral fitness, R_{oo} (Supplementary Figs. 7.2.8, 7.2.9). As the intracellular half-lives of several NRTIs are not definitively established, we tested a range of half-lives for lamivudine (3TC), azidothymidine (AZT), stavudine (d4T), ddI and tenofovir disoproxil fumarate (TDF) (Supplementary Fig. 7.2.10). Against a strain with higher R_{oo}, higher adherence levels were required for treatment success, and there was a wider range of adherence levels for which mutant virologic failure occurred. The effect of increasing half-life was drug-dependent, but for most NRTIs simulated, it increased the likelihood of mutant virologic failure.

Equipped with a model of drug interaction, we were able to extend the simulations to combination therapy (Supplementary Methods and Supplementary Fig. 7.2.11). For proof of concept, we use a two-drug combination of the boosted protease inhibitor darunavir (DRV/r) with the integrase inhibitor raltegravir (RAL). The combined effect of these two drugs is given by a Bliss-independent [30] interaction pattern [144], which describes drugs acting on different targets, therefore reducing viral replication multiplicatively. In a recent DRV/r-RAL clinical trial [300], patients experiencing virologic failure had their plasma viral population genotyped. Although 17% of patients tested positive for RAL-resistance mutations in the gene encoding integrase, no patients tested positive for DRV resistance in the gene encoding protease [300]. Our simulation is consistent with this study: treatment failure occurred without DRV resistance (Fig. 2.2.6a).

RAL-resistant mutants were selected for only when the concentration of DRV/r was low and the concentration of RAL was moderate to high (Supplementary Fig. 7.2.11). This state of "effective monotherapy" (ref. [17]) can occur if the drugs are administered as separate pills. If, however, dual therapy were administered as a combination pill, then the two concentrations would rise and fall roughly together, reducing the chance that they reach the discordant levels that select for resistance. Simulation of dual therapy as a single combination pill verified this hypothesis. However, this protection from

Figure 2.2.6 *(following page)*: Outcomes of DRV/r plus RAL dual suppression therapy simulations, considering resistant mutants for both drugs. (a) Each drug is taken independently, and adherence may differ between them. The brightness of each color at a particular point indicates the probability of the corresponding outcome, with the black contours showing where each outcome occurs 95% of the time. Success depends largely on adherence to DRV/r (success is almost certain if adherence is >50%), whereas the type of failure is determined by adherence to RAL (resistance is almost certain if adherence is >30%). All failure via resistance is due to RAL mutant-based virologic failure. DRV mutant-based virologic failure (virologic failure) never occurs in the simulations. (b,c) Drugs are taken with equal average adherence. The height of the area shaded indicates probability of the corresponding outcome at that adherence level. (b) Drugs are taken as separate pills. Average adherence is the same, but pills are taken independently. (c) Drugs are packaged as a combination pill and are always taken together. Mutant virologic failure occurs only when the two drugs are given in separate pills; combination pills eliminate mutant virologic failure but increase the adherence required for near-certain success.

a Dual therapy, differential adherence

b Matched adherence, separate pills

c Combination pill

Therapy success   Failure with resistance   Failure with wildtype only
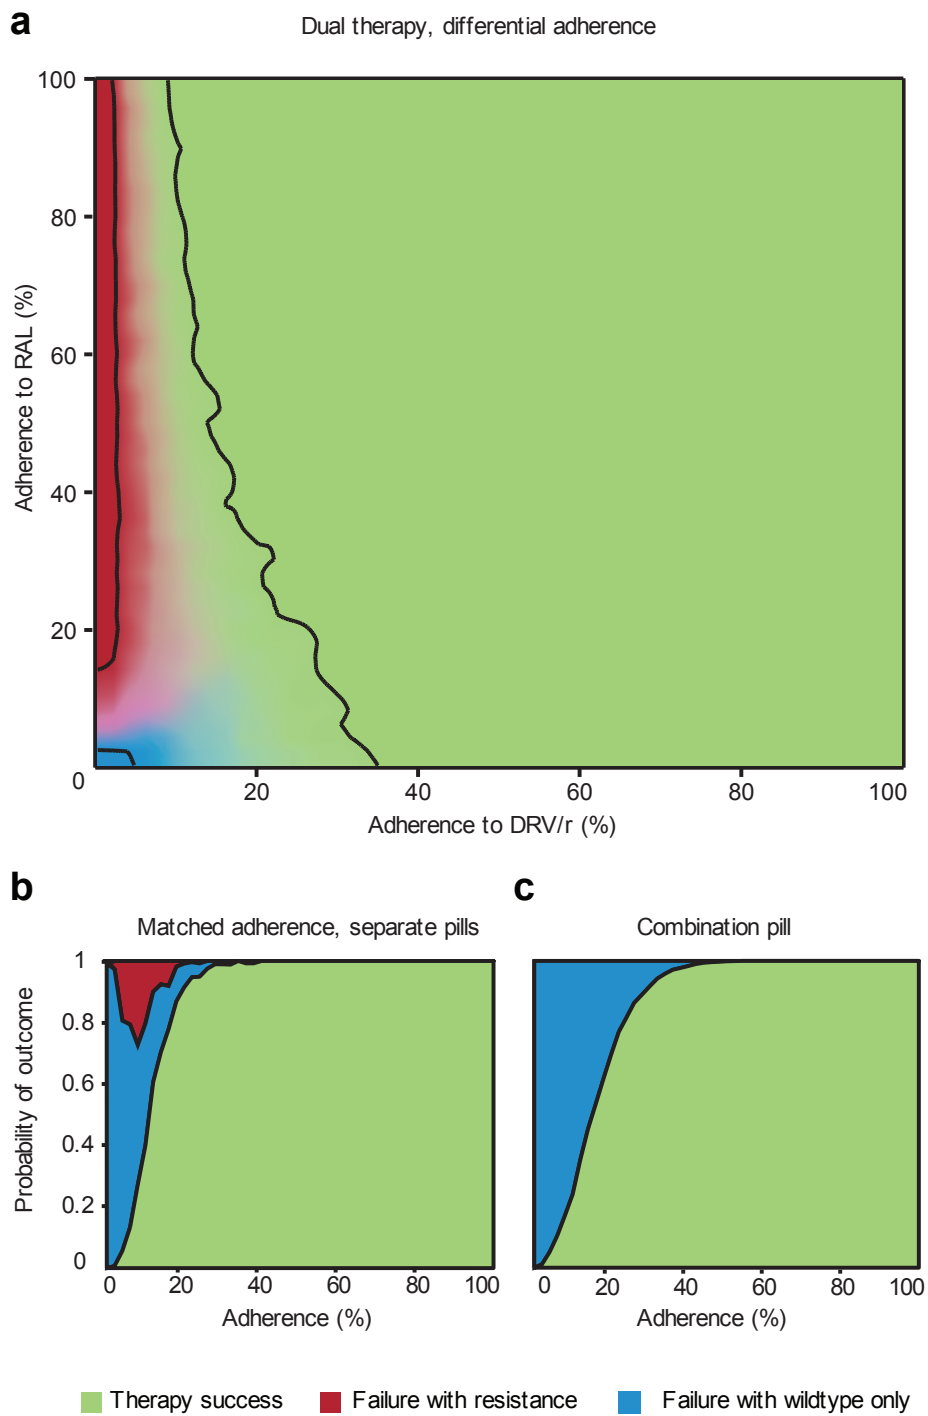
Figure 2.2.6 *(continued)*

27

resistance came at a cost: higher adherence was required to prevent wild-type virologic failure. For example, to ensure a 95% chance of success in the simulation, a patient taking separate pills must be 25% adherent to each pill (Fig. 2.2.6b), but 35% adherent to a combination pill (Fig. 2.2.6c). We expect this trend to apply to other drug combinations.

2.3 Discussion

Recent efforts to quantify pharmacodynamics [265, 279, 280], combined with insights into patients' drug-taking behavior [16], have enabled us to develop what is to our knowledge the first explanatory model of virologic failure in agreement with clinical trials. All parameters in our model have direct physical interpretations, and their values were taken directly, or derived from, previous literature. The model was not fit or trained to match clinical data. Despite our model's simplicity, it can explain the clinically observed drug-class-specific relationship between adherence and outcome [15] (Fig. 2.2.5). Even without full viral dynamic simulations, a straightforward analysis of the mutant selection window can explain why certain drugs are more likely to select for resistance (Figs. 2.2.2f and 2.2.5b).

In addition, we address a long-standing mystery of antiretroviral therapy. Even when failure of protease inhibitor-based regimens is documented, mutations that confer resistance to the protease inhibitor appear infrequently [12, 118, 247, 300]. Although it is possible that mutations may occur outside the

protease-encoding gene [62, 112, 209, 235] and escape routine detection, our model provides a more straightforward explanation: due to the sharp slope of protease inhibitor dose-response curves [279], even relatively strong protease inhibitor resistance mutations are selected only in a narrow range of drug concentrations. Moreover, as protease inhibitor concentrations decay rapidly compared to other drugs, they traverse this narrow range quickly, leaving little time for a resistant strain to grow before wild-type-based virologic failure. We predict that patients who fail protease inhibitor therapy with wild-type virus should be able to re-suppress the virus if the same drug is taken with improved adherence. A previous study [155] observed this outcome in patients who failed LPV/r without detectable resistance. Even with protease inhibitors that are more susceptible to resistance, only wild-type virus is detectable when adherence dips below the level guaranteeing success, providing an antiresistance 'buffer' that may warn clinicians of resistance risk. NFV is the sole exception to this pattern, owing to its having the lowest slope and second-highest IC_{50} of the protease inhibitors and consistent with its documented vulnerability to resistance [155].

The tradeoff between protection from resistant and susceptible strains occurs not only between drug classes but also between different formulations of the same drugs. We predict that a new combination pill containing DRV/r and RAL would not lead to resistance, even though the current separate-pill formulation does. This result suggests that some combination pills may be 'resistance proof', but their known benefit of increasing patient adherence must be weighed against the fact that they require higher adherence to prevent wild-type-based virologic

failure. This tradeoff results from the possibility that a patient who is prescribed multiple pills may at times take only some of them [99], providing partial protection from the virus but allowing entry into a 'zone of monotherapy' [17] that can select for resistance.

We can extend our model to a broader range of combination therapies once interactions between drugs [144] are characterized; these interactions affect the evolution of resistance [202]. Our monotherapy results are a first step for examining how pharmacokinetics and pharmacodynamics determine treatment outcomes. These results can inform innovations in lower-cost maintenance therapy among highly adherent patients, for whom monotherapy shows promise but also poses resistance risks [241]. Specifically, on the basis of our simulations, we propose that EFV and ETV monotherapy may be promising avenues for further study, despite the disheartening performance of monotherapy with the first approved NNRTI, NVP [44], and the ambiguous performance of ETV-based HAART for patients with resistance to the NRTI backbone [264].

Simulations that start with a high viral load (suppression phase) and simulations that start with an undetectable viral load (maintenance phase) generally showed similar outcomes; however, for several drugs, failure with resistance was more likely during the suppression phase. Such differences are often attributed to the presence of preexisting mutants when viral load is high [33, 149, 232, 287]. However, in our model, frequent reactivation from the latent reservoir provides a sufficient source of mutants during both phases (Supplementary Tables 7.1.4, 7.1.5), and ongoing replication is an additional

common cause of resistance (Supplementary Figs. 7.2.6, 7.2.7). The key difference between the two phases is in how virologic failure is defined. As patients remained in suppression simulations until the predefined endpoint, wild-type growth sometimes preceded (and contributed to) growth of the mutant. More frequent measurement of viral load in maintenance simulations improved the chance that virologic failure was diagnosed before resistance reached detectable levels, consistent with clinical meta-analysis [111]. Also consistent with clinical observations [137], continuation of maintenance trials after rebound allowed the possibility of re-suppression, but it sometimes led to emergence of resistance (Supplementary Fig. 7.2.5).

It is difficult to quantitatively compare our simulations to clinical trials, as adherence is rarely precisely known. We suspect that our results are biased toward success for several reasons. First, we considered only single-point mutations, but strains with multiple mutations may lead to failure at higher adherence levels. Second, we considered neither correlations between consecutive missed doses nor variations in the time of day when a dose is taken, both factors that lead to longer treatment interruptions and increase the chance of virologic failure [153, 182, 233, 234, 318]. Third, as is common in models of viral dynamics, we assumed that the virus population is homogeneous and well mixed. Actual infections may include subpopulations that grow faster (higher R_0, for example, owing to cell-to-cell transmission [281]) or that reside in tissues that drugs do not fully penetrate [156, 189, 269]. For example, the concentration of EFV in the cerebrospinal fluid is only 0.5% of plasma concentrations [26]. As our

predictions rely on plasma drug concentrations, they may be optimistic in the case of EFV (see ref. [123] for further discussion). In the absence of strong evidence for these effects, suboptimal adherence is the most likely cause of treatment failure. Given the above limitations, our modeling results should not be taken as clinical recommendations at this stage.

Patients experiencing virologic failure may not respond to a similar regimen in the future [124, 180, 233], but the precise reasons for this are not clear. The simplest explanation is that growth of a resistant strain during prior treatment makes it more likely this strain will exist in the future [32]. This explanation assumes that, in the absence of prior growth, most resistant mutants are relatively rare. If the diversity (effective population size) of the latent reservoir is not severely depleted over time, then our calculations contradict this assumption for single mutations: even in the absence of prior treatment, a majority of mutations exit the reservoir every few weeks. Resistance is then available to be selected regardless of prior growth. The occurrence of multiple mutations within the same viral genome is unlikely, however, without prior growth. To explain generally how prior virologic failure undermines future treatment, we need to model the long-term accumulation of multistep mutations in the viral population [219, 277]. To build such models, it will be important to understand interactions between mutations (including compensatory mutations [123]) and account for recombination [206].

We have emphasized here the variable nature of anti-HIV drug resistance. Common practice classifies a genotype as resistant if it is associated with

virologic failure in a meta-analysis of clinical outcomes; otherwise it is sensitive. This categorization is misleading: a mutation's ability to promote viral growth depends on all of the drugs in a regimen, adherence and the other mutations present. As standards of care evolve and study populations change, a mutation may gain or lose resistant status as a result of shifts in these confounding variables. Our model provides a rigorous alternative for evaluating resistance, by using mechanistic parameters to predict clinical outcomes. Our framework can help researchers prioritize drugs for clinical trials and select regimens for personalized HIV treatment.

2.4 METHODS

2.4.1 PHARMACOKINETICS, PHARMACODYNAMICS AND THE MUTANT SELECTION WINDOW

Viral fitness followed equation (2.1) with parameters R_{00}, IC_{50} and m. Fitness of resistant mutants followed equation (2.2) with parameters s, ρ and σ. (Supplementary Tables 7.1.1–7.1.3). Relative wild-type and mutant viral fitness values $R_0(D)/R_{00}$ and $R'_0(D)/R_{00}$ were measured using *in vitro* assays and were fit to Hill curves to determine the parameters IC_{50}, m, σ, ρ and s; these values were reported previously [265, 279]. We estimated absolute *in vivo* viral fitness in the absence of drugs (R_{00}) using measurements from previous studies (Supplementary Methods). We modeled drug concentration as instantaneously increasing after a dose to the steady-state peak concentration (C_{max}) and then

decaying exponentially (with half-life $T_{1/2}$) to the trough concentration (C_{min}) before the subsequent dose. When doses were missed (representing suboptimal adherence), the concentration continued to decay, and a subsequent dose increased the concentration by $\Delta C = C_{max} - C_{min}$.

We determined the bounds of the MSW by solving for D in $R_o(D) = R'o(D)$ and $R'o(D) = 1$. We determined the upper bound of the WGW by solving $R_o(D) = 1$. We computed the time after a single dose when a particular concentration D was reached by solving for t in $D = C_{max}2^{-t/T_{1/2}}$.

The MSW concept as applied here to antiretroviral therapy was adapted from the extensive literature on antibiotic resistance. Both *in vitro* and *in vivo*, drug concentrations that fluctuate within the MSW lead to the development of resistance, but those outside it do not (reviewed in ref. [80]). Although some studies of antibiotic-resistant *Escherichia coli* have found no upper limit to the MSW [329], no such results are known for antiretroviral resistance. The definition of the MSW most commonly used in antibiotic work is slightly different from the one we use, with the lower limit defined as $R_o(D) = 1$ because of experimental constraints [79]. We have chosen to modify this definition, as selection for the mutant can occur even at lower drug concentrations where $R_o(D) > 1$ (ref. [110]). The MSW and WGW can be described for each drug during combination therapy (Supplementary Methods).

Our model for HIV dynamics during antiretroviral drug treatment uses equations common in the literature [215]. These equations track the number of uninfected $CD4^+$ cells, amount of free virus and number of infected $CD4^+$ cells. A constant number of uninfected cells are produced each day, and they die at a constant rate. Cells are infected at a rate proportional to the number of uninfected cells, the amount of virus, and the viral fitness. Virion production from infected cells is described by the burst rate, and virions are cleared at a constant rate. Infected cells have a higher death rate than uninfected cells. Additionally, we include a population of long-lived infected cells in the latent reservoir, which activate at a constant daily rate regardless of viral fitness. Because we are interested only in viral dynamics during treatment and at the initial stages of failure, we have ignored the effects of the immune response. Viral fitness, and hence the rate of infection of new $CD4^+$ cells, is determined by the baseline R_o and the drug concentration. All equations and parameters are given in the Supplementary Methods and Supplementary Table 7.1.6. In the Supplementary Methods, we also derive a simplified form of HIV dynamics that requires fewer parameters and only one state variable per viral strain; we used this simplified model to design our simulations. More detailed models that explicitly track multiple stages of the viral life cycle may more accurately reflect some short-term dynamics, such as lags in viral growth during acute infection or lags in viral decay during the early days of treatment [256, 273]. Summarizing viral fitness by a single parameter (R_o) smoothes out these dynamics.

There may be multiple strains of virus (wild-type and mutants) and consequently multiple types of infected cells. Even in the absence of drug, mutations will arise due to random errors in replication, though they will be selected against due to their fitness cost (s). Each mutation appears at a rate u that depends on the particular nucleotide changes required to effect the desired amino acid substitution (Supplementary Tables 7.1.2, 7.1.3, 7.1.7). The balance between these two processes results in all mutations being present in the population at an expected low level u/s, called mutation-selection equilibrium [215, 255]. We assume that the plasma virus population reaches this equilibrium in each patient before treatment (that is, that sufficient time has passed between initial infection and treatment initiation and that no prior treatment has selected for resistance to the particular drug being studied) and that the population in the latent reservoir is representative of the plasma population (Supplementary Tables 7.1.4, 7.1.5). *De novo* mutations occur with a probability u during replication.

We used stochastic simulations to study the dynamics of the system described. Many mutations have been characterized for each drug, and to model a realistic worst-case scenario we considered a single synthetic mutant defined as having the highest benefits (ρ, negative σ), lowest cost (s), highest mutation rate and highest equilibrium frequency (due to mutation-selection balance) of all the single-nucleotide mutants known for that drug. Each monotherapy simulation therefore tracked only two strains, wild-type and mutant. For dual therapy, we considered three strains: wild-type, resistant to drug 1, and resistant to drug 2. Simulations modeled 48-week trials, using discrete time-steps of $\Delta t = 30$ min. All

simulations were done in Matlab R2010b. The full details of the algorithm for simulating a single patient are given in the Supplementary Methods.

In maintenance trials, patients began with full viral suppression (2 RNA copies per ml, $c\ ml^{-1}$) and underwent monotherapy for 48 weeks or until virologic failure, whichever occurs first. virologic failure is defined as 'confirmed rebound': two consecutive weekly measurements (starting at week 5) with viral load above $200\ c\ ml^{-1}$. In suppression trials, patients began with a realistic distribution of treatment-naive viral loads (between 3,000 and $10^6\ c\ ml^{-1}$) (Supplementary Fig. 7.2.13a) and underwent monotherapy for a full 48 weeks. We tracked measurements every 2 weeks. Virologic failure is defined as a viral load above 50 $c\ ml^{-1}$ at week 48. In both types of trials, virologic failure is classified as 'with resistance' if at least 20% of the viral population at the time of detection is mutant.

We simulated imperfect adherence by allowing each dose to be missed with a constant probability given by the expected adherence level parameter. In reporting outcomes versus time, we simulated patients with a distribution of adherence levels taken from a study using unannounced pill counts30. For simulations with two drugs, the value of adherence may be different for each drug, allowing for "differential adherence," which has been observed in many studies [99]. Even when adherence to the two drugs has the same average value, the drugs can be simulated as two separate pills (allowing each pill to be taken or forgotten independently) or as a single combination pill (causing the two drug concentrations to rise and fall in lockstep).

2.5 SUPPLEMENTARY METHODS

2.5.1 VIRAL DYNAMICS MODEL

The following system of equations models the dynamics of multiple strains $(i = 1, 2, \ldots, n)$ of HIV in a patient:

$$
\begin{aligned}
\dot{x} &= \lambda - \sum_{i=1}^{n} \beta_i x v_i - d_x x \\
\dot{y}_i &= \beta_i x v_i + A_i - d_y y_i \\
\dot{v}_i &= k_i y_i - d_v v_i
\end{aligned}
\tag{2.3}
$$

where state variables x, y_i, and v_i are the number of infectable CD4$^+$ T-cells, the number of actively infected cells of strain i, and the number of free virus particles of strain i, respectively. The number of latently infected cells is considered to be constant, as it doesn't decay significantly over the course of a clinical trial, and so latently infected cells of strain i activate at a constant rate A_i. Active cells produce virus at rate k_i and die at rate d_y, and virus is cleared at rate d_v. The infectivity parameter β_i determines the rate at which virus of strain i infects susceptible host cells. Host cell dynamics are determined by production rate λ and death rate d_x.

When $A_i = 0$ for a strain i, this model reduces to the traditional viral dynamics model [215]. For that model we can describe the *basic reproductive ratio*, which is defined as the number of new infections generated by a lone infected cell before it

dies. Strain i will only have a positive growth rate and be capable of sustaining an infection if its basic reproductive ratio, $R_{oi} := \lambda \beta_i k_i / (d_x d_y d_v)$, is greater than 1. In the model we present here the latent reservoir provides a constant source of virus (A_i), which removes the threshold criteria for R_o, although this value still describes viral fitness and the amount of ongoing viral replication.

For a single strain, the unique non-negative steady-state solution to our model is

$$
y_1 = \frac{\lambda}{2 d_y R_{o1}} \left[R_{o1} \left(\frac{A_1}{\lambda} + 1 \right) - 1 + \sqrt{R_{o1}^2 \left(\frac{A_1}{\lambda} + 1 \right)^2 + 2 R_{o1} \left(\frac{A_1}{\lambda} - 1 \right) + 1} \right]
$$

$$(2.4)$$

In our model, for $R_{oi} > 1$, strain i grows to a high steady state that depends on availability of host cells and the abundance of other strains. There are several limiting cases that can be derived from equation (2.4). In the absence of other strains (or if $R_{oj} \ll 1$ for all $j \neq i$), and for small reactivation $A_i \ll \lambda$, strain i grows to the steady state $y_i \approx \tilde{Y}_i := \lambda (R_{oi} - 1) / (d_y R_{oi})$. The value $\tilde{Y}_i$ is the setpoint viral load that is maintained by replication alone, without additional contribution from the latent reservoir. The residual active infection maintained by the latent reservoir in complete absence of viral replication $(R_{oi} = 0)$ is $\tilde{y}_{oi} := A_i / d_y$. For positive $R_{oi} < 1$, strain i reaches a low steady state $y_i \approx \tilde{y}_i := \tilde{y}_{oi} / (1 - R_{oi})$. Since anti-HIV drugs act by decreasing β_i and k_i, the value of R_{oi} is understood to depend on the current drug concentration(s).

To eliminate some of the model parameters and smooth the high-frequency

fluctuations that may have little clinical impact over the course of a drug trial, we study a simplified version of the model in equation (2.3). We assume that v_i and x are at equilibrium relative to y_i. This allows us to derive a reduced n-dimensional model:

$$\dot{y}_i = A_i + d_y y_i \left[\frac{\lambda R_{oi}}{\lambda + \sum_{j=1}^{n} R_{oj} d_y y_j} - 1 \right] \qquad (2.5)$$

When the total infection is small, the summation term vanishes, and $\dot{y}_i \approx A_i + d_y y_i (R_{oi} - 1)$. For $R_{oi} \ll 1$, nearly all of strain i is produced by exit from the reservoir; y_i therefore approaches a value near $\tilde{y}_{oi}$. As the total infection grows (assuming $R_{oi} > 1$ for one or more i), the fractional term approaches 1, describing saturation of the limiting resource, at which point new infection events are balanced precisely by death of infected cells and y_i approaches a value near $\tilde{Y}_i$. This reduced model has identical steady state values of virus and $CD4^+$ cells as the full model, but smooths out fluctuations in infection size caused by the dynamics of total $CD4^+$ cells. Because we focus on initial virologic failure, which occurs at relatively low viral loads, the fluctuations in $CD4^+$ cell levels are minor, and the approximation captures the full dynamics (equation (2.3)) well.

We can account for mutation by including the mutation rate matrix Q, where Q_{ij} describes the probability that an infected cell of type j gives rise to one of type i:

$$\dot{y}_i = A_i + \frac{\lambda d_y \sum_{j=1}^{n} y_j R_{oj} Q_{ij}}{\lambda + \sum_{j=1}^{n} R_{oj} d_y y_j} - d_y y_i \qquad (2.6)$$

2.5.2 MODEL PARAMETERS

The value of R_{oi} at each point in time depends on the baseline basic reproductive ratio ($R_{oo} = 10$, see below), the current drug concentration(s), and parameters describing resistance of the strain, as described by equations (2.1) and (2.2) in the main text. The death rate of actively infected cells, d_y, is 1 per day [188]. Supplementary Table 7.1.6 summarizes the parameters used in the model.

BASIC REPRODUCTIVE RATIO

The basic reproductive ratio (R_o) combines various components of viral fitness into a single number. $R_o > 1$ is required for the virus to have a positive growth rate and sustain an infection. The baseline R_o, which we denote R_{oo}, is defined in the absence of drug and has been estimated in past studies by measuring the increase in viral load during the early days of acute infection or during planned treatment interruption. During the acute phase, before the CTL response develops, typical values for R_{oo} are 10-20(ref. 64;67). After this initial phase, R_{oo} declines to 2-5, with some outliers as high as 6-11 (ref. 68-72) . Based on these findings, we chose a value of R_{oo}=10 to present our results. We also checked sensitivity to this parameter by using larger and smaller R_{oo} values (Supplementary Figures 7.2.8 -7.2.9).

We can also double-check that our value of R_{oo} from the literature is consistent with an independent set of measurements. The growth rate of a mutant strain in the absence of drug is $R_{oo} * (1 - s)$ (see equation (2.1) in the main text), where s is the reduction in the replication capacity of the mutant virus. If $R_{oo} * (1 - s) > 1$, then a mutant strain will expand in the absence of drug. If this condition fails, then the mutant strain would never be detected at high abundance (ignoring secondary or compensatory mutations). Since all the resistance mutations that we study do occur clinically, we expect that $R_{oo} > 1/(1 - s)$ should almost always hold. 95% of the mutations studied have $s < 0.9$, for which the positive growth condition is satisfied for the value $R_{oo} = 10$.

To maintain consistency with the chosen value $R_{oo} = 10$, we capped the cost of mutations used in the viral dynamics simulation at $s = 0.9$, guaranteeing that no mutant's baseline R_o would be less than 1. Values of s that are negative are also inappropriate for our model, as they imply that the resistant mutant is more fit than the wild type even in the absence of the drug, causing the mutant to be prevalent at baseline. Measurements of s that were close to 0 or negative were assumed to be caused by experimental error, and so we set these values to $s = 0.05$ to represent a small cost to these mutations.

Latent reservoir exit rate

Based on the following argument, we estimate the total reservoir exit rate $\sum_i A_i$ to be 3000 cells per day. The exit rate for a particular mutant strain is determined by multiplying by the equilibrium frequency of pre-existing mutants, u/s. (Our

simulation treats each exit as an independent event; use of this modeling approach implicitly assumes that the reservoir was seeded by a large, diverse population, and that its diversity, or effective population size, is maintained over time.) Viral loads of around 2 RNA copies per mL are maintained in patients on maximally suppressive HAART [74]. The rate of exit from the reservoir must be enough to account for this residual viral load, since ongoing replication is negligible. This viral load corresponds to $\approx 3 \times 10^3$ plasma virions (for a 70 kg person with 3 L plasma). It has been shown, for a wide range of viral loads, that the total number of infected cells in a patient is roughly equal to the number of plasma virions [126]. The infection size $\sum y_i \approx \left(\sum A_i \right) / d_y$ is therefore 3×10^3, implying a total reservoir exit rate of 3000 cells per day.

Alternately, we can estimate the number of infected cells by noting that total viral production (burst from infected cells) must balance total viral clearance (breakdown of free virus in lymphatic tissue). Using parameters previously established [64], free virus in lymph tissues is 100 times as abundant as virus in the extracellular fluid, and so would be about 1.5×10^6 virions (based on 15 L ECF) for this example. This paper also determined that the ratio of viral burst size to viral clearance rate is typically 500 virions per cell (e.g., $k_i = 10{,}000$ virions per day per cell; $d_v = 20$ per day). These figures again imply an infection size of 3000 cells.

Our calculations also agree with the results of a model which examined the many years-long decay of the latent reservoir in HAART patients [272]. Although this model used different sources for parameter values, it is consistent

with an exit rate of 3000 cells per day, as long as the reservoir is not significantly depleted.

HOST CELL PRODUCTION RATE

For a single wild-type strain in the absence of drug, the model (equation (2.5)) provides $\lambda = \tilde{Y} d_y R_{oo} / (R_{oo} - 1)$, where $\tilde{Y}$ is the total number of infected cells at infection setpoint. As established above, this value is approximately equal to the number of plasma virions at setpoint. We considered setpoint viral loads from 3000 to 10^6 RNA copies per ml plasma, or 4.5×10^6 to 1.5×10^9 total plasma virions. These values give a range of 5×10^6 to 1.7×10^9 cells per day for λ.

RESISTANCE MUTATION RATES

The mutation rate matrix entry Q_{ij} describes the probability that strain j reproduces to create strain i. We include only single step mutations from the wild type $(j = 1)$ to another strain i (at rate u_i) and ignore back-mutation. Therefore $Q_{i1} = u_i$ for $i > 1$, $Q_{1,1} = 1 - \sum_{k=2}^{n} u_k$, $Q_{ii} = 1$ for $i > 1$ and $Q_{ij} = 0$ for all other entries.

The overall mutation rate for HIV is 3×10^{-5} per base per replication cycle [187], and recent work has shown that the rate varies considerably depending on the specific base changes involved. The nucleotide mutation matrix used in this study was derived by normalizing mutation accumulation data from a study of HIV replication of lacZa reporter sequence [2]. The normalized data was then rescaled to convert from the lacZa base composition to the HIV consensus

sequence base composition [170]. Specifically:

1. Define the variables:

   - $u = 3 \times 10^{-5}$ is the average per-site mutation rate of HIV.

   - s_{xy} is the total number of single-nucleotide substitutions from base x to base y, combining data from both the forward and reverse orientations of lacZa in Table 3A of Abram et al. [2].

   - s_{x*} is the total number of single-nucleotide substitutions from base x to any other base.

   - S is the total number of single-nucleotide substitutions overall.

   - n_x and n'_x are the abundance of base x in the reporter sequence and in the HIV consensus sequence, respectively. N and N' are the lengths of the two sequences, respectively.

     - $n_T = 37, n_C = 56, n_A = 36, n_G = 45; N = 174$

     - $n'_T = 2163, n'_C = 1772, n'_A = 3411, n'_G = 2373; N' = 9719$

2. Calculate the relative mutability of each base x in the reporter sequence, $r_x = (s_{x*}/n_x) / (S/N)$. A value $r_x > 1$ indicates that base x is more mutable than the average, while $r_x < 1$ indicates the opposite.

3. The per-site mutation rates from all bases x, denoted u_{x*}, are assumed to be proportional to the relative mutabilities r_x. To compute the values u_{x*}, scale the relative mutabilities so that the sum $n'_T u_{T*} + n'_C u_{C*} + n'_A u_{A*} + n'_G u_{G*}$ equals $N'u$, the genomic mutation rate

of HIV (about 0.3 substitutions per replication). The correct scaling factor is $u_{x*}/r_x = N'u/\left(\sum r_x n'_x\right)$.

4. To determine the individual rates u_{xy}, partition each value u_{x*} proportional to the substitutions counted in the reported sequence; that is,

$$u_{xy} = u_{x*}\left(s_{xy}/s_{x*}\right).$$

Supplementary Table 7.1.7 gives the resulting per-site probability $\left(u_{xy}\right)$ for each nucleotide substitution in a single round of viral replication.

Mutation rates were calculated only for those amino acid substitutions which could be achieved via a single nucleotide change. All drugs studied had at least one such substitution that conferred resistance. For each possible starting codon, the rate of substitution equals the sum of all rates of nucleotide substitutions that achieve the desired amino acid change. The mutation rate u then equals the average of rates for all possible starting codons, weighted by the probability of finding that codon (based on the HIV consensus sequence base composition) (used in Supplementary Tables 7.1.4, 7.1.5).

2.5.3 Simulation algorithm

We used stochastic simulations to study the dynamics of the system described in equation (2.5) with mutation. Multiple mutations have been characterized for each drug, and to model a realistic worst-case scenario, we considered a single "synthetic" mutant defined as having the highest benefits $(\rho,$ negative $\sigma)$, lowest cost (s), and highest mutation rate of all the single-nucleotide mutants known for that drug. Each monotherapy simulation therefore tracked only two strains, wild

type y_1 and mutant y_2. Simulations modeled 48-week trials, using discrete timesteps of $\Delta t = 30$ minutes. All simulations were done in Matlab R2010b. The following steps describe the simulation for a single patient on monotherapy, with expected adherence value a:

1. Draw from the viral load setpoint distribution in Supplementary Figure 7.2.13a. This setpoint is used to determine the value of the λ parameter, assuming that the patient has 3 L plasma.

   - In the suppression phase of therapy, the initial infection size is the setpoint, rounded to the nearest integer number of cells.

   - In the maintenance phase of therapy, the initial infection size is the fully-suppressed infection size $\sum y_i \approx \left(\sum A_i \right) / d_y = 2$ c.ml^{-1}(RNA copies per ml).

2. Assign each infected cell to the mutant population (y_2) with probability u/s; otherwise the cell is in the wild-type population (y_1).

3. Identify all scheduled doses for the entire trial. All scheduled doses are evenly spaced, with the first dose occurring at the beginning of the trial. The patient takes each scheduled dose with probability a.

   - Exception: in the maintenance phase, the patient is always assumed to take the first scheduled dose.

4. Calculate the drug concentration every timestep, as described in Methods.

- In the suppression phase, the initial drug concentration is zero.

- In the maintenance phase, the initial drug concentration is C_{max}.

5. Calculate the basic reproductive ratios for the wild type and the mutant every timestep, as described in equations (2.1) and (2.2) of the main text and the Methods.

6. For each timestep:

   (a) The number of infected cells of strain i to exit the reservoir is drawn from a Poisson distribution with mean value $A_i \Delta t$.

   (b) The number of newly infected cells generated by strain i is drawn from a Poisson distribution with mean value $d_y y_i \Delta t \left[\frac{\lambda R_{oi}}{\lambda + \sum_{j=1}^{n} R_{oj} d_y y_j} \right]$.

   (c) Each cell newly infected by the wild type enters the mutant population with probability u; otherwise it remains wild type. Cells infected by the mutant do not back-mutate.

   (d) Each infected cell dies with probability $1 - \exp(-d_y \Delta t)$.

7. Determining outcome at 48 weeks:

   - In the suppression phase, the patient's status is observed at the end of the 48-week trial. If viral load is below 50 c.ml^{-1}, the trial is declared successful; otherwise virologic failure occurs.

   - In the maintenance phase, the patient's status is observed each week for 48 weeks, beginning at Week 5. If any two consecutive

observations show a viral load of at least 200 c.ml^{-1}, virologic failure occurs; otherwise the trial succeeds.

- A failed trial is considered a mutant-based failure if at least 20% of the viral population is mutant; otherwise it is considered a wild type-based failure.

8. Determining outcome over time:

- Patient's status was evaluated every 2 weeks, for 48 weeks.

- In the suppression phase, if viral load is below 50 c.ml^{-1} at the evaluation, the patient is classified as having "suppressed viral load;" otherwise the patient has "detectable viral load."

- In the maintenance phase, the patient's viral load is measured each week for 48 weeks, beginning at Week 5. If any two consecutive measurements at or before the evaluation show a viral load of at least 200 c.ml^{-1}, the patient is declared to have "detectable viral load," and is then removed from the trial, retaining this classification for all future time-points. Otherwise, the patient is declared to have "suppressed viral load."

- In the maintenance phase allowing recovery, the patient's viral load is measured as in the maintenance phase above. If viral load is at least 200 c.ml^{-1} both at the evaluation and at the immediately preceding measurement, the patient is declared to have "detectable viral load." Patients who were previously "detectable" remain in the trial and

may re-suppress.

- A measurement of "detectable viral load" is considered "via resistance" if at least 20% of the viral population is mutant; otherwise it is considered to be "via wild type."

By using a well-mixed population and by assuming that the processes of reservoir exit, replication, and death are Poisson, this method implicitly sets the effective population size of the infection equal to the census size of infected cells. Population structure, selection on linked loci, and variations in burst size among infected cells are all mechanisms that could increase variance in viral offspring number, decreasing the effective population size [3, 93]. Estimating the relevant population size to use for a model of drug resistance is difficult, as most approaches define an effective population size only for neutral loci. Simply "plugging in" a population size derived from a model without selection would be misleading in this context [167], and in lieu of a more informed value, we simply use the census size. This approach likely overestimates probabilities of mutant emergence and underestimates variability among patients [123, 238].

For dual therapy, we consider three strains: wild type, resistant to Drug 1, resistant to Drug 2. The two drugs can be simulated as two separate pills (allowing each pill to be taken or forgotten independently) or as a single combination pill (forcing the two drug concentrations to rise and fall in lockstep). In the case of two separate pills, the value of a may be different for each drug, allowing for "differential adherence" – which has been observed in some studies [99].

2.5.4 Graphing outcome versus adherence

For each monotherapy, 25,250 patients were simulated, with expected adherence a ranging from 0 to 1 (roughly equal numbers of patients were simulated for each 1% increment, including 50 patients with $a = 0$ and 50 patients with $a = 1$). The x-axis measures the *ex post* adherence for patients — that is, the actual percentage of doses taken, which may differ from the expectation a. Results were plotted for overlapping 2% windows, centered every 1% between 0 and 1, as well as for the points 0 and 1 themselves.

Analysis of dual therapy with a combination pill was similar to that of monotherapy, but with 126,250 patients (including 250 patients with $a = 0$ and 250 patients with $a = 1$).

For dual therapy with separate pills, 169,000 patients were simulated, with expected adherences a_1, a_2 ranging from 0 to 1 (roughly equal numbers of patients were simulated for each 4% $\times$ 4% increment, including 25,000 patients on the border of the distribution where at least one a_i is equal to 0 or 1.) As with monotherapy, the axes measure *ex post* adherence. Results were plotted for overlapping 4% $\times$ 4% windows, centered every 2% between 0 and 1; points plotted on the border of the distribution show patients with at least one a_i exactly equal to 0 or 1.

Note that, for maintenance therapy, the axes do not include zero, as each patient is guaranteed to take the first dose (adherence is never zero).

2.5.5 Graphing outcome over time

Analysis was performed separately for each overlapping 2% adherence window, centered every 1% between 0 and 1, as well as for the points 0 and 1 themselves. The resulting graph shows a weighted average of these results, using the adherence distribution in Supplementary Figure 7.2.13. Measurements were taken every two weeks, and the graphs show the proportion of the population with each outcome. As there is no censoring of data, the analysis is equivalent to the Kaplan-Meier method [151].

2.5.6 MSW for combination therapy

For calculations involving combination therapy (limited to two drugs in this paper), viral fitness is influenced by the dose-response curves of all drugs. DRV and RAL belong to different classes and have been shown to reduce fitness in a multiplicative (Bliss-independent) fashion, which is often expected for drugs acting on different targets [30, 144]. The equation describing viral fitness with two Bliss-independent drugs is given by:

$$R_0\left(D_1, D_2\right) = \frac{R_{00}}{\left(1 + \left(\frac{D_1}{IC_{50,1}}\right)^{m_1}\right)\left(1 + \left(\frac{D_2}{IC_{50,2}}\right)^{m_2}\right)} \tag{2.7}$$

where D_1, D_2 are the concentrations of each drug in the relevant compartment, $IC_{50,1}, IC_{50,2}$ are the concentrations at which 50% inhibition occurs, and m_1, m_2 are the slope parameters. The numerator R_{00} is the baseline basic reproductive ratio in the absence of drug treatment. Mutations that confer resistance to a given

drug change the IC_{50}, slope and drug-free fitness similarly to the way described in equation (2.2) (main text).

For a two-drug combination where we assume that a viral strain may only be resistant to a single drug, there are now eight potential selection windows. Drug levels may be high enough for guaranteed treatment success; in the MSW for one or both drugs; in the overlapping region for one or both of the MSWs and the WGW, or strictly in the WGW. Supplementary Figure 7.2.11 shows the possible windows for the RAL+DRV/r combination.

2.5.7 DERIVATION OF FIG. 2.2.2F: COMPARING RISK OF WILD TYPE-BASED AND MUTANT-BASED VF FROM SELECTION WINDOW DATA

Fig 2.2.2f ranks drugs by the relative risk of mutant versus wild-type failure, regardless of the total risk of failure, based on the time spent in each selection window. The ranks are plotted along a line with values ranging from -1 (DRV/r and d4T, highest relative risk of wild-type failure) to 1 (FTC, the highest relative risk of mutant failure). This plot was constructed based on the data in Fig. 2.2.2a.

To devise this scale, we let

$$x = \text{time until entry into MSW (days) / time until entry into WGW (days)}$$

$$= \text{length of green bar / length of green + dark red bars,}$$

$$y = \text{time spent in MSW (days)}$$

$$= \text{length of both red bars.}$$

$$(2.8)$$

If the drug immediately enters the WGW at day 0, or if it skips the MSW completely, then x is defined to be 1.

Then the scale value to be plotted, a, is calculated as

$$a = \frac{y}{y_{max}} - x, \qquad (2.9)$$

where $y_{max} \approx 16.5$ days, the maximum time that a drug spends in the MSW (obtained for FTC). Since both x and y/y_{max} range between 0 and 1, the scale ranges between -1 (failure via wild type only) and 1 (largest relative risk of resistance).

In this formula, x is a proxy for the rapidity of wild type-caused virologic failure ("wild-type risk") relative to mutant-caused virologic failure ("mutant risk"). When x is small, the MSW window is reached long before the WGW, meaning that "mutant risk" is high and "wild-type risk" is low. When x is high, the

54

WGW is reached soon after the MSW, or without ever entering the MSW, and so "wild-type risk" is high and "mutant risk" is low. While x considers how quickly the infection can start to grow, it does not consider the length of time in the MSW. Even if the MSW begins as soon as a dose is taken (so that $x = 0$), one still needs to consider for how long the mutant strain is selected over the wild-type to determine whether mutant-based or wild type-based virologic failure is more likely to occur. Figure 7.2.12 shows a scatter plot of y versus x.

2.6 Acknowledgements

2.7 Manuscript information

2.7.1 Previously published as

This manuscript appeared in [261]:

Rosenbloom DIS, Hill AL, Rabi SA, Siliciano RF, Nowak MA (2012). Antiretroviral dynamics determines HIV evolution and predicts therapy outcome. *Nature Medicine*, 18(9):1378–1385.

3

Predicting outcomes of treatments to eradicate the HIV latent reservoir

3.1   Introduction

THE LATENT RESERVOIR (LR) FOR HIV-1 IS A POPULATION of long-lived resting memory CD4$^+$ T cells with HIV-1 DNA integrated into their genomes [50, 51]. After the reservoir is established during acute infection [53], it increases to $10^5 - 10^7$ cells and then remains stable. As only actively replicating virus is targeted by current anti-HIV drugs, latently infected cells persist even after years

of effective treatment [52, 92, 284, 323]. Cellular activation leads to virus production and, if treatment is interrupted, a rapid increase in viremia is observed within weeks of discontinuation of therapy. This dynamic prevents cure of HIV-1 by HAART alone and thus necessitates lifelong adherence to HAART. Drugs that specifically target LR cells for rapid activation may reduce the size of the LR [49, 81]. Collectively called *latency reversing agents* (LRA), this class includes the histone deacetylase inhibitors such as vorinostat [8, 10] and valproic acid [9, 285], the alcoholism drug disulfiram [327], protein kinase C activators prostratin [27] and bryostatin [197, 249], and quinoline derivatives [328]. The hope is that following treatment with these drugs, patients might be able to discontinue HAART with minimal risk of viral rebound. While the discovery of LRAs is now the subject of an intense research effort, it is unclear how much the LR must be reduced to enable patients to discontinue HAART safely. Mathematical models of treatment dynamics are urgently needed to inform clinical trial design, interpret outcomes, and guide further drug discovery.

Mathematical models have been instrumental in understanding the dynamics of HIV-1 infection, including the LR. Multi-compartment models of HIV-1 infection have successfully been used to describe the phases of viral decay after initiation of HAART [161, 215, 240], the role of ongoing replication in slowing the rate of LR decay [272], and the appearance of viral "blips" during treatment [59, 257, 258]. Recent studies have considered the role of the LR in the development of drug resistance to HAART[261]. However, no model has been developed to study the effect of treatment with LRAs. Here we build and

analyze such a model to predict how *in vitro* drug efficacy translates to patient outcomes, measured as the length of time following discontinuation of HAART before viral rebound occurs.

3.2 MODEL AND RESULTS

We consider proposed therapy protocols for latency reversing agents (LRAs) that administer the treatment while a patient is on fully suppressive HAART [81]. It is believed that LRA therapy will reactivate transcription at the HIV-1 LTR, leading to production of cytotoxic viral products and cell death. While plasma HIV-1 levels may increase during this phase, strict adherence to HAART is expected to prevent new cellular infections. After a period of treatment with both LRAs and HAART, both therapies would be interrupted. Here we provide a model of the viral infection immediately after this therapy ends (Fig. 3.1.1(A)). Our goal is to estimate the probability of cure, or, barring that best-case outcome, to estimate the distribution of times until virologic rebound, in terms of parameters describing the underlying infection dynamic and the LRA therapy.

The model tracks two cell types: productively infected active $CD4^+$ T cells, and latently infected resting $CD4^+$ T cells. Cells carrying nonviable HIV-1 provirus (which may vastly outnumber cells carrying replication-competent proviruses) are excluded from these two quantities. The level of plasma virus is not tracked explicitly, but is assumed to be proportional to the productively infected cells. Four types of events can occur in this model, which is described

Figure 3.1.1: Schematic of LRA therapy and stochastic model of rebound following therapy. A) Proposed treatment protocol, illustrating possible viral load and size of latent reservoir before and after LRA therapy. When HAART is started, viral load decreases rapidly and may fall below the limit of detection. The latent reservoir is established early in infection (not shown) and decays very slowly over time. When LRA is adminstered (either continuously, as shown, or in intervals), the latent reservoir declines. Depending on the efficacy of LRA therapy, the infection may be cleared, or viremia may eventually rebound. B) LRA efficacy is defined by the parameter Q, the number of $\log_{10}$-reductions in LR size. C) Stochastic model of viral dynamics following LRA therapy, tracking both latently (rectangles) and actively (ovals) infected cells. Each arrow represents a type of event that occurs in the model and its rate, described in the text. D) The expected number of "offspring" for each productively infected cell is the basic reproductive ratio $R_0 = 2$. The "infectivity variance" parameter λ determines variance of the offspring distribution. The offspring distribution conditional on event b occurring is given by a Poisson: $P(c) = \lambda^c e^{-\lambda}/c!$

formally as a two-type branching process (Fig. 3.1.1(C)). A latently infected cell can either activate at rate a, or die at rate d_z. An actively infected cell can either produce a burst of virions at rate b, resulting in the infection of c other cells, after which it dies, or it can die without producing virions (at rate d). The number of cells infected after each burst event is a Poisson random variable with mean and variance λ (Fig. 3.1.1(D)), see Sec 3.4.3.

The total death rate of productively infected cells ($d_y = b + d$) has been well characterized to be 1 day^{-1} from treatment initiation studies. The rate of reactivation of cells from the latent reservoir can be estimated based on the size and composition of the LR (Fig. 3.2.1) and the level of residual plasma virus for patients on fully suppressive HAART ($\approx$ 2 copies HIV-1 RNA per milliliter plasma, c ml^{-1}). We estimate a to be 7×10^{-5} day^{-1} and consider values in the range 10^{-5} to 3×10^{-3}. The death rate of latently infected cells is estimated from studies of the rate of decay of the LR to be $d_z = 4.6 \times 10^{-4}$ day^{-1} (corresponding to a 44 month half-life). We also present results for two extremes: a half-life of only 6 months, and, $d_z = 0$. The basic reproductive ratio for this model, defined as the expected number of new infected cells that a single actively infected cell produces, is $R_0 = \lambda b/(b + d)$. The average R_0 value is estimated from time-to-rebound in HAART-interruption studies to be $R_0 = 2$ (which is lower than the values estimated for acute infection). R_0 does not uniquely determine the dynamics of the stochastic model because b and λ cannot be simultaneously identified. Holding R_0 constant, the parameter λ controls the strength of random drift in the infection: for high λ, reproduction resembles a

"jackpot" event where a few infected cells give rise to many new infection events, while many other infected cells die before infecting additional cells. Here we vary λ in the range 2 – 50 (see Sec. 3.4.2).

The initial conditions depend upon the number of latently infected cells that survive LRA therapy. This quantity is defined by the latent reservoir size prior to LRA therapy and the log-efficacy of LRA therapy (Fig.3.1.1(B)). The model aims to determine whether or not this population of cells will escape drift and restart the infection before all the cells die. These results are independent of the details of how the drug is administered or the mechanism of action. As the model does not provide for any limitation of growth as the infection becomes large, it is valid only for the initial stages of viral rebound. Since clinical viral rebound thresholds (viral load above 50 – 200 c ml^{-1}) are well below carrying capacity (typical setpoint viral load of 10^4 – 10^6 c ml^{-1}), this model suffices to analyze the probability and timing of rebound following LRA therapy and HAART interruption.

We used several experimental findings to estimate the size range of the LR, which we define as the number of resting CD4$^+$ T cells with integrated HIV that are capable of producing infectious virus upon reactivation (Fig. 3.2.1). We considered three cases for the LR size distribution among patients. Limiting dilution co-culture assays [283] are currently the gold standard for LR size measurement. In case (i), we assume all patients have a reservoir size equal to the average measured in these assays ($\approx 1 \times 10^{-6}$ cells), and in case (ii) we

Figure 3.2.1 *(following page)*: Experimental scheme for classifying rest-ing CD4$^+$ T cells based on HIV-1 infection and viral production, using data from [89, 125]. Scheme starts at the top, with purified resting CD4$^+$ T cells, and proceeds downward through the experimental analyses listed. The f vari-ables represent fractions of this resting CD4$^+$ T cell pool with the charac-teristics listed. "PCR": Digital droplet PCR identified cells containing HIV-1 DNA, nearly all of which is expected to be integrated. "Co-culture": PHA was used to induce viral replication in latently infected cells. "Seq. defect": Non-induced cells were analyzed for genetic defects preventing production of replication-competent virus. A fraction i of these non-induced cells had no ob-servable defects (all open reading frames intact); this fraction constitutes f_{NII} of all resting CD4$^+$ T cells. Question marks indicate that it was not possible to determine by this analysis what fraction of cells would produce replication-competent virus *in vivo*, due to integration site effects and undetectable se-quence defects. Even defective provirus may be able to produce defective virons that contribute to residual viremia (gray arrow). The latent reservoir (shaded box) consists of induced and replication-capable non-induced cells. Values shown are averages and ranges of $\approx$ 30 patients.
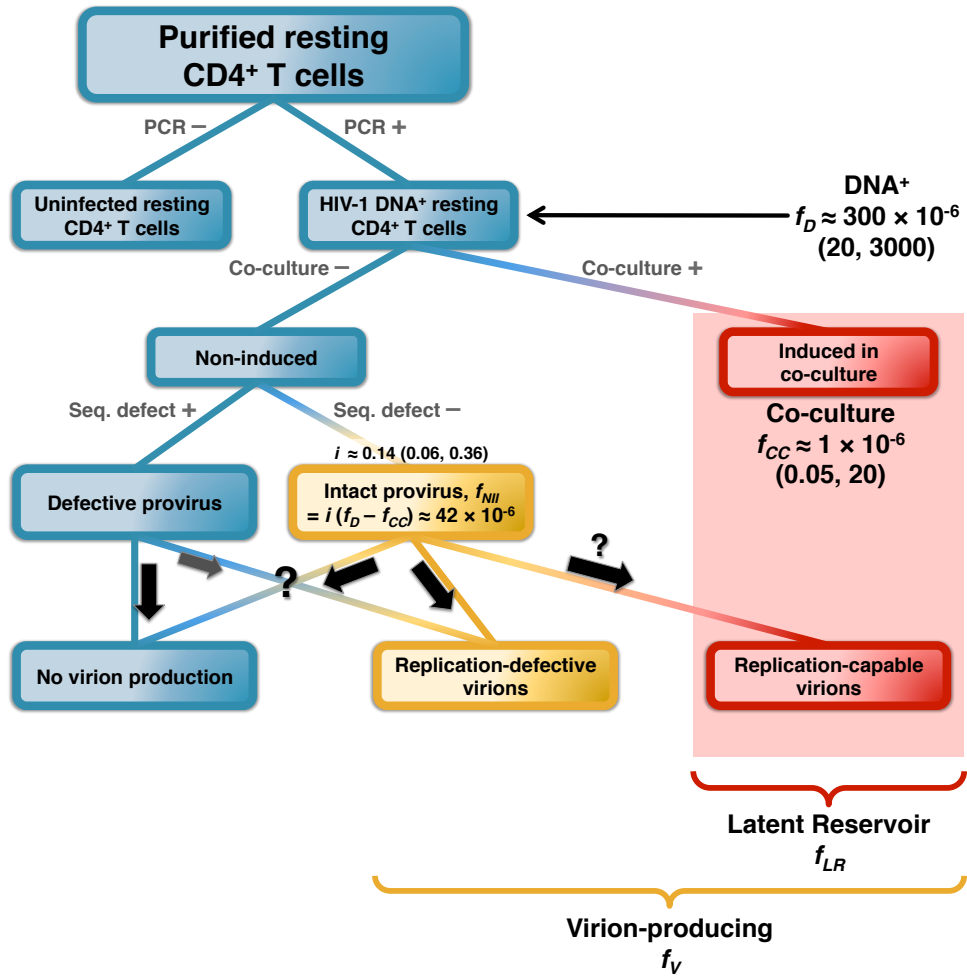
Figure 3.2.1 *(continued)*

incorporate inter-patient variability in assay results [89]. PCR measurements detect cells with HIV-1 DNA at a higher frequency than co-culture assays, but much of this virus may be defective. Full genome clonal sequencing of provirus from cells not induced in co-culture identifies a portion with all open reading frames intact [125]. In case (iii), we include these cells in the LR. For all cases we assume a total resting $CD4^+$ cell count of 10^{12}. The resulting distributions are shown in Fig. 3.2.2(A), see Sec. 3.4.1.

The best-case outcome of LRA therapy, barring complete eradication of the reservoir, is that none of the surviving latently infected cells activate and lead to a resurgent infection. In this case, we say that LRA has cleared the infection. We used the model to predict the relationship between LRA log-efficacy (denoted Q) and clearance probability (Sec. 3.4.3). Fig. 3.2.2(B) shows results for the three possible reservoir distributions (i) – (iii) described above.In cases (i) and (ii), where the average pretreatment reservoir size is 10^6, the reservoir must be reduced by three to five orders of magnitude before half of patients clear the infection. Including inter-patient variability only causes the clearance probability to increase more gradually with Q. If co-culture does not detect all cells in the latent reservoir (case iii), then Q of four to six is required for 50% clearance. In all three cases, the clearance probability decreases with reservoir half-life. Clearance probability also increases with infectivity variance λ, as this parameter controls the likelihood of viral lineage extinction by drift (Fig. 7.3.1B).

If LRA therapy fails to clear the infection, the next-best outcome is substantial extension of the time until virologic rebound, defined as a viral load of 200 c ml^{-1}

(Sec. 3.4.2). We computed the relationship between efficacy of LRA therapy and the median time until rebound, among the patients who do not clear the infection (Fig. 3.2.2(C)). For an LR size of 10^6 (case i), only modest increases in median rebound time are predicted for up to 100-fold reductions in the size of the reservoir ($Q \leq 2$). In this range, the rebound time is independent of latent cell lifespan, and it is driven mainly by the reactivation rate and the viral reproductive ratio. The curve inflects upward at $Q = 2$ (on a log scale) and reaches a ceiling as clearance of the infection becomes the dominant outcome (Fig. 3.2.2(C)(i)). If cells in the reservoir are extremely long-lived, it is possible for rebound to occur even after decades of apparent cure. If the LR size is larger (case iii), then the median rebound time curve is shifted rightwards, requiring higher LRA efficacy for the same outcomes (Fig. 3.2.2(C)(iii)). In all three cases, the inflection point decreases in λ. In case (i), this point varies between $Q = 1.5$ (for $\lambda = 50$) and $Q = 3$ (for $\lambda = 2$). Accordingly, the median rebound time increases in λ (Fig. 7.3.1C).

The upward inflection observed in median rebound time results from an important change in the forces governing viral dynamics. If the reservoir is large enough (low Q), a surplus of cell activation occurs such that the dominant component of rebound time is the time that it takes for virus from one of the many activated cells to grow exponentially to rebound levels; the system is in a *growth-limited regime*. If the reservoir is small (high Q), the expected waiting time until activation of the first cell fated to establish a rebounding lineage is the

Figure 3.2.2 *(following page)*: Clearance probabilities and rebound times following LRA therapy predicted from model. A) Three cases for the population-level distribution of LR size (Sec. 3.4.1). Case i) All patients have the same latent reservoir size, $M_{LR} = 10^6$, estimated from the geometric mean number of cells that are capable of producing infection in laboratory co-culture assays. Case ii) Latent reservoir size is distributed according to variation observed in co-culture assays, with geometric mean 10^6. Case iii) The latent reservoir includes many cells that fail to be detected in co-culture but have intact viral genomes. B) Probability that the reservoir is cleared by LRA. Clearance occurs if all cells in the reservoir die before a reactivating lineage leads to viral rebound. C) Median viral rebound times, among patients who do not clear the infection. D) Survival curves for patients following LRA therapy. The percentage of patients who have not yet experienced viral rebound is plotted as a function of the time after interruption of LRA therapy and HAART. Curve color indicates the efficacy of LRA in reducing the size of the LR ($Q = 0$ to 6, see legend). Results are shown for a half-life of 44 months; other half-lives are shown in Fig. 7.3.3 Solid lines represent simulations, and open circles represent approximations from a branching process calculation (Sec. 3.4.3). All simulations included $10^4 - 10^5$ patients and used parameters $a = 7 \times 10^{-7}$, $\lambda = 20$ and $R_0 = 2$.
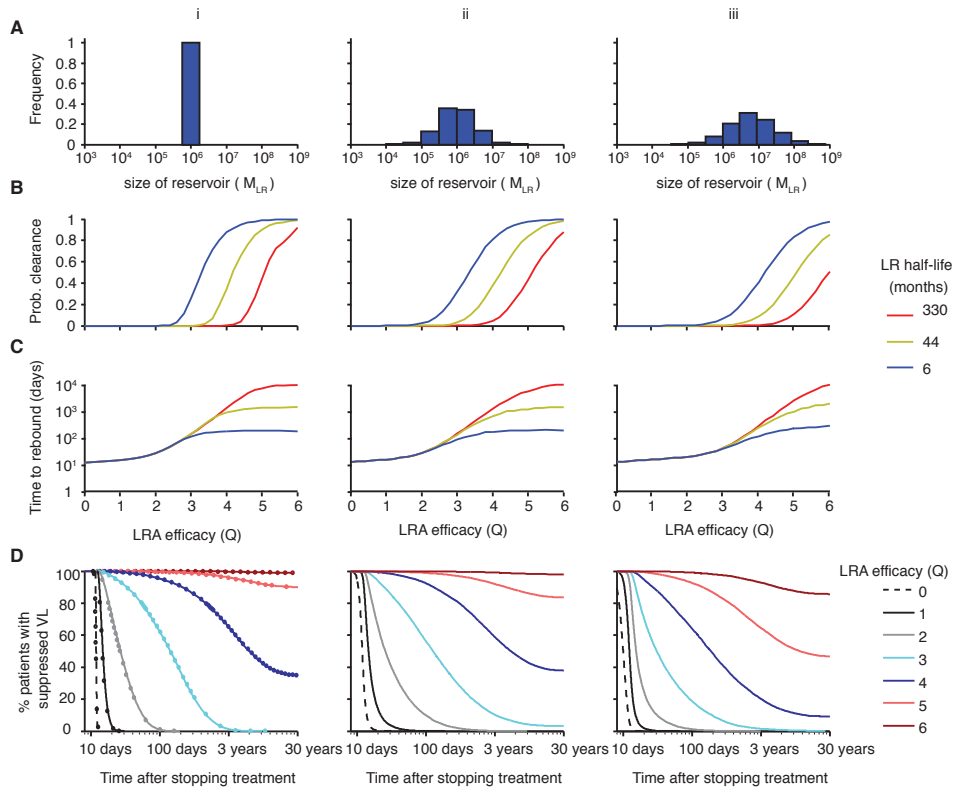
Figure 3.2.2 *(continued)*

68

dominant component of rebound time; the system is in an *activation-limited regime*. Since this waiting time is roughly exponentially distributed, times to rebound in this regime can vary widely among patients. The "threshold Q" defining the boundary of these two regimes matches the value where upward inflections are observed (Fig. 7.3.2(A)).

The scope of predicted interpatient variability in outcomes can be seen in survival curves, plotting the fraction of simulated patients maintaining virologic suppression over time (Fig. 3.2.2D). For small reductions in LR size ($Q \leq 2$) patients uniformly rebound within a few months, since rebound dynamics are not in the activation-limited regime (Fig. 3.2.2(D)(i)). If LRA therapy manages to decrease the reservoir size 1,000-fold ($Q = 3$), then about 20% of patients remain rebound-free for at least a year. Higher Q leads to clearance in many patients, though rebound can still occur after a decade without viremia. Interpatient variation in LR size makes it more likely to observe long periods without rebound (Fig. 3.2.2(A)(ii)). Rebound delays of over three years are achieved for 10% of simulated patients at $Q = 3$, versus only 1% in case (i). The fortunate few who completely clear the infection started with an LR smaller than the average size of 10^6 prior to treatment. Survival curves decline more rapidly if the average LR size is larger (case iii) or if λ is lower (Fig. 7.3.4), indiciating a less beneficial outcome for LRA therapy.

For all three reservoir size distributions considered, rebound may occur even after long periods of virologic suppression. Taking case (ii) for example, among the patients without rebound at six months post-therapy, 71%, 38%, or 11% suffer

Figure 3.2.3: Efficacies required for successful LRA therapy. Target LRA efficacy values based on three different treatment goals are plotted versus reservoir half-life $\left(\frac{\ln(2)}{a+d_z}\right)$ and the reservoir reactivation rate (a). A) The target efficacy Q at which at least 50% of patients still have suppressed viral load one month after treatment interruption. B) The target efficacy Q at which at least 50% of patients still have suppressed viral load one year after treatment interruption. C) The target efficacy Q at which at least 50% of patients have eradicated the reservoir without experiencing viral rebound. Because some patients may go for a year without rebound but then rebound later, the target Q for one year off therapy is always less than that for a lifetime off therapy. Results were calculated from the branching process description of the stochastic process, which agrees with simulation (Sec. 3.4.3). All calculations used $\lambda = 20$ and varied R_0 to ensure baseline rebound time was constant (12 days). Worst-case: $a = 3 \times 10^{-3}, R_0 = 1.43$, high: $a = 4.6 \times 10^{-4}, R_0 = 1.7$, estimated: $a = 7 \times 10^{-5}, R_0 = 2$, best-case: $a = 10^{-5}, R_0 = 2.21$.

rebound in the following six months, for efficacies $Q = 2$, $Q = 3$, or $Q = 4$, respectively (Fig. 3.2.2(A)(ii)).

To set goals for treatment efficacy, we calculate three "target values" of Q, assuming a pretreatment reservoir size of 10^6. One month is near the upper limit of rebound times without LRA therapy, and we first calculate define the log-efficacy for which 50% of patients exceed this limit (Fig. 3.2.3(A)). This value is insensitive to the estimate for reservoir half-life and increases logarithmically with the activation rate. For a broad range of activation rates, a

1.5- to 3-log reduction is needed for a one-month delay. We next calculate the target efficacy for one-year delays (Fig. 3.2.3(B)). Using estimated parameter values, this goal requires a 3-log reduction in reservoir size. This value is only mildly sensitive to reservoir half-life (declining only for very short half-lives, at which reservoir clearance is likely) and also scales logarithmically with activation rate. If activation exceeds the estimated rate, then the target Q is closer to 4.5 – a 30,000-fold reduction. Finally, since the ultimate goal of LRA therapy is to clear the reservoir completely, we determine the log-efficacy for which at least half of patients clear the infection (Fig. 3.2.3(C)). This value consistently exceeds the more modest one-year target and scales logarithmically with the product of activation rate and half-life. Parameter scaling relationships follow from a generating function analysis of the branching process (Sec. 3.4.3). All target Q values also scale with the reservoir: a 1-log increase in reservoir size would necessitate a unit increase in Q.

We evaluated the robustness of our conclusions to simultaneous changes in latent cell activation and death rates, pretreatment reservoir size distributions, and infectivity variance. For a worst-case analysis, the latent cell death rate was set to zero (such that the reservoir decays only via activation), the pretreatment reservoir size distribution was set to that of case (iii), and infectivity variance was set to a low value. The resulting target Q for a one-month delay increased by 2, and the clearance target Q increased by 3 versus baseline (Fig. 7.3.5(i) versus (iii)). For a best-case analysis, the latent cell activation rate was set to 7 times lower than baseline and the death rate was increased to yield a reservoir half-life

of only six months, the pretreatment reservoir size distribution was set to that of case (i), and infectivity variance was set to a high value. The results were similar to that of the baseline case, but with Q shifted upwards by 1, and with a lower ceiling rebound time (200 days versus over 1,000 days) (Fig. 7.3.5(i) versus (iii)).

Throughout our analysis, we have characterized LRA drug efficacy by the log-reduction in reservoir size following therapy, which may not be observable. Laboratory studies of cellular models of latency may estimate Q, but their relevance *in vivo* remains unknown. It may not be possible to measure Q values above 1 by quantifying reservoir size following LRA therapy, as current co-culture assays cannot detect reservoirs smaller than about 10^5 cells [89]. Since current approaches to LRA therapy seek to reduce reservoir size by inducing activation of latently infected cells, Q may be estimated by measuring the dynamics of viral load during simultaneous HAART/LRA therapy (Sec. 3.4.6). Since the effect of LRA therapy on resting $CD4^+$ T cell phenotype is not fully characterized [275, 276], there is considerable uncertainty in this relationship; nonetheless a sharp, transient peak viral load of at least several hundred c ml^{-1} is expected for highly effective therapy ($Q \geq 2$) (Table 3.4.1).

3.3   Discussion

Our model is the first to describe the action of investigational latency reversing agents and set quantitative goals for LRA therapy, offering guidance for the design and testing of treatment protocols. There is currently little understanding of the degree of reservoir activation required to provide meaningful benefit. We

analyzed experimental findings regarding reservoir size and composition to describe three plausible settings for LRA therapy (Fig. 3.2.2A). In each setting, our model translates a reactivation measure of therapy efficacy (parameter Q), which may be estimated *in vitro*, to a prediction of clinical benefits.

For a wide range of parameters, we find that LRAs must reduce the reservoir by at least 1.5–3 orders of magnitude to see a meaningful increase in the time to virologic rebound after HAART interruption (upward inflection in Fig. 3.2.2C and Fig. 7.3.1C), and that 3–5 orders of magnitude are needed for half of patients to clear the infection (Fig. 3.2.3C and Fig. 7.3.2). Standard deviations in rebound times of many months are expected to be the norm for successful therapy, owing to variation in pretreatment reservoir size and roughly exponentially-distributed reactivation times after effective LRA therapy brings the infection to an activation-limited regime. While the required LRA efficacy for these beneficial outcomes is almost certainly beyond the reach of current drugs, our results do permit some optimism: reactivation of all cells in the reservoir is unlikely to be necessary for complete reservoir clearance and safe cessation of HAART. This is due to the reasonably high probability that a cell in the LR will either die before reactivating or, following activation, fail to produce a chain of infection events leading to rebound. On a more cautionary note, the wide distribution in reactivation times implies that continual monitoring of patients is essential, as rebound is possible even after long periods of viral suppression.

Clinical and laboratory findings constrain the basic reproductive ratio R_o and active cell death rate d_y in a relatively narrow range. The rebound population size,

while it may vary by over order of magnitude, has little effect on model outcomes, especially in the activation-limited regime. Remaining parameters — pretreatment reservoir size, latent cell half-life $\frac{\ln(2)}{a+d_z}$, latent cell activation rate a, and infectivity variance λ — have profound impacts on treatment outcome, and they are not well-established to within an order of magnitude. While LR size is generally estimated at 10^6, recent studies have shown that even co-culture assays — recognized as the gold standard for latency measurement — may drastically underestimate reservoir size (Fig. 3.2.2(A)(ii) versus 3.2.2(A)(iii)). We accounted for these studies to show that this underestimate may cause expectations of LRA therapy outcomes to be unduly optimistic (Figs. 3.2.2(B,C)(ii) and 3.2.2(ii) versus Figs. 3.2.2(B,C)(iii) and 3.2.2(iii)). Considering both variation in pretreatment reservoir size and latent cell half-lives, the log-reduction needed to delay rebound for one year in half of patients is ≈ 3 to ≈ 4.3 — a reduction of 1,000- to 20,000-fold (Fig. 7.3.3(C)(ii) vs. (A)(iii)). When broad variation in a and λ is also considered, the range expands to ≈ 2 to ≈ 5.2 (Fig. 7.3.5(C)(ii) vs. (C)(iii)).

Our analysis characterizing the required efficacy of LRA therapy does not rely on the specific mechanism of action of these drugs, only the amount by which they reduce the reservoir. We have assumed that the reservoir is a homogeneous population without variation in activation and death rates. The presence of reservoir compartments with different drug penetrations does not alter our results, as they are stated in terms of total reservoir reduction. If, however, these compartments vary in activation or death rates, or if viral dynamics of activated

cells depends on their source compartment, then our model may need to be modified. In the absence of clear consensus on compartments constituting the LR, we have considered the simplest scenario with fewest assumptions, which may be able to fit future LRA therapy outcomes.

To date, laboratory and clinical studies of investigational LRAs have generally found weak potential for reservoir reduction, with $Q < 1$ [9, 54, 327]. We predict that much higher efficacy will be required for any hope of eradication. While we have focused on the role of LRA therapy in reducing the reservoir size, our main findings may also serve to interpret viral eradication or delays in rebound caused by early treatment initiation [29, 243, 296] or stem cell transplantation [121, 139], both of which also reduce the LR size.

3.4 SUPPLEMENTARY METHODS

3.4.1 CHARACTERIZING THE LATENT RESERVOIR

This section details the calculations summarized in Fig. 3.2.1.

COMPONENTS OF THE LATENT RESERVOIR

Let M be the total number of resting $CD4^+$ T cells in an individual. The fraction of these cells containing HIV-1 DNA, written f_D, is measured using PCR on purified resting $CD4^+$ T cells [89]. The total number of resting $CD4^+$ T cells with HIV-1 DNA is $M_D = f_D M$.

Many of these cells may not be capable of producing infectious virions upon reactivation. Limiting-dilution co-culture assays [283] measure the fraction f_{CC} of resting $CD4^+$ T cells that produce replication-competent virus following cellular activation. The total number of these cells is $M_{CC} = f_{CC} M$.

Experiments show that f_D exceeds f_{CC} by more than two orders of magnitude, but the reasons for this discrepancy are unclear. Some portion of cells with HIV-1 DNA may harbor defective provirus that cannot produce viral proteins. Others may contain intact provirus, but evade co-culture detection for other reasons. These two scenarios can be distinguished by full genome clonal sequencing of integrated provirus, from which defects preventing viral replication may be identified [125]. Among those cells infected with provirus but undetected in co-culture, let i be the fraction that nonetheless are shown by sequencing to have all open reading frames intact. Then f_{NII}, defined as the frequency of

non-induced, intact provirus-containing cells, is $f_{NII} = i \left(f_D - f_{CC} \right)$.

The quantity of interest for our model is the number of resting CD4$^+$ T cells harboring proviral DNA capable of causing infection of other cells. This latent reservoir contains M_{LR} cells, or a fraction f_{LR} of M. At minimum, it includes those cells tested positive by co-culture, and at maximum, it includes all cells containing intact proviruses. It is likely to be less than the upper bound, since even cells with intact provirus may be incapable of re-starting infection due to integration into a transcriptionally silent site, or due to other defects that escaped detection.

Size of the latent resevoir and interpatient variability

The size and composition of the latent reservoir can vary significantly between patients. This variation is relevant to variation in rebound time. Based on results in [89, 125], we can calculate the average parameters and estimate the population level distributions:

Quantity	Average	Distribution
f_D	300×10^{-6}	$LogNormal(-3.5, 0.5)$
f_{CC}	10^{-6}	$LogNormal(-6, 0.5)$
i	0.14	$LogNormal(-0.85, 0.2)$

The distributions were estimated by assuming the ranges observed in these study of approximately 30 patients represent the center 95% of the distribution (ie 2 standard deviations). We examinine three different cases for inter-patient variability:

Case 1: All patients have the same LR size, which equals the average value measured in co-culture assays ($f_{LR} = f_{CC}$).

Case 2: The fraction f_{CC} varies among patients as in the table above, and $f_{LR} = f_{CC}$ for each patient.

Case 3: The fractions f_D, f_{CC}, and i are sampled from the above distributions, subject to the constraint $f_{CC} \leq f_D$. Then f_{NII} is computed. The value $\log(f_{LR})$ is then sampled uniformly from the interval $\left[\log \left(f_{CC} \right), \log \left(f_{NII} \right) \right]$.

Inter-patient variability in outcomes is lowest for Case 1 and highest for Case 3. Throughout this paper, we assume the only parameters that vary between patients are those related the composition of the reservoir.

ESTIMATING THE NUMBER OF VIRION-PRODUCING CELLS

The value f_V is defined as the fraction of resting memory CD4$^+$ T cells that are capable of producing virions that are detected in viral load assays, regardless of whether these virions are infectious. This value is used to calibrate the reactivation rate a from observed residual viral load values (Eq. (3.2), below). While $f_{LR} \leq f_V \leq f_D$, we cannot establish the relationship between f_V and f_{NII}, as some cells in f_{NII} may harbor transcriptionally silenced provirus, and conversely some cells in f_V may harbor provirus with detectable defects. As an intermediate estimate we assume $f_V = f_{CC} + f_{NII} = 43 \times 10^{-6}$, though we also consider the two extreme values in the best and worst case scenarios discussed in the main text.

3.4.2 Estimating viral dynamics parameters

Death rate of actively infected cells d_y

The parameter d_y has been measured in many treatment initiation studies to be 1 day^{-1} [188]. In the stochastic model we constrain $b + d = d_y = 1\mathrm{day}^{-1}$.

Activation rate of latently infected cells a

Parameter a is estimated from observed viremia during fully supressive HAART treatment $(R_0 = 0)$. Among patients, residual viral load is highly correlated to LR size [11]. Since reservoir decay is slow compared to the dynamics of actively infected cells, the residual viral load reaches a quasi-steady-state relative to the size of the latent reservoir (Section 3.4.3). At this level, the number of actively infected cells is

$$y_0 \approx \frac{az_0}{d_y} \tag{3.1}$$

Here z_0 is the number of cells in the LR capable of causing productive infection at the time HAART is interrupted, which is equal to $f_{LR}M$, as described in Section 3.4.1, above.

The value y_0 may be estimated from residual viral load measurements. Let σ be the proportionality constant such that v_0 (residual number of *infectious* virions during HAART) equals σy_0. Since viral load measurements do not distinguish between infectious and noninfectious particles, the only quantity that may be ascertained is the total number of virions w_0, which exceeds v_0. Assuming that

infectiousness of the released virion affects neither viral burst rate nor decay rate, the same proportionality constant σ applies to the number of cells capable of producing (infectious or noninfectious) virions, written $\hat{y}_o$, resulting in the relationship $w_o = \sigma \hat{y}_o$. Since $\hat{y}_o \approx a f_V M / d_y$, the value of a equals:

$$a = \frac{d_y w_o}{\sigma f_V M} \tag{3.2}$$

A reasonable estimate for the proportionality constant is $\sigma = 1$, as discussed in Suppl. Materials of [261]. We use $w_o = 2$ HIV-1 RNA copies per milliliter plasma, c ml^{-1}, corresponding to 3,000 plasma virions for a 70 kg person with 3L plasma, and we use $M = 10^{12}$. Based on the observed averages $f_D = 3 \times 10^{-4}$, $f_{CC} = 10^{-6}$, and $i = 0.14$, we estimate $f_V = 4.3 \times 10^{-5}$ and $a = 7 \times 10^{-5}$ day^{-1}. This activation rate is below previous estimates [254, 272], which would predict a higher residual viral load than observed for patients on fully suppressive HAART. However, because the size and composition of the LR are still a matter of debate, we consider a range of values. For a worst-case scenario of high activation rate, we suppose that only cells testing positive in co-culture contribute to residual viremia ($f_V = f_{CC} = 10^{-6}$ and $a = 3 \times 10^{-3}$). For a best-case scenario of low activation rate, we suppose that residual viremia is seeded from all cells harboring HIV-1 DNA ($f_V = f_D = 3 \times 10^{-4}$ and $a = 10^{-5}$). Results for these scenarios are shown in Figs. 3.2.3 and 7.3.5.

Viral rebound threshold n

A viral rebound threshold of 200 c ml^{-1} corresponds to 3×10^5 plasma virions (for a 70 kg person with 3 L plasma). Using the estimate $\sigma \approx 1$ above, the number of actively infected cells at rebound is $n \approx 3 \times 10^5$. Model results are not sensitive to this value, as rebound probability depends on the logarithm of n (Section 3.4.3, below).

Death rate of latently infected cells d_z

Resting memory CD4$^+$ T cells die at a rate d_z, which may be estimated from studies measuring the total decay rate of the reservoir, $a + d_z$. Given a mean half-life of 44 months [284], we estimate d_z to be

$$
\begin{aligned}
d_z &= \frac{ln(2)}{1320 \text{ days}} - a \\
&\approx 4.6 \times 10^{-4} \text{day}^{-1}
\end{aligned}
\tag{3.3}
$$

This parameter is varied to explore a range of half-lives in Figs. 3.2.2, 3.2.3, 7.3.3.

Viral fitness R_0

The combined viral fitness parameter R_0 can be estimated from rebound times measured in HAART-interruption studies [63, 263]. The average rebound time in these studies is twelve days, which consists of both the time needed for drug

levels to decay such that the infection can grow and the subsequent time needed for exponential viral growth to rebound levels computed in Equation (3.23). For estimating the drug decay time, we used the method of [144, 279] with a typical regimen of AZT, 3TC, and SQV. For Equation (3.23), we used a rebound factor (described in Section 3.4.5) of $r = 4{,}300$, based on residual viral load $w_0 = 2$ c ml^{-1}, *infectious* residual viral load $v_0 = w_0 \frac{f_{CC}}{f_V} = 0.0465$ c mL^{-1} (using average values of f_{CC} and f_V in Section 3.4.1, above), and detection at 200 c ml^{-1}. We compute a value of $R_0 = 2$, at which it takes about four days for drug levels to decay and another eight for viral growth to rebound levels.

INFECTIVITY VARIANCE λ

Based on the rate at which patients fail therapy due to drug resistance, a previous study estimated the rate at which cells that are fated to establish a lineage activate from the latent reservoir to be about 4 per day, in the absence of treatment [238]. This estimate is highly uncertain, as it is sensitive to measured mutation rates and fitness costs of resistance mutations.

Using our baseline values of $a = 7 \times 10^{-5}$ and $M_{LR} = 10^6$, the number of cells activating per day is 70. An extinction probability of $h_1 = 1 - 4/70 = 0.94$ (defined in Equation (3.8)) would make our baseline values consistent with the above estimate. Using $R_0 = 2$, the implied λ to obtain this lineage extinction probability is 28. To account for uncertainty, we consider λ between 2 and 50.

R_o and d_y do not uniquely determine the dynamics of the stochastic model because b and λ cannot be simultaneously identified. After choosing a value for λ between 2 and 50, the parameter b is calculated using the relationship $R_o = b\lambda/d_y$. The parameter d is then obtained from the relationship $b + d = d_y$.

3.4.3 Stochastic model of viral dynamics

Stochastic process

The stochastic model of viral dynamics described in the text can be formally represented as the reactions below:

$$
\begin{aligned}
Z &\to Y\text{... rate constant: } a \\
Z &\to o\text{... rate constant: } d_z \\
Y &\to cY\text{... rate constant: } b * p_\lambda(c) \\
Y &\to o\text{... rate constant: } d
\end{aligned}
\tag{3.4}
$$

In this notation Y and Z are individual actively or latently infected cells, respectively, and the arrows represent events that lead one type of cell to become the other type. We assume that an actively infected cell can either die (at rate d) or produce a burst of virions (at rate b) that results in the infection of c other cells, where c is Poisson-distributed random variable with parameter λ, $p_\lambda(c) = \frac{\exp(-\lambda)\lambda^c}{c!}$. After a burst event, the original cell dies. Since each birth event

causes the death of the parent cell, the total death rate is $d_y := b + d$. This model

describes a two-type branching process. The reproductive ratio for this model is

$R_0 = \frac{b\lambda}{b+d}$, where $\lambda \geq R_0$. We do not explicitly track free virus, but assume it is at

a level proportional to the number of infected cells. This assumption is valid

because the rates governing the production of virus from infected cells and the

clearance rate of free virus are much higher than other rates, allowing a separation

of time scales. Because we are not interested in blips or other intraday viral

dynamics, this assumption does not influence our results. A method for

calculating the proportionality between free virus and infected cells is provided

in Section 3.4.2, above.

GENERATING FUNCTION ANALYSIS OF THE MODEL

Let

$$
\begin{aligned}
f_1(\xi, \zeta, t) &= E\left[\xi^{y(t)} \zeta^{z(t)} \mid y(0) = 1 \,\&\, z(0) = 0\right], \\
f_2(\xi, \zeta, t) &= E\left[\xi^{y(t)} \zeta^{z(t)} \mid y(0) = 0 \,\&\, z(0) = 1\right]
\end{aligned}
\tag{3.5}
$$

be the basic generating functions for the stochastic process, starting with one

active cell and starting with one latent cell, repsectively. Dummy variables ξ and ζ

correspond to active and latent cells, respectively. The backward Kolmogorov

equations [152] can be represented by the system of coupled ordinary

differential equations

$$\frac{\partial f_1}{\partial t} = b \left(\exp\left[\lambda \left(f_1 - 1 \right) \right] - f_1 \right) + d \left(1 - f_1 \right),$$
$$\frac{\partial f_2}{\partial t} = a \left(f_1 - f_0 \right) + d_z \left(1 - f_0 \right), \tag{3.6}$$

with boundary conditions $f_1(\xi, \zeta, 0) = \xi$ and $f_2(\xi, \zeta, 0) = \zeta$. The birth term $\exp\left(\lambda \left(f_1 - 1 \right) \right)$ follows from the Poisson-distributed offspring distribution with parameter λ.

After LRA therapy, the initial reservoir size is z_0. The residual viremia y_0 is determined by activation-death equilibrium during HAART, and so it is Poisson-distributed with mean $\frac{a z_0}{d_y}$. The probability generating function corresponding to this initial condition is then

$$g_{z_0}(\xi, \zeta, t) = f_2(\xi, \zeta, t)^{z_0} \exp\left[\frac{a z_0}{d_y} \left(f_1(\xi, \zeta, t) - 1 \right) \right]. \tag{3.7}$$

If the initial reservoir size is heterogeneous, then the relevant generating function is given by the sum $\sum_{z_0} p_{z_0} g_{z_0}(\xi, \zeta, t)$, where p_{z_0} is the probability that a patient has z_0 latently infected cells following LRA therapy.

PROBABILITY THAT LRA THERAPY CLEARS THE INFECTION.

The fixed points of the differential equations (3.6) give the ultimate extinction probabilities h_1 and h_2, starting with a single active cell and a single latent cell,

respectively. The probability h_1 is the smallest root of the equation

$$b\left(h_1 - \exp\left[-\lambda\left(1 - h_1\right)\right]\right) = d\left(1 - h_1\right),\qquad(3.8)$$

and the probability h_2 of extinction starting from one latent cell follows,

$$h_2 = \frac{d_z}{a + d_z} + \frac{a}{a + d_z}h_1.\qquad(3.9)$$

Using the above initial condition $\left(y_0, z_0\right)$, the clearance probability P_{clr} is determined by substituting these fixed points into the generating function g_{z_0}:

$$P_{clr} = h_2^{z_0} \exp\left[-\frac{a z_0}{d_y}\left(1 - h_1\right)\right].\qquad(3.10)$$

Define $\mathit{Æ} = \frac{a\left(1 - h_1\right)}{a + d_z}$, the probability that a latently infected cell is fated to activate and establish a rebounding lineage (assuming no interference from any prior rebounding lineage). If this probability is small, then clearance probability is approximately

$$P_{clr} \approx e^{-\frac{\mathit{Æ} z_0\left(a + d_z + d_y\right)}{d_y}}$$
$$\approx e^{-\mathit{Æ} z_0},\qquad(3.11)$$

where the second approximation follows from the fact that active cell dynamics are faster than latent cell dynamics, $d_y \gg a + d_z$. The key parameter determining clearance probability is therefore $\mathit{Æ} z_0$, the expected number of latent cells fated to activate and establish a rebounding lineage.

The observation in the main text that the target value of Q (Fig. 3.2.3A) scales logarithmically with the product of a and reservoir half-life follows from this computation, since reservoir half-life is $\frac{\ln(2)}{a+d_z}$, and $Q \propto \log\left(\frac{1}{z_0}\right)$.

APPROXIMATION OF REBOUND PROBABILITY WITH SIMPLIFIED BIRTH PROCESS.

By approximating the Poisson birth event with a simpler process where only two new infections result per birth, closed forms for the generating functions can be derived, allowing easier computation of rebound probabilities. In this approximation, active cells die at "effective death rate" d_e and give rise to two new active cells at "effective birth rate" b_e. Latent cell dynamics are the same as in the full model. The two parameters are chosen to have the same extinction probability as in the original model,

$$\frac{d_e}{b_e} = h_1, \tag{3.12}$$

and the same expected growth rate (and thus the same R_0)

$$b_e - d_e = b(\lambda - 1) - d. \tag{3.13}$$

The differential equation for active cells, analogous to the first line of (3.6), is then

$$\frac{\partial f_1}{\partial t} = b\left(f_1^2 - f_1\right) + d\left(1 - f_1\right), \tag{3.14}$$

while the differential equation for latent cells is the same as in (3.6). The basic generating functions for this simplified process can be solved in closed form,

$$f_1(\xi, \zeta, t) = \frac{b_e\xi - d_e + d_e(1 - \xi)\exp\left[(b_e - d_e)t\right]}{b_e\xi - d_e + b_e(1 - \xi)\exp\left[(b_e - d_e)t\right]},$$

$$f_2(\xi, \zeta, t) = \frac{1}{\exp\left[a + d_z\right]t}\left[\zeta\int_0^t \exp\left[a + d_z\right]\tau\left(d_z + af_1(\xi, \zeta, \tau)\right)d\tau\right], \qquad (3.15)$$

for the supercritical case $b_e > d_e$.

The generating function $g_{z_0}(\xi, \zeta, t)$ for the process starting at the initial condition described above is again defined as in Eq. (3.7), now using the new functions f_1, f_2. The probability that there are y active cells at time t, written $P(y, t)$, is equal to the coefficient of ξ^y in the Taylor expansion of $g(\xi, \zeta, t)$ around $\xi = 0, \zeta = 1$. Repeated differentiation is computationally costly and subject to compounded rounding errors, and Cauchy's integral formula provides an effective alternative. Following $[59]$, the probability that there are more than n active cells at time t equals

$$P(> n, t) = 1 - \frac{1}{\pi}\int_0^\pi \mathrm{Re}\left[g_{z_0}(e^{i\theta}, 1, t)\frac{1 - e^{-i(n+1)\theta}}{1 - e^{-i\theta}}\right]d\theta. \qquad (3.16)$$

Let $f(\theta)$ be the real integrand in this equation, which presents damped oscillations of period $\frac{2\pi}{n+1}$, where $f\left(\frac{2k\pi}{n+1}\right) = 0$ for integers $1 \leq k \leq \frac{n+1}{2}$. We treat each half-period separately, defining a sequence of approximations a_k $(k = 0, 1, \ldots, n)$ where $a_k \approx \int_{k\frac{\pi}{n+1}}^{(k+1)\frac{\pi}{n+1}} f(\theta)d\theta$. Each a_k $(k > 0)$ is defined by a trapezoidal rule, splitting the interval $\left[k\frac{\pi}{n+1}, (k+1)\frac{\pi}{n+1}\right]$ into 150 equal-width

segments. Since a sharp peak appears at $\theta = 0$ for certain parameter values, the value a_0 is defined using a global adaptive numerical integration routine (quadgk in Matlab R2012b).

Not every a_k must be computed to approximate the integral. In Section 3.4.4 below, we describe an efficient method for selecting a small fraction of the a_k for computation, yielding a several hundred-fold speedup. This method is more reliable for this integral than the Euler summation approach presented in [1].

Estimating "target Q" for 50% probability of rebound

A rough estimate of the initial reservoir size z_0 such that half of patients have rebounded at time t can be obtained directly from the generating function $g_{z_0}(\xi, \zeta, t)$ above. We used this rough estimate as an initial guess for the search algorithm described in Section 3.4.4 below to identify the target Q values in Fig. 3.2.3 B,C.

The probability that there are no actively infected cells at time t is $g_{z_0}(0, 1, t)$. If sufficient time has passed to allow for substantial exponential growth (i.e., $e^{(b_e - d_e)t} \gg 1$), then the integral in Eq. (3.15) is dominated by values at large τ, and so the fraction within the integral may be treated as a constant, $\frac{a d_e}{b_e}$. The probability is then

$$
P(0, t) = g_{z_0}(0, 1, t)
$$
$$
\approx \exp\left[-\frac{a z_0 (1 - h_1)}{d_y (1 - h_1 e^{-rt})}\right] \left[1 - \frac{a}{a + d_z} (1 - h_1) \left(1 - e^{-(a + d_z)t}\right)\right]^{z_0},
$$
$$
(3.17)
$$

where the abbreviation $r = b_e - d_e$ is shorthand for the expected growth rate. If, furthermore, not so much time has passed that the reservoir may be depleted in many patients, then $\left(1 - e^{-(a+d_z)t}\right) \ll 1$ and the expression can be approximated:

$$P(0, t) \approx \exp\left[-az_0 \left(1 - h_1\right)\left(t + \frac{1}{d_y}\right)\right]. \qquad (3.18)$$

Note that d_z has dropped out of this estimate; the rate at which active cells become present is roughly $az_0 \left(1 - h_1\right)$.

To account for growth to the rebound threshold n, we can use a point estimate for the rebound time based on deterministic exponential growth starting at one active cell, $t_{\text{grow}} \approx \frac{\ln(n)}{r}$. The rebound estimate then becomes

$$P(> n, t) \approx 1 - \exp\left[-az_0 \left(1 - h_1\right)\left(t - \frac{\ln(n)}{r} + \frac{1}{d_y}\right)\right]. \qquad (3.19)$$

The required initial reservoir size — and the corresponding target Q — can then be solved for, using $P(> n, t) = 1/2$ and the desired value of t. It follows from this computation that the required z_0 scales inversely with a. Since $Q \propto \log\left(\frac{1}{z_0}\right)$, the target Q scales logarithmically with a and is relatively insensitive to d_z, as observed in the main text.

SIMULATION OF THE MODEL

We use the Gillespie algorithm to track the number of latently and actively infected cells in a continuous time stochastic process. We start with an intial

number of latent cells $z_0 = 10^{-Q}M_{LR}$ and an initial number of actively infected cells y_0 chosen from a Poisson distribution with parameter $a\,10^{-Q}M_{LR}/d_y$, where M_{LR} is the pretreatment latent reservoir size (described in Section 3.4.1). The simulation proceeds until the number of actively infected cells reaches the threshold for clinical detection given by a viral load of 200 c ml^{-1} (equivalent to $y = 3 \times 10^5$ cells total). Because stochastic effects are important only for small y, we switch to a faster deterministic calculation when y reaches a level where the probability of extinction is very low. This switch occurs when the probability that no active cell currently alive establishes a growing infection, h_1^y, declines below 10^{-4}. For each Q value we perform 10^4 simulations.

3.4.4 Additional numerical methods used in analysis of stochastic model

Approximation of the integral in Eq. (3.16)

The sum $\sum_{k=0}^{n} a_k$ described following Eq. (3.16), above, was approximated using the following algorithm:

1. Compute a_k for $k = 0$ to $k = \text{InitialBlockSize} - 1$. Store the sum of these values as S.

2. Set $\text{NextK} = \text{InitialBlockSize}$. Set $\text{BlockSize} = \text{MinBlockSize}$. Set $\text{NumToCompute} = \text{MinNumToCompute}$.

3. While $\text{NextK} < n$, do:

   (a) Initialize $A = 0, B = 0$.

(b) Split the block containing the values a_k from $k = $ NextK to $k = $ NextK + BlockSize $- 1$ into three sections:

    i. FirstPart contains a_k for k from NextK to NextK + NumToCompute $- 1$;

    ii. LastPart contains a_k for k NextK + BlockSize $-$ NumToCompute to NextK + BlockSize $- 1$;

    iii. MiddlePart contains all a_k in between. *These values of a_k are not computed.*

(c) Compute an UpperEstimate for the sum $\sum_{k=\text{NextK}}^{\text{NextK}+\text{BlockSize}-1} a_k$ by adding together all a_k in FirstPart and LastPart, and then approximating the value of each a_k in MiddlePart as the average of the *final two values* in FirstPart.

(d) Compute a LowerEstimate for the sum similarly, except now approximating the value of each a_k in MiddlePart as the average of the *first two values* in LastPart. Since the sequence a_k decreases in an alternating manner, LowerEstimate $<$ UpperEstimate.

(e) If LowerEstimate and UpperEstimate are too far apart (see Notes below), increase NumToCompute by 2 and return to Step 3b. Otherwise, add the average of the two estimates to S and continue.

(f) If NumToCompute $=$ MinNumToCompute (indicating that the

error between upper and lower estimates was never so far apart as to require increasing NumToCompute), increment A by 1; otherwise increment B by 1.

(g) Set NextK = NextK + BlockSize − 1. Reset NumToCompute to MinNumToCompute.

(h) If A ≥ 2B + NumRecentBlocks, then increase BlockSize by a multiplicative factor of BlockIncreaseFactor, rounding to the nearest even integer. Reset A = 0, B = 0.

Notes. LowerEstimate and UpperEstimate are required to have a difference of less than $\frac{2\pi}{n} \times 10^{-5}$ (ensuring a total error in S of under 10^{-5}), and a log-ratio of less than 0.04. The parameter BlockIncreaseFactor is itself adaptive, increasing by $0.4/(1 + 0.5 * (\text{BlockIncreaseFactor} - 1))$ immediately after Step 3h if B = 0; but decreasing by 0.4 (never dropping below 1.05) if B ≥ 1.

Parameters used. InitialBlockSize = 200. MinBlockSize = 40. MinNumToCompute = 6. BlockIncreaseFactor = 1.15. NumRecentBlocks = 6.

Binary search for Q for 50% probability of rebound

The estimate in Section 3.4.3 is used as an initial guess for the post-therapy reservoir size z_0 that would produce a 50% chance of rebound. Survival probability P_{surv} was computed for this initial guess, using the method of

Section 3.4.3. While P_{surv} was not within 10^{-4} of 0.5, a new guess for z_o was made using a bisection method: If P_{surv} is too low (high), but a previously computed value was too high (low), then linear interpolation was used between the current and previous values to select a new z_o that is estimated to have $P_{surv} = 0.5$. If P_{surv} is too low (high), but *no* previously computed value was too high (low), then the guess for z_o was divided (multiplied) by 10. For all results reported, between 3 and 8 iterations were required to obtain the desired P_{surv}. Results were then converted to Q values for a given pre-treatment reservoir size.

3.4.5   Deterministic model of viral dynamics

A deterministic model was used for two purposes: to provide an estimate of R_o (described in Section 3.4.2, above) and to estimate the threshold Q separating the growth-limited and activation-limited regimes. The threshold value is defined as that which equalizes the deterministic rebound time and the expected waiting time until activation of the first cell fated to establish.

Model definition

A complete model of viral dynamics including the latent reservoir is shown in the flow diagram of Fig. 3.4.1. All variables respresent total amounts present in the body. State variables *x, y, v,* and *z* are the number of infectable CD4$^+$ T cells, the number of productively infected cells, the number of free virus particles, and the number of latently infected cells, respectively. Productively infected cells produce virus at rate *k*, die at rate d_y, and transition into latency at a rate γ. Virus is cleared

at rate d_v. The infectivity parameter β determines the rate at which virus infects susceptible host cells. Host cell dynamics are determined by production rate θ and death rate d_x. Latently infected cells reactivate at a rate a and die at a rate d_z. When population sizes are large, this model can be described with a set of differential equations [161, 215, 240, 272]:

$$\dot{x} = \theta - \beta x v - d_x x$$
$$\dot{y} = \beta x v - d_y y + az - \gamma y$$
$$\dot{v} = ky - d_v v$$
$$\dot{z} = \gamma y - az - d_z z$$

(3.20)

Here we make a number of simplifying assumptions valid for understanding dynamics leading to rebound. Because the terms k and d_v are an order of magnitude larger than other rates in the system, we can apply a separation of timescales and assume that free virus particle levels change so quickly that they track infected cell levels. Formally, this is accomplished by setting $\dot{v} = 0$, leading to $v = ky/d_v$ and only three differential equations. For viral loads at or below rebound levels, uninfected CD4$^+$ T cells do not become limited and can be assumed to remain at their pre-infection/post-long-term-HAART steady-state level of $x = \theta/d_x$, and hence $\dot{x} = 0$. Moreover, at low viral loads, new influx into the latent reservoir can be ignored ($\gamma \approx 0$).
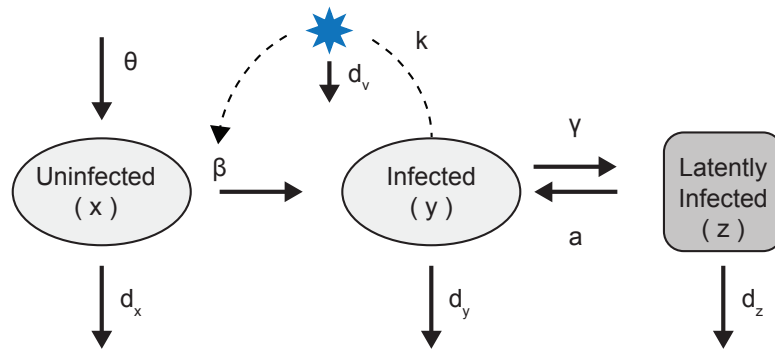
Figure 3.4.1: Schematic of the deterministic viral dynamics model including the latent reservoir. State variables x, y, v, and z are the number of infectable CD4$^+$ T cells, the number of productively infected cells, the number of free virus particles, and the number of latently infected cells, respectively. Productively infected cells produce virus at rate k, die at rate d_y, and transition into latency at a rate γ. Virus is cleared at rate d_v. The infectivity parameter β determines the rate at which virus infects susceptible host cells. Host cell dynamics are determined by production rate θ and death rate d_x. Latently infected cells reactivate at a rate a and die at a rate d_z.

These assumptions lead to the reduced set of equations:

$$\dot{y} = az + (R_\circ - 1)\, d_y y$$
$$\dot{z} = -(a + d_z)\, z,$$

(3.21)

where $R_\circ$ is the combined viral fitness parameter $(\theta k \beta) / (d_v d_x d_y)$ describing the expected number of secondary infected cells produced by a single infected cell introduced to an uninfected host. As in the stochastic model described in Section 3.4.3, the initial conditions $(y_\circ, z_\circ)$ are such that $z_\circ$ is the reservoir size following LRA therapy, and — assuming that latent cell dynamics are much slower than active cell dynamics $(a + d_z \ll d_y)$ — the residual active infection is $y_\circ \approx a z_\circ / d_y$.

CALCULATING REBOUND TIME FROM THE DETERMINISTIC MODEL

Following HAART interruption (with or without LRA therapy), the number of infected cells (and thus viral load) grows according to Equation (3.21) with latent cells at a transiently constant value $z_\circ$. Let $y_r = r y_\circ$ be the infection size at which rebound is detected (e.g., ≈ 200 c ml^{-1}). The parameter r is the "rebound factor", the amount by which the infection must grow in order for rebound to be detected.

This equation can be solved exactly:

$$y(t) = y_\circ \frac{R_\circ e^{d_y (R_\circ - 1) t} - 1}{R_\circ - 1},$$

(3.22)

giving the time to rebound as

$$t_0 = \frac{\ln\left(\frac{1+r(R_0-1)}{R_0}\right)}{d_y(R_0-1)}.$$ (3.23)

This equation is appropriate when cells exit the reservoir frequently (i.e. without highly effective LRA therapy). Incorporating reservoir decay negligibly changes the results because $a + d_z \ll 1$. In Section 3.4.2 above, this equation was used to calibrate R_0 based on observed rebound time t_0.

Following LRA therapy, the size of the reservoir is reduced to $10^{-Q}M_{LR}$. Let $q = 10^{-Q}$. Since the residual active infection y_0 also scales by q, the rebound factor increases to r/q. The time to rebound is now increased to

$$t_1 = \frac{\ln\left(\frac{1+(r/q)(R_0-1)}{R_0}\right)}{d_y(R_0-1)}$$ (3.24)

Eradication therapy therefore extends the rebound time by an amount Δt:

$$\Delta t = t_1 - t_0$$

$$= \frac{1}{d_y(R_0-1)}\ln\left(\frac{1+(r/q)(R_0-1)}{1+r(R_0-1)}\right)$$ (3.25)

$$\approx \frac{\ln(1/q)}{d_y(R_0-1)}$$

where the approximation is valid for $r(R_0-1) \gg 1$. This order relationship is very likely to hold as the rebound factor is ≈ 100 or more. The only way for the relationship to fail would be for R_0 to be in a very narrow range just above 1.

Under the deterministic model, rebound time increases only logarithmically with the decrease in LR size. The stochastic model agrees with this relationship only in the *growth-limited regime*, where activation of fated-to-rebound cells is common. At higher Q, the waiting time until this activation occurs exceeds the deterministic growth time; in this *activation-limited regime*, the stochastic model predicts rebound times well in excess of those predicted by the deterministic model (Fig. 3.4.2). As a rough estimate, the waiting time in the stochastic model is $\frac{1}{a(1-h_1)z_0}$, where h_1 is the probability that a reactivating lineage goes extinct (defined in Equation (3.8)). The threshold drug efficacy Q_T defines the boundary between the two regimes. It can be estimated by solving numerically for the log-efficacy that sets deterministic growth time from Equation (3.24) equal to the stochastic waiting time:

$$\frac{\ln\left(\frac{1 + r\,10^{Q_T}(R_0 - 1)}{R_0}\right)}{d_y(R_0 - 1)} = \frac{1}{a\left(1 - h_1\right) M_{LR}\, 10^{-Q_T}}. \tag{3.26}$$

This threshold can be observed in the upward inflection in the rebound time curves in Figs. 3.2.2C and 3.4.2 occurring at $Q \approx 1 - 2$.

Fig. 7.3.2 shows that Q_T increases with pretreatment reservoir size M_{LR} and decreases with variance parameter λ, since higher values of λ increase the extinction probability h_1. For pretreatment reservoir size $M_{LR} = 10^6$, reduction

99

Figure 3.4.2: A fully deterministic model is a poor predictor of rebound times. A) Equation (3.25) was used to calculate the time to rebound for a given LRA drug efficacy, Q. The deterministic model assumes that LRA reduces the size of the reservoir and hence the residual viral load by Q orders of magnitude and then tracks the time for the viral population to grow to 200 c ml^{-1}. B) Rebound times calculated from the deterministic model for $R_0 = 2$ are compared to median rebound times calculated from the stochatic process. The models agree only for small Q. The Q values where the models diverge corresponds to the transition between the *growth-limited regime* and the *activation limited regime*.

of approximately 30- to 1,000-fold is required to reach the activation-limited regime, in which substantial increases in rebound time may be achieved.

3.4.6   VIREMIA DURING ADMINSTRATION OF LRA THERAPY

If LRA therapy reduces reservoir size by inducing activation of latently infected cells, then an increase in viral load during therapy can be expected. The precise dynamics of viral load during LRA therapy depend on the efficacy and duration of therapy. To model these dynamics, we assume LRA therapy is administered continuously for a period of time τ (ignoring pharmacokinetics) and increases the reactivation rate of latently infected cells. During this time HAART is co-administered, and we assume no new infections can occur. Since transcription

Figure 3.4.3: Schematic of viral dynamics during LRA therapy. (A) The model tracks latently infected cells, actively infected cells, and free virus, distinguishing between cells in which HIV-1 transcription is reactivated by LRA (y') or naturally (y). (B) Illustration of typical viral load dynamics during course of LRA therapy. LRA therapy increases the reactivation rate of cells from the LR, causing residual viral load to increase. The timing and magnitude of this peak allow for an estimation of the efficacy Q. We assume that HAART is administered for a short while beyond the end point of LRA therapy, preventing the reactivated cells from starting new infections. All symbols are defined in the text.

at the HIV-1 LTR may be reactivated in resting CD4$^+$ T cells while the cell otherwise retains a resting phenotype, cellular functions in these LRA-activated cells may proceed at a slower rate [275]. We therefore track those cells reactivated by LRA separately from those reactivated normally by antigenic stimulus.

The model we use to consider this scenario is shown schematically in Fig. 3.4.3A and described formally in the next section. Since our goal is to track viral load during a period when HAART is administered, we consider all cells capable of producing virions, regardless of replication capability. As this collection of resting cells likely exceeds the LR by an order of magnitude or more (f_V versus f_{LR}, Fig. 3.2.1), and since therapy increases activation rate, large numbers of cells can be assumed to activate daily, and a deterministic model

similar to the one presented in Section 3.4.5, above, suffices. Latent cells (z) die at a rate d_z and can be reactivated naturally (a) or by LRA-induced mechanisms (a'). Naturally reactivated cells (y) die at rate d_y, and LRA-reactivated cells (y') die at rate d_y' (which may be lower than d_y) [275]. Free virus is produced at a rate k from naturally reactivated cells and a rate k' from LRA-reactivated cells. Since induced cells may have smaller burst sizes [276], k' may be less than k.

We use this model to relate change in residual viremia over time to the log-efficacy of LRA therapy, Q. Define $a_z = a + a' + d_z$, the total rate at which z decays during treatment. Assuming that this decay affects latent cells regardless of whether the proviral sequence is able to replicate, the fraction of the latent reservoir that remains after treatment duration τ is $q = e^{-a_z \tau}$. Log-efficacy of therapy is $Q = log_{10}(q)$, resulting in the relationship $a_z = \ln(10) \cdot Q/\tau$. Viral load approximately follows a biexponential curve, generally reaching a peak quickly (time determined by the faster of the two rates a_z and d_y') and then decaying slowly (at the slower of the two rates) (Fig. 3.4.3B).

The height of this peak and the time after treatment initiation at which it occurs are

$$
\begin{aligned}
v(t_{max}) &\approx w_\circ \frac{k'}{k} \frac{a'}{a} \frac{d_y}{d_y'} \\
&\approx w_\circ \frac{k'}{k} \frac{d_y}{d_y'} \frac{\ln(10) * Q}{a\tau},
\end{aligned}
\tag{3.27}
$$

$$
t_{max} = \frac{\ln(d_y') - \ln(a_z)}{d_y' - a_z},
\tag{3.28}
$$

102

where $w_0 \approx 2$ c ml^{-1} is the residual viral load (not necessarily infectious) before LRA therapy. This peak is reached before treatment ends if $\tau > 1/d_y'$ and $Q > \frac{1}{\ln(10)} \approx 0.43$. Derivations for these results are provided in the next section, and examples for selected treatment parameters are provied in Table 3.4.1. Even mildly effective therapy $(Q = 1)$ can result in large increases in residual viremia from the baseline level, appearing after a few days of treatment. Viremia is lower if the same reservoir reduction Q is achieved over a longer treatment time τ. Viremia is also decreased if LRA-induced cells have a lower burst rate than normally activated cells $(k'/k < 1)$. If LRA-induced cells have a longer lifespan $(d_y/d_y' > 1)$, then peak viremia increases but is delayed.

In vivo estimates of LRA therapy efficacy may be obtained from measurement of viral load during therapy. Highly effective therapy is predicted to result in large, observable increases in residual viremia during continuous administration (Table 3.4.1). It is important to note that this conclusion applies only to forms of LRA that reactivate latently infected cells without damaging viral production in these cells.

DETAILED CALCULATIONS

We extend Equations 3.20 and 3.21, tracking latently infected cells z, productively infected cells reactivated naturally y, productively infected cells

103

Table 3.4.1: The timing and size of peak viremia, as well as viremia at treatment end, depend on the efficacy and duration of LRA therapy, the change in burst size, and the change in infected cell lifespan. Therapy protocol is described in the text and illustrated in Fig. 3.4.3. Time to peak viremia was calculated using Eq. 3.28. Peak viremia was calculated using Eq. 3.34, of which Eq. 3.27 is an approximation. Symbol definitions are provided in the text. Parameters used: $a = 7 \times 10^{-5}$ day^{-1}, $d_z = 4.6 \times 10^{-4}$ day^{-1}, $z_0 = M_v = 43 \times 10^6$ cells, $d_y = 1$ day^{-1}, and $w_0 = 2$ c ml^{-1} (implying $\frac{k}{d_v} = \frac{1}{1500}$ c ml^{-1} cell^{-1}).

Log-efficacy LRA (Q)	Treatment time (τ) (days)	Fold increase in burst size (k'/k)	Fold increase in lifespan (d_y/d'_y)	Time to peak (days)	Peak viremia (c ml^{-1})	Viremia at treatment end (c ml^{-1})
1	15	1	1	2	3×10^3	500
1	60	1	1	3	10^3	100
1	180	1	1	4	300	30
1	180	0.1	1	4	30	4
2	15	1	1	2	5×10^3	100
2	16	1	1	2	5×10^3	100
2	180	1	1	4	700	7
4	15	1	1	1	8×10^3	5
4	15	1	10	4	2×10^4	8×10^3
4	15	0.1	1	1	800	< 2
4	15	0.1	10	4	2×10^3	800
4	60	1	1	2	3×10^3	< 2
4	180	1	1	3	10^3	< 2

reactivated by LRA therapy y', and free virus v:

$$\dot{z} = -(a + a')z - d_z z$$

$$\dot{y} = az - d_y y$$

$$\dot{y'} = a'z - d'_y y'$$

$$\dot{v} = ky + k'y' - d_v v$$

(3.29)

These equations incorporate the simplifying assumptions that fully effective HAART yields $R_0 = 0$, that uninfected $CD4^+$ T cells do not become limited and remain at their pre-infection/post-long-term-HAART steady state level, and that new influx into the LR can be ignored.

Using the separation of timescale for virus dynamics (v) and the initial conditions $z(0) = z_0$, $y(0) = az_0/d_y$, $y'(0) = 0$, these equations can be solved:

$$z(t) = z_0 e^{-a_z t}$$

$$y(t) = \frac{az_0}{d_y - a_z}\left(e^{-a_z t} - e^{-d_y t}\right) + \frac{az_0}{d_y}e^{-d_y t}$$

$$y'(t) = \frac{a'z_0}{d'_y - a_z}\left(e^{-a_z t} - e^{-d'_y t}\right)$$

$$v(t) = \frac{ky(t) + k'y'(t)}{d_v},$$

(3.30)

where $a_z = a + a' + d_z$. Note that the timescale separation for $v(t)$ is based on the assumption that k, k', and d_v greatly exceed other rates, which may be violated if LRA-reactivated cells have much smaller burst rate k' and/or much higher death rate d'_y than normally reactivated cells. In general this approximation may only slightly overestimate viremia on the first day after LRA is started.

The amount by which LRA therapy increases the reactivation rate (a') can be related to the fraction to which the reservoir is reduced (q), or equivalently, to the log-efficacy $Q = log_{10}(q)$, after a treatment time τ:

$$z(\tau) = z_0 e^{-a_z \tau}$$

$$= qz_0,$$

(3.31)

and so the reactivation rate is:

$$a_z = \frac{-ln(q)}{\tau} - (a + d_z)$$
$$= \frac{ln(10) \cdot Q}{\tau}$$

(3.32)

Note that some natural activation and cell death contribute to Q, so that the absence of treatment does not correspond to $Q = 0$ but to $Q = log_{10}(e)(a + d_z)\tau$, which is generally small.

From these equations, we can determine the expected changes in residual viral load over time during LRA therapy. The contribution to residual viremia from naturally activating cells (initially $w_0 = 2 \text{ c ml}^{-1}$) only decreases during treatment, as the reservoir is depleted. If any increases in observable residual viremia occur in response to LRAs, it will come from compartment y' activated by the drug. This contribution follows a biexponential curve, increasing to peak value quickly (at roughly the smaller of the two times $1/d'_y$, $1/a_z$ days), then decaying more gradually to zero (at the smaller of the two rates d'_y, a_z per day).

We can calculate precisely the time of the peak in residual viremia due to contributions from newly reactivated cells :

$$t_{max} = \frac{ln(d'_y) - ln(a_z)}{d'_y - a_z}.$$

(3.33)

106

The viral load at this peak is:

$$
\begin{aligned}
v(t_{max}) &= \frac{k'a'z_0}{d_v d_y'} u^{\frac{u}{1-u}} \\
&= \frac{k'a'}{d_v a}\frac{d_y}{d_y'} y_0 u^{\frac{u}{1-u}} \\
&= w_0 \frac{k'}{k}\frac{a'}{a}\frac{d_y}{d_y'} u^{\frac{u}{1-u}} \\
&\approx w_0 \frac{k'}{k}\frac{a'}{a}\frac{d_y}{d_y'} \\
&\approx w_0 \frac{k'}{k}\frac{d_y}{d_y'}\frac{\ln(10)\cdot Q}{a\tau},
\end{aligned}
\tag{3.34}
$$

where $w_0 = ky_0/d_v$ is the residual viral load before LRA therapy and $u = a_z/d_y'$. The first approximation holds when u is small (if Q/τ is small compared to d_y'), and the second holds when $a + d_z \ll a'$ (Q/τ is large compared to $a + d_z$). The peak viral load occurs during treatment when $t_{max} < \tau$, which holds if and only if both $\tau d_y' > 1$ and $\ln(10)\cdot Q > \tau d_y'$, which places requirements on both treatment time ($\tau > 1/d_y'$) and treatment strength ($Q > \frac{1}{\ln(10)} \approx 0.43$). When LRA is started, viral load approaches the peak linearly with an initial slope of $a'z_0$. Viremia at the end of treatment can be found by substituting in τ for t in Eq. (3.30).

3.5 MANUSCRIPT INFORMATION

3.5.1 IN SUBMISSION AS

This manuscript is being submitted as:

Hill AL, Rosenbloom DIS, Fu F, Nowak MA, Siliciano RF. Predicting outcomes of treatments to eradicate the HIV-1 latent reservoir.

4

Imitation dynamics of vaccination behavior on social networks

4.1 INTRODUCTION

PREEMPTIVE VACCINATION IS A FUNDAMENTAL STRATEGY for controlling infectious diseases [43]. While there is vigorous debate about the civil liberties implications of mandatory versus voluntary vaccination policies [57], mounting evidence shows that voluntary vaccination plans fail to protect populations adequately [22–25, 36, 38, 56, 91, 98, 253, 316]. A recent example of this failure

is the sharp decline in take-up of the combined measles-mumps-rubella vaccination in Britain soon after administering it to children was made voluntary [142]. Because of declining familiarity with the disease and rising fears of vaccine complications, parents hoped to avoid the alleged vaccination health risk to their own children while implicitly relying on enough other children getting vaccinated to provide herd immunity. The "public good" created by herd immunity gives rise to an enduring social dilemma of voluntary vaccination.

Classical game theory predicts that, when individuals act in their own interests with perfect knowledge of their infection risk, their vaccination decisions converge toward a Nash equilibrium, at which no individuals could be better off by unilaterally changing to a different strategy [24, 25]. Although this equilibrium is the result of each individual following her self-interest, it may lead to suboptimal vaccination coverage for the community [98]. The collective result of vaccination decisions determines the level of population immunity and thus the severity of an epidemic strain. With increasing levels of vaccination coverage in the community, even the individuals who are unvaccinated are less likely to become infected; therefore, they have less incentive to get the vaccine. This scenario naturally leads to the "free riding" problem that is commonly observed in public goods studies [116].

Previous studies of vaccinating dynamics have typically combined a game-theoretic model assuming full rationality and complete information with a model of disease transmission in either homogeneously mixed populations [24, 25] or random networks [242]. In studies where the

assumption of rationality is relaxed, deterministic evolutionary dynamics still recover equilibrium states equivalent to those predicted by models of rational agents [23]. It is worth noting that aggregate population models have been parameterized with empirical data to quantitatively predict vaccinating behavior in some cases [22, 23, 98]. Here we extend this previous work by accounting for decision-makers' social networks and their use of anecdotal information in making vaccination choices. Individuals have incomplete information and tend to rely on salient anecdotes from friends and the media in order to form opinions of disease risk and prevention [145, 228, 314]. The rise to prominence in the British media of isolated cases linking the pertussis vaccine and brain damage triggered a sharp decline in coverage in the late 1970s, demonstrating the power of the anecdote [23, 208]. Apart from these prominent cases, each person can encounter different anecdotal evidence, depending on her social network [84, 242]. Illness of a close friend can impact one's perception of infection risk and the importance of prevention in far more powerful ways than media reports can [228].

Motivated by the above considerations, we propose a simple agent-based model in the spirit of evolutionary game dynamics [195, 213, 217] to study the voluntary vaccination dilemma. In order to make precise predictions, we couple the vaccination dynamics with an epidemiological model, in particular the *SIR* model, which tracks populations of *s*usceptible, *i*nfected, and *r*esistant/vaccinated individuals over time, within a single season or epidemic. Such models have been used, for example, to design clinical trials of vaccines or

to predict whether a vaccination program will halt an epidemic before it spreads to much of the population [72, 176].

Our model captures the strategic interaction between vaccinating and free-riding individuals in the following way. Individuals decide whether to vaccinate during a vaccination campaign, before the seasonal epidemic begins. The epidemiological model then determines whether each susceptible (unvaccinated) individual becomes infected at some point during the season. Once the epidemic ends, individuals can revise their vaccination decision for the next season. Such a model is most appropriate for describing infections such as influenza. Flu vaccines are typically available prior to a predicted outbreak and are effective for only one season due to mutation of pathogens and waning immunity [38, 316].

4.2   Model & methods

Consider a well-mixed population of individuals with a voluntary vaccination option. We model the vaccination dynamics as a two-stage game (as illustrated in figure 4.2.1). The first stage is a public vaccination campaign, which occurs before any infection. At this stage, each individual decides whether or not to vaccinate. Vaccination incurs a cost, V, to the vaccinated individual, and it grants perfect immunity from the seasonal infectious disease. The total cost of vaccination includes the immediate monetary cost, the opportunity cost of time spent to get the vaccine, and any perceived or actual adverse health effects. In the second stage, the epidemic strain infects an initial number of individuals I_0 and then

spreads according to SIR dynamics, with per-day transmission rate r and recovery rate g (see the supplementary materials for model details). The epidemic continues until there are no more newly infected individuals (which occurred in under 200 days for all cases simulated). The final size equation [72] gives the infection risk for an infinite population (see supplementary materials for derivations):

$$w(x) = \frac{R(\infty)}{1-x} = 1 - e^{-R_0 R(\infty)},\qquad (4.1)$$

where $R(\infty)$ is the final size of the epidemic (fraction that have been infected at some point in the season), which satisfies $R(\infty) = (1-x)(1 - e^{-R_0 R(\infty)})$; R_0 is the basic reproductive ratio; and x is the fraction of vaccinated individuals.

The infection cost I includes health care expenses, lost productivity, and the possibility of pain or mortality. After the epidemic, the individuals with the highest payoffs are those who declined vaccination but avoided infection. We call these lucky individuals successful free-riders, as they benefit from others' vaccination efforts. The game dynamics remain unchanged if we rescale the payoffs by defining the relative cost of vaccination $c = V/I$ ($0 < c < 1$). The values of c appropriate for modeling a particular disease can be estimated from surveys of health opinions, behaviors, and outcomes, as done by, e.g., Galvani et al. [98], but in general vaccination cost should be low relative to the cost of infection. The Nash equilibrium of this game can be solved by setting the expected cost of vaccination equal to that of non-vaccination, which implies the
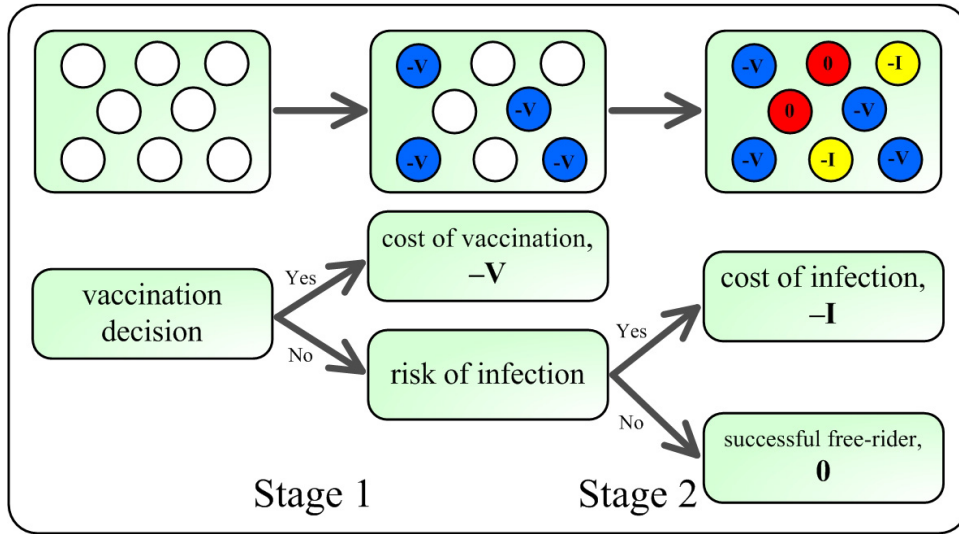
Figure 4.2.1: Schematic illustration of our model. We model the vaccination dilemma as a two-stage game. At Stage 1 (vaccination choice), a proportion x of the population decides to vaccinate. Vaccination costs V and provides perfect immunity from the infectious disease. At Stage 2 (health outcome), we use the Susceptible-Infected-Recovered model to simulate the epidemiological process. Each unvaccinated individual faces the risk of infection during the seasonal epidemic outbreak. The cost of infection is I. Those unvaccinated individuals who remain healthy are free-riding off the vaccination efforts of others, and they are indirectly protected by herd immunity.

mixed strategy

$$x^* = 1 + \frac{\ln(1 - c)}{cR_0}. \tag{4.2}$$

This level of vaccination uptake falls short of the social optimum $x_h = 1 - 1/R_0$, which both achieves herd immunity (full protection of unvaccinated individuals) and minimizes the sum of all individuals' costs related to both vaccination and infection (see supplementary materials). The misalignment between individual and group interests leads to a social dilemma.

Here, we relax the assumption of rationality and study this vaccination dilemma from an evolutionary perspective. Each season, an individual adopts a pure strategy, which determines whether or not she vaccinates. At the end of the season, each individual decides whether to change her strategy for the next season, depending on her current payoff. Specifically, individual i randomly chooses individual j from the population as role model. The strategy of a role model with higher payoff is more likely to be imitated. We suppose that the probability that individual i adopts individual j's strategy is given by the Fermi function [31, 298, 311, 312]

$$f(P_j - P_i) = \frac{1}{1 + \exp[-\beta(P_j - P_i)]},\qquad(4.3)$$

where β denotes the strength of selection ($0 < \beta < \infty$).

This updating dynamic diverges from a fully rational model in two ways. First, individuals adjust their strategies retrospectively, in response only to the observed payoff outcomes and not the expected payoffs of strategies. In a population with low vaccination uptake, most non-vaccinators fall ill, but if individual i happens to choose one of the few successful free-riders as a role model, then she will be more likely to imitate the free-rider's strategy. Second, the strength of selection parameter introduces a stochastic element to the model: for small β (weak selection), individuals are less responsive to payoff differences, and an individual with a high payoff may adopt the strategy of a less successful role model. Large values of β (strong selection) diminish this stochastic effect, and individuals reliably switch to (or keep) the strategy with the higher observed

payoff, even if the payoff difference is small. Previous work using the same update dynamic has characterized agents with high β as being more rational [298]. This characterization is not appropriate in our context, as higher β only increases an agent's sensitivity to the (perhaps unrepresentative) observed payoff, not the expected payoff.

The model presented here can be conveniently extended to structured populations by restricting the neighborhood of individuals whom one can infect or imitate. In addition to the well-mixed case, we simulated populations structured as square lattices, Erdős-Rényi random graphs [88], and Barabási-Albert scale-free networks [20] (see supplementary materials). The initial state consists of equal fractions vaccinators and unvaccinators, randomly distributed throughout the population. Each two-stage iteration (vaccination strategy updating followed by an epidemic process) updates the frequencies of each strategy. Since we are interested primarily in the effect of population structure on vaccination coverage (rather than on infection risk), we calibrated epidemic parameters to ensure that the infection risk in an unvaccinated population is equal across all population structures [242] (see supplementary materials). Each simulation was run for 3,000 iterations. The long run equilibrium results shown in figures 4.3.1–4.3.3 represent the average of frequencies over the last 1,000 iterations in 100 independent simulations. Our presented results are robust to population size N for all population structures examined, as long as $N > 200$.

4.3 RESULTS

In the vaccination game, if all of one's neighbors adopt one strategy, then it is advantageous to adopt the opposite strategy. We therefore always find persistent polymorphisms of vaccinated and unvaccinated individuals for intermediate values of c. Figure 4.3.1 plots both the equilibrium frequency of (a) vaccinated and (b) infected individuals for different values of c and β in the well-mixed imitation dynamics. We find qualitative agreement between stochastic simulations and an analytical prediction that uses both the equation for infection risk (4.1) and an infinite-population approximation of the imitation dynamics (described in supplementary materials).

For weak selection ($\beta = 1$ in figure 4.3.1), the imitation dynamics approximate the rational equilibrium x^* given in equation (4.2). One can understand this observation analytically by noting that the strategy update equation (4.3) is roughly linear for small β. First-order approximation of the imitation dynamics closely approximates the replicator dynamics [133, 271, 306], which in this game converge to the unique evolutionarily stable strategy–the Nash equilibrium (see supplementary materials). As vaccination falls with increasing c, the final size of the epidemic grows. Above a high cost threshold $c_H \approx 0.893$, no one chooses vaccination and the epidemic reaches its maximum size.

Strong selection in the imitation dynamics (represented by $\beta = 10$ in figure 4.3.1) can decrease vaccination uptake below the level predicted by the

Figure 4.3.1 *(following page)*: Vaccination dynamics in well-mixed populations. The fractions (a) vaccinated and (b) infected are shown as functions of the relative cost of vaccination, c, for the intensity of selection $\beta = 1$ and 10. The lines are analytical predictions from deterministic equations (see supplementary materials). The deviation between simulation and theory is largely due to stochasticity in disease transmission: holding vaccination constant at some level below the herd immunity threshold ($1 - 1/R_0 = 0.6$), simulated infection risk is smaller than the prediction in equation (4.1) (see suppl. figure 4.5.1b). Individuals in the simulation respond to this decreased risk by vaccinating less than in the theory, which in turn leads to a larger epidemic versus the theory. Strong selection magnifies individuals' responses, producing larger deviations. For vaccination coverage above the theoretical herd immunity level, the deterministic approximation underestimates infection risk, leading to an opposite deviation at low c. Parameters: population size $N = 5000$, $R_0 = 2.5$ (realized by setting $r = 5/(6N)$ and $g = 1/3$), number of infection seeds $I_0 = 5$.
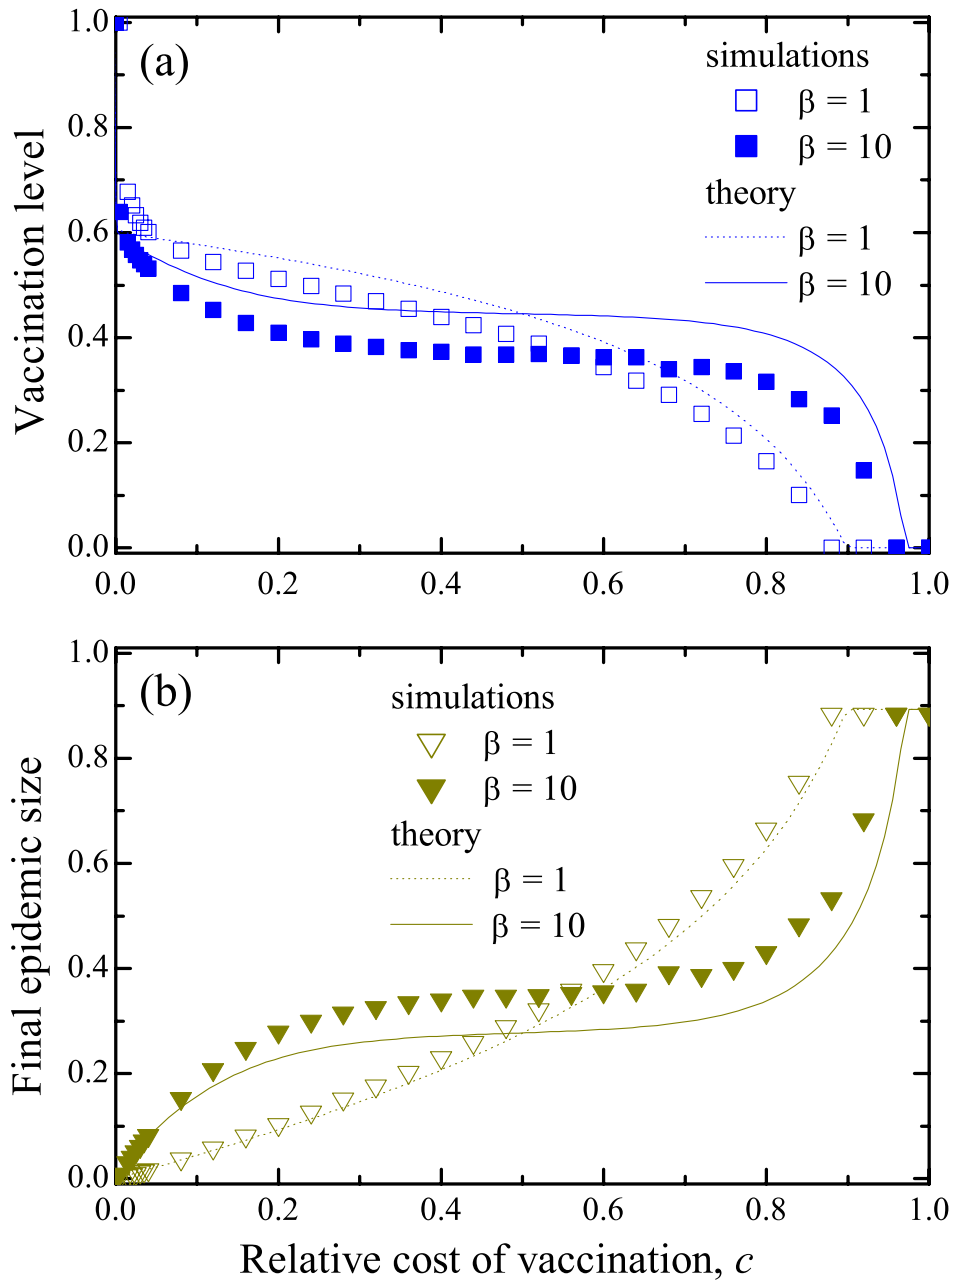
Figure 4.3.1 *(continued)*

rational equilibrium. In other words, individuals who carefully attend to peers'

health outcomes and reliably copy the behavior of successful peers will end up

attempting to free-ride more than they rationally "ought" to. If, for example,

infection is twelve times as costly as vaccination (namely, $c = 0.08$, a reasonable

assumption for influenza, see supplementary materials), then strong selection in

our model lowers vaccination coverage by 8 percentage points versus weak

selection (figure 4.3.1a), which increases the epidemic size from 4% of the

population to 15% of the population (figure 4.3.1b). With increasing cost, the

equilibrium vaccination coverage follows a rotated "S" curve, dropping rapidly

(slope $\approx -\frac{\beta}{2}$) from the herd immunity threshold at low values of c, reaching a

plateau near $1 - \frac{2\ln 2}{R_o}$ for intermediate values of c, and then dropping rapidly to

zero as c grows large. The threshold c_H increases with selection strength

(figure 4.3.1a).

Results are qualitatively similar for any basic reproductive ratio R_o of the

infection. suppl. figures 4.5.5, 4.5.6 compare the cases $R_o = 2.5$ and $R_o = 6$. The

higher value increases infection risk, making the population respond with

increased vaccination. Increasing R_o also raises the threshold c_H.

Restricting interaction to local neighborhoods partly ameliorates the

free-riding problem, but introduces greater sensitivity to the cost parameter c

(figure 4.3.2). We consider a population of individuals arranged on a square

lattice where each individual has four immediately adjacent neighbors. While the

vaccination coverage in well-mixed populations drops from herd immunity levels

Figure 4.3.2: Vaccination dynamics in lattice populations. Left panels (a), (b) show the fractions vaccinated and infected, respectively, as functions of c for the intensity of selection $\beta = 1$ and 10. right panels (c), (d) show snapshots of the system at equilibrium frequencies with weak and strong selection, respectively. Blue denotes vaccinated individuals, red successful free-riders, and yellow infected individuals. Strong selection breaks apart clusters of vaccinators: 54% of vaccinated individuals' neighbors are also vaccinated in (c), versus only 49% in (d). Parameters: population size $N = 100 \times 100$ with von Neumann neighborhood, disease transmission rate $r = 0.46$, recovery rate $g = 1/3$, number of infection seeds $I_0 = 10$, (c)(d) $c = 0.08$, (c) $\beta = 1$, (d) $\beta = 10$. The lines in (a) and (b) are visual guides.

as soon as c increases above zero, restricted spatial interaction promotes near-universal coverage at a range of positive c, preventing the epidemic. To give a simple operational definition, we say that vaccination "prevents the epidemic" in a structured population if the average final epidemic size is less than twice the size of the initial inoculum. Define as c_L the critical vaccination cost below which the epidemic is prevented. For weak selection on the lattice ($\beta = 1$ in figure 4.3.2), we get $c_L \approx 0.022$. Above this threshold, the vaccination level drops precipitously, causing an epidemic that is even larger than in the well-mixed case.

At higher selection strength, the threshold c_L is lower, and vaccination coverage is even more sensitive to costs rising above c_L (figure 4.3.2a). The high cost threshold c_H rises with selection strength, meaning that the transitional region between c_L and c_H, where vaccinated and unvaccinated individuals coexist, widens with larger β. Holding c constant at a value above c_L, increasing the strength of selection leads to more free-riding attempts, breaking apart clusters of vaccinators, thus allowing a larger epidemic to occur (figure 4.3.2c versus 4.3.2d).

Most actual populations are heterogeneous in the sense that different individuals may have different numbers of neighbors (i.e., degree) [20]. To account for this feature, we consider vaccination dynamics on Erdős-Rényi random graphs, which have moderate degree heterogeneity; on scale-free networks, which have an even more variable degree distribution, our results are similar (see supplementary materials).

Higher vaccination coverage is typically required to achieve herd immunity in populations with greater degree heterogeneity [237] (see also suppl.

figures 4.5.2–4.5.4). This increased vulnerability to epidemic attacks reduces the temptation to free-ride, actually making it easier for a population of selfish imitators to achieve the high vaccination threshold required for herd immunity. The threshold cost c_L therefore increases versus the lattice case. Vaccination coverage drops after cost exceeds this threshold, although the effect is not quite as extreme as in lattice populations (figures 4.3.3a and 4.3.3b). Similarly to lattice populations, increased selection strength increases the size of the intermediate region between c_L and c_H.

Degree heterogeneity triggers a broad spectrum of individual vaccinating behavior. Specifically, an individual's vaccination strategy is now influenced by her role in the population, and "hubs" who have many neighbors are most likely to choose to be vaccinated, as they are at greatest risk of infection (figures 4.3.3c and 4.3.3d). Hubs that do manage to free-ride successfully become victims of their own success, as their vaccinated neighbors of smaller degree are likely to imitate them and switch strategies, potentially infecting the hubs in the following season.

4.4 Discussion & conclusion

Our model shows how incomplete information and strong selection (high payoff-sensitivity, parameterized by β) in a population of imitators causes the vaccination coverage to fall well short of the social optimum, and even below the Nash equilibrium. Weak selection in a well-mixed population recapitulates the
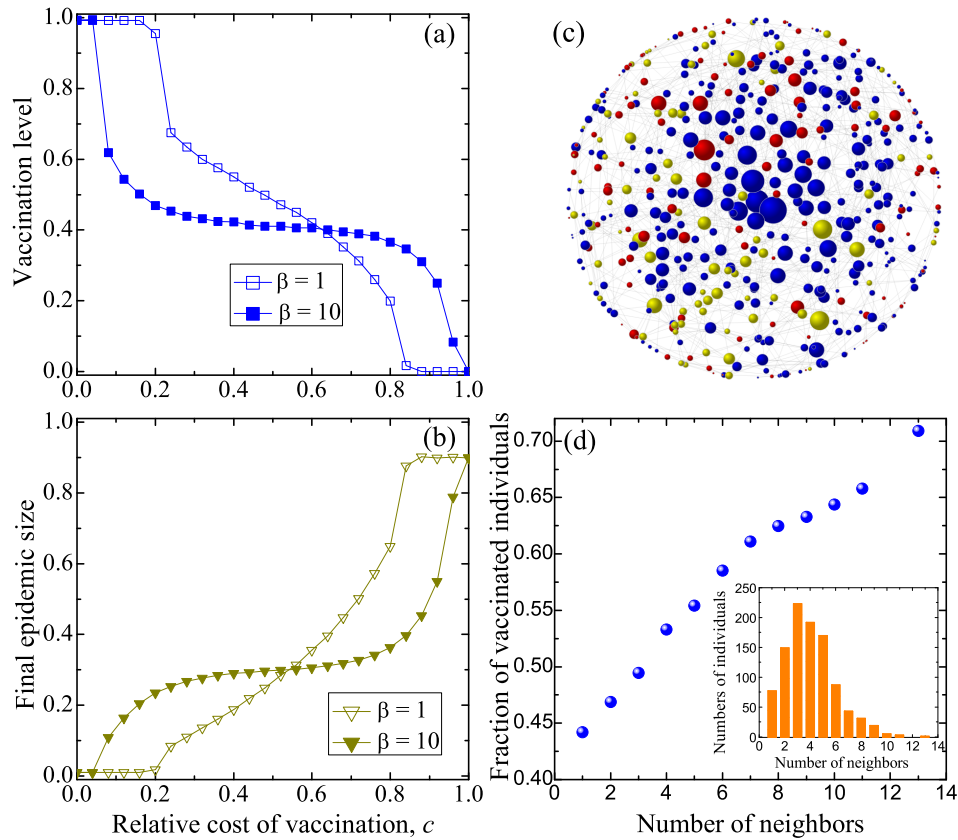
Figure 4.3.3: Vaccination dynamics in random network populations. Left panels (a), (b) show the fractions vaccinated and infected, respectively, as functions of c for the intensity of selection $\beta = 1$ and 10. Right panels: (c) Snapshot of a single simulation on a random network at equilibrium frequencies. The size of a node corresponds to its degree (number of neighbors). Blue nodes are vaccinated, yellow are infected, and red are successful free-riders. (d) The frequency of vaccination on a random network, as a function of the number of neighbors an individual has. The inset in panel (d) shows the degree distribution of the random network. Parameters: (a)-(d) average degree $\bar{k} = 4$, disease transmission rate $r = 0.51$, recovery rate $g = 1/3$; (a)(b)(d) $N = 1000$, $I_0 = 10$; (c) $N = 500$, $I_0 = 5$; (c)(d) $c = 0.1$, $\beta = 10$. The lines in (a) and (b) are visual guides.

replicator dynamics, converging to the Nash equilibrium. Strong selection, on the other hand, drives individuals to imitate successful free-riders based on a single observation, even when a rational agent with complete information would realize that attempted free-riding does poorly in expectation. This "paradox of imitation" is a very general phenomenon [267] and may in part explain cases where public vaccination levels are low. In particular, for the range of vaccination cost appropriate to influenza (i.e., $c \approx 0.002$ to 0.08, see supplementary materials), the imitation dynamics with strong selection in the well-mixed case falls well short of the rational optimum, leading to over-exploitation of herd immunity and an increase in preventable infections. Our model describes the admittedly extreme case in which each individual observes only one randomly chosen role model each round. Allowing imitators to learn from a somewhat larger group of peers could lessen the sampling error, but would not eliminate it.

This kind of error is reminiscent of, but distinct from, the phenomenon of "information cascades" that generate rationalized conformism or "groupthink" [18, 28]. Such cascades may also be obstacles to high vaccination coverage [21]. To explore conformism (or, alternatively, stubbornness) in the context of our model, one might include an additional cost τ of switching strategy in the thermal updating rule [297, 312]; that is, $f(\Delta P) = 1/[1 + \exp(\beta(\Delta P + \tau))]$. A large negative (positive) τ would then represent the tendency to copy one's peers (stick with the current strategy), regardless of payoff comparisons. Previous studies have shown in detail how this sort of payoff-neglecting imitation can lead to widespread conformism and adoption of sub-optimal strategies [18, 28].

It is widely known that population structure can promote the evolution of cooperative behavior [211, 214, 222, 223, 303, 304]. We have shown, however, that population structure is a "double-edged sword" for public health: It can promote high levels of voluntary vaccination and herd immunity, but small increases in cost beyond a certain threshold c_L cause vaccination to plummet – and infections to rise – more dramatically than in well-mixed populations. For example, the random network population under strong selection ($\beta = 10$) can prevent the epidemic completely for costs up to $c = 0.04$, but 11% of the population become infected at cost $c = 0.08$. In the well-mixed population, the epidemic grows gradually, from 8% to 15%, over the same cost range. This threshold effect is robust to changes in population structure and exists in lattice (figures 4.3.2a and 4.3.2b) and scale-free network (suppl. figures 4.5.7a and 4.5.7b) populations as well.

In social networks, individuals' degrees vary greatly, and highly-connected individuals (hubs) can spread disease to a large number of peers if infected. The vaccination of hubs can play a vital role in containing infections [237], and public health programs often try to promote herd immunity by allocating vaccinations preferentially to these hubs [19]. Physicians who are hubs in a disease-transmission network, for instance, have high rates of vaccine uptake [41]. Our model shows that even individuals with incomplete information can self-organize to achieve this pro-social outcome (figure 4.3.3). Since hubs generally face greater infection risk than small-degree individuals do, they have increased incentive to vaccinate; hubs' self-interest is therefore

relatively well-aligned with overall welfare.

Recent work with a detailed model designed to mimic a smallpox outbreak on a random network [242] reaches a complementary conclusion about the fragility of high-coverage equilibria: voluntary vaccination can contain a disease in low-degree networks, but as the average degree increases, the system reaches a critical threshold past which it behaves like a well-mixed population and the epidemic spreads. This work focused on vaccination decisions made during the course of an epidemic in response to local disease prevalence, as opposed to season-by-season updating of preemptive vaccination decisions. Taken together, our current work and this previous result demonstrate how local disease transmission and decision-making based on local context change the character of vaccination dynamics. Voluntary vaccination can be a viable policy for achieving high coverage and eradicating disease, but the final outcome is sensitive to small changes in (actual or perceived) vaccination cost and in the social network. This sensitivity may in part explain how anecdotal evidence of vaccine-related health risks has been able to trigger steep declines in coverage and loss of population immunity [23, 142, 208]. Policy levers that subsidize vaccination can take advantage of these threshold effects to promote disease containment and eradication.

Achieving socially optimal coverage through voluntary vaccination is a problem of cooperation with limited information and uncertainty about outcomes. The problem is similar to public goods games studied by economists [230], as herd immunity provides a communal benefit. Individuals' use of salient

anecdotes to cope with uncertainty, however, is not a typically studied feature of public goods games. Two sources of uncertainty face an individual deciding whether to vaccinate: uncertainty about contracting the infection if unvaccinated, and uncertainty regarding adverse reactions to the vaccine itself. Our current work focuses on the former uncertainty, treating the vaccine cost as a fixed quantity, which is a summary of all expected costs. It may also be instructive to treat vaccine cost as a random variable, as a way of explicitly modeling public fears concerning vaccine safety. These fears often have a tremendous impact on vaccine take-up and public health [77, 208].

4.5 SUPPLEMENTARY MATERIALS

4.5.1 EPIDEMIC MODELING

We use the Susceptible-Infected-Recovered (SIR) model for the disease transmission process. The SIR model is appropriate for a large class of infectious diseases such as influenza and measles, and is widely studied in epidemiology [154]. In this model, the population is divided to three classes: susceptible individuals (S), who are healthy but can catch the disease if exposed to infected individuals; infected individuals (I), who have the disease and can pass it on; and recovered individuals (R), who acquire immunity to the disease.

Suppose the disease transmission rate is r, and the rate of recovery from infection is g. The fraction of susceptible, infected, and recovered individuals is S, I, and R, respectively, in a population of size N. For well-mixed populations, the time evolution of the population states can be expressed as the following deterministic ordinary differential equations:

$$\frac{dS}{dt} = -rNSI, \qquad (4.4)$$

$$\frac{dI}{dt} = rNSI - gI, \qquad (4.5)$$

$$\frac{dR}{dt} = gI. \qquad (4.6)$$

The initial condition for an epidemic introduced by one infected individual is $S(0) = 1 - 1/N \approx 1$, $I(0) = 1/N$, and $R(0) = 0$. Denote rN/g by R_0, commonly called the "basic reproductive ratio" [119]. Here R_0 is the mean number of secondary infections caused by a single infected individual, during his/her entire infectious period, in a completely susceptible population.

Dividing Eq. (4.4) by Eq. (4.6), we obtain

$$\frac{dS}{dR} = -R_0 S. \qquad (4.7)$$

Integrating above equation from time 0 to ∞, we get the transcendental equation

for the final epidemic size $R(\infty)$:

$$S(\infty) = S(0)e^{-R_0[R(\infty)-R(0)]}. \tag{4.8}$$

Using the initial condition $S(0) \approx 1$ and $R(0) = 0$, and the final state $I(\infty) = 0$ and $S(\infty) = 1 - R(\infty)$, we obtain:

$$R(\infty) = 1 - e^{-R_0 R(\infty)}. \tag{4.9}$$

$R(\infty)$ is the final fraction of individuals who had been infected during the epidemic outbreak, *i.e.*, the final epidemic size, which can be calculated numerically from the above equation. Differentiating both sides of Eq. (4.9) with respect to $R(\infty)$, we can see that the final size is positive if and only if $R_0 > 1$. If $R_0 < 1$, the disease does not spread.

If we consider preemptive vaccination by supposing that a portion x of the population initially vaccinated, Eq. (4.9) can be rewritten as

$$R(\infty) = (1 - x)(1 - e^{-R_0 R(\infty)}). \tag{4.10}$$

Increasing vaccination decreases the final size of the epidemic, and if $x > x_h = 1 - \frac{1}{R_0}$, we have $R(\infty) = 0$. The critical value x_h is called the "herd immunity threshold," above which the infection does not spread through the population. For vaccine-preventable diseases, herd immunity therefore grants indirect protection to unvaccinated individuals; it is a public good that vaccinated
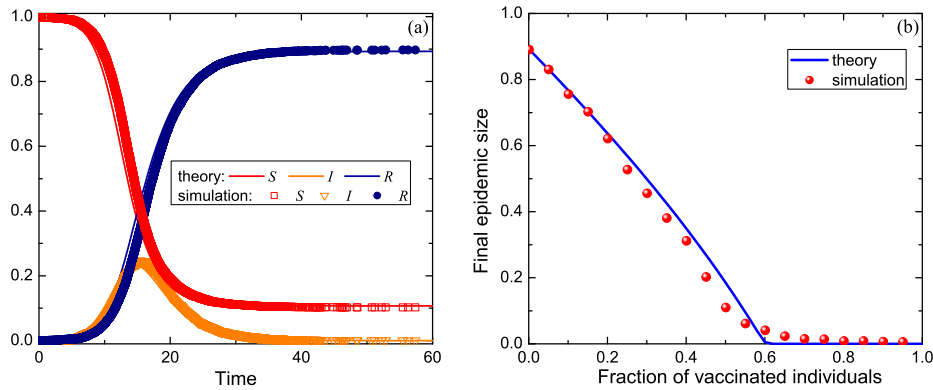
Figure 4.5.1: Epidemic spreading in well-mixed populations. (a) The fractions of susceptible, infected, and recovered individuals as a function of time. (b) The fraction of individuals who had been infected as a function of preemptive vaccination level. The solid line in panel (b) is numerically determined from Eq. (4.10). Parameters: (a) $N = 10^4$, $g = 1/3$, $R_o = 2.5$, the number of initial infection seeds $I_o = 10$; (b) $N = 1000$, $R_o = 2.5$, $I_o = 5$, results averaged over 50 runs.

individuals create and unvaccinated individuals can free-ride on.

We use the Gillespie algorithm in our stochastic simulations (detailed in section 4.5.1). To lower the chance that an epidemic outbreak fails merely due to stochastic effects, we make the initial number of infection seeds I_o more than one. Figure 4.5.1 shows the epidemic spreading in well-mixed populations. For $R_o = 2.5$, the final epidemic size with zero vaccination is ~ 0.893, and the herd immunity threshold is $x_h = 0.6$. Our simulation results agree with the deterministic model given by Eqs.(4.4)–(4.6) (figures 4.5.1a and 4.5.1b). Note that for intermediate initial fractions of vaccinated individuals ($0.2 < x < 0.6$), the final epidemic size resulting from simulations is lower that than the analytical prediction (Eq. 4.10). Stochastic effects due to finite infection size cause this

deviation.

EPIDEMIC SPREADING IN STRUCTURED POPULATIONS

It is not typically possible to derive explicit equations for epidemic spreading in structured populations [154], so we use stochastic simulations. It is widely accepted that population structure can substantially alter epidemiological dynamics from the well-mixed case [207]. To identify only the effects of population structure, we must calibrate epidemic parameters to ensure that infection risk is equal in all structures examined [242]. We use as the base case a well-mixed population with $R_0 = 2.5$, which is within the typical R_0 values for influenza ($1.5 < R_0 < 3$). We fix the recovery rate at $g = 1/3$ (rates may be interpreted as per-day, so the mean infectious period is 3 days), and choose the transmission rate r such that the final epidemic size is that of the well-mixed population without vaccination.

We simulated populations structured as square lattices, Erdős-Rényi random graphs [88], and Barabási-Albert scale-free networks [20]. To account for the increased risk that individuals with many connections face, we assume that the infection probability of a susceptible individual i is proportional to the number of her infected neighbors $N_I(i)$. The transition rate from S to I for individual i is then $rN_I(i)$ [154].

For lattice populations, the final epidemic size shows a clear phase transition from zero to one with increasing r values (figure 4.5.2a). For low r values, the epidemic spreading is inhibited due to local spatial clustering effects. We select

Figure 4.5.2: Epidemic spreading in lattice populations. (a) The final epidemic size is shown as a function of the transmission rate r with zero vaccination coverage. (b) The final epidemic size as a function of vaccination level (preemptive, random vaccination). The arrow notes where vaccination brings the final epidemic size below 0.2% (twice the size of initial inoculum), representing an approximate herd immunity threshold. Parameters: $N = 100 \times 100$ with von Neumann neighborhood (four adjacent neighbors), (a)-(b) $g = 1/3$, $I_0 = 10$, (b) $r = 0.46$. Results are averaged over 100 runs.

$r = 0.46$, which gives a final epidemic size of ~ 0.893, approximately equal to the base case. Using this transmission rate, we simulate the effect of preemptive, random vaccination on the epidemic (figure 4.5.2b). The final epidemic size decreases more precipitously than in the well-mixed case (cf. figures 4.5.1b and 4.5.2b). At vaccination levels greater than about 0.3, the disease cannot persist in the population (marked by the arrow in figure 4.5.2b).

Compared to spatial lattices, both the absence of local clustering and the presence of degree heterogeneity (different individuals can have different numbers of neighbors) in random graphs and scale-free networks make it easier for the disease to spread at lower transmission rates and higher vaccination rates (figures 4.5.3 and 4.5.4). Using the same method as above, we choose $r = 0.51$

Figure 4.5.3: Epidemic spreading in Erdős-Rényi random networks. (a) The final epidemic size as a function of the disease transmission rate r with zero vaccination coverage. (b) The final epidemic size as a function of the vaccination level (preemptive, random vaccination). The arrow notes where vaccination brings the final epidemic size below 2%, representing an approximate herd immunity threshold. Parameters: (a)-(b) $N = 1000$, average degree $\bar{k} = 4$, $I_0 = 10$, $g = 1/3$; (b) $r = 0.51$. Results are averaged over 100 runs.

for random graphs (figure 4.5.3a) and $r = 0.55$ for scale-free networks (figure 4.5.4a). Notice that even for vanishingly small r values, scale-free networks are fragile to epidemic attacks, consistent with previous findings [236]. Accordingly, the vaccination level needed to contain the disease is the highest among all the population structures we studied (figure 4.5.4).

STOCHASTIC SIMULATION PROCEDURE: GILLESPIE ALGORITHM

We use the Gillespie algorithm to simulate the epidemiological process [106]. The simulation procedure works as follows:

Step 1: At time t, calculate each susceptible and infected individual's transition rate, $p_i(t)$. The rate at which a susceptible individual becomes infected is

Figure 4.5.4: Epidemic spreading in Barabási-Albert scale-free networks. (a) The final epidemic size as a function of the disease transmission rate r with zero vaccination coverage. (b) The final epidemic size as a function of the vaccination level (preemptive, random vaccination). The arrow notes where vaccination brings the final epidemic size below 2%, representing an approximate herd immunity threshold. Parameters: (a)-(b) $N = 1000$, average degree $\bar{k} = 4$, $I_0 = 10$, $g = 1/3$; (b) $r = 0.55$. Results are averaged over 100 runs.

$p_i(t) = r \times$ number of infected neighbors. The rate at which an infected individual recovers from the disease is $p_i(t) = g$. The total transition rate is $\lambda(t) = \sum_i p_i(t)$.

Step 2: The time at which the next transition event occurs is $t' = t + \Delta t$, where Δt is sampled from an exponential distribution with mean $\frac{1}{\lambda(t)}$. (Generate a uniform random number $u \in [0, 1)$. Then the time interval is $\Delta t = -\frac{\ln(1-u)}{\lambda(t)}$.)

Step 3: Choose the individual whose state changes at time t' by sampling proportional to $p_i(t)$. Generate a uniform random number $v \in [0, 1)$. If $\sum_{j=1}^{k-1} p_j(t)/\lambda(t) < v < \sum_{j=1}^{k} p_j(t)/\lambda(t)$, then individual k is chosen to change state. (Define $\sum_{j=1}^{0} p_j(t)/\lambda(t) = 0$.)

Step 4: Repeat Steps 1–3 until the number of infected individuals $I(t)$ is zero, or stop after a predetermined time period.

Values of c are supported by data from [98], which estimates the cost of health outcomes: Vaccination costs \$37 on average, and the expected cost of infection for non-vaccinated individuals is given for four cases:

- Young individuals (< 65 years) during normal seasons: \$570

- Elderly individuals (≥ 65 years) during normal seasons: \$4,160

- Young individuals during pandemics (*e.g.,* 1918 influenza): \$21,220

- Elderly individuals during pandemics: \$16,170

To represent the imperfect effectiveness of vaccination (roughly 80% effectiveness for the young and 60% for the elderly), we scaled the vaccination cost by $\frac{1}{80\%} = 1.25$ and $\frac{1}{60\%} = 1.67$ for each group, respectively. The relative vaccination costs (as a fraction of infection costs) are then:

- Young, normal seasons: $c = 0.08$

- Elderly, normal seasons: $c = 0.01$

- Young, pandemics: $c = 0.002$

- Elderly, pandemics: $c = 0.004$

Based on these estimates, we propose that it is reasonable to use values of c in the range 0.001 to 0.1 in our model to discuss influenza.

4.5.3 Nash equilibrium and social optimum

Nash equilibrium

Here we demonstrate the existence and uniqueness of the Nash equilibrium (*i.e.*, individual rational optimum) in the vaccination game, provided that individuals have perfect knowledge of the vaccination coverage level and infection risk. Propose strategy x (representing the probability of vaccination) to be a Nash equilibrium: if most of the population plays strategy x, then individuals adopting a different strategy y can do no better than the resident. For an ε-size invasion ($\varepsilon \ll 1$), the new vaccination coverage is $p := x(1 - \varepsilon) + y\varepsilon$. The expected payoff to strategy y is then

$$E(y, p) = -yc + (1 - y)\left\{[1 - w(p)] \cdot 0 + w(p) \cdot (-1)\right\}, \qquad (4.11)$$

where $w(p)$ is the infection risk for an unvaccinated individual, given that a proportion p of the population is vaccinated. Strategy x is Nash if it is a best response to itself, which requires the conditions

$$\left.\frac{\partial E(y, p)}{\partial y}\right|_{y=x} = w(x) - c - \varepsilon(1 - x)w'(x) = 0, \qquad (4.12)$$

$$\left.\frac{\partial^2 E(y, p)}{\partial y^2}\right|_{y=x} = 2\varepsilon w'(x) - (1 - x)\varepsilon^2 w''(x) \leq 0. \qquad (4.13)$$

Note that $w(x)$ strictly decreases with x, until x reaches the herd immunity threshold x_h. For $x < x_h$, the inequality (4.13) is therefore strict for a sufficiently small invasion ε, and so higher-order conditions are not required. Also, for small

invasions, the ε term in Eq. (4.12) can be safely neglected. The vaccination cost falls into one of three ranges:

- *Case 1*, $0 < c \leq w(0)$. Since $w(x)$ strictly decreases, there is a unique x^* that solves $w(x^*) = c$. This value x^* is the Nash equilibrium.

- *Case 2*, $c > w(0)$. As the derivative in Eq. (4.12) is negative, the best response is x as small as possible; that is, the pure Nash equilibrium $x^* = 0$.

- *Case 3*, $c \leq 0$. As the derivative in Eq. (4.12) is positive, the best response is x as large as possible; that is, the pure Nash equilibrium $x^* = 1$.

Moreover, strictness of the inequality (4.13) in Case 1 implies that an alternative strategy $y \neq x^*$ does strictly worse, meaning that the Nash equilibrium is also evolutionarily stable $[132]$.

Furthermore, we can show that the unique Nash equilibrium x^* in this game is globally stable. For any proportion $\varepsilon \in (0, 1)$ of individuals playing strategy $y \neq x^*$, we always have

$$\Delta E = E(x^*, p) - E(y, p) = (x^* - y)\left[w(p) - c\right] > 0, \qquad (4.14)$$

which means that the strategy x^* is favored against any alternative strategy at any frequency.

CALCULATING THE NASH EQUILIBRIUM FROM EPIDEMIOLOGICAL PARAMETERS

For well-mixed populations, $w(x)$ is the ratio of the number of individuals who

acquired disease, $R(\infty)$, to the total number of susceptible individuals, $S(0)$. Therefore we obtain

$$w(x) = \frac{R(\infty)}{1 - x} = 1 - e^{-R_0 R(\infty)}. \tag{4.15}$$

Using Eq. (4.10) and the Nash condition $w(x^*) = c$, we have $e^{-R_0 R(\infty)} = 1 - c$ and $R(\infty) = (1 - x^*)c$. We then obtain the Nash Equilibrium

$$x^* = 1 + \frac{\ln(1 - c)}{cR_0}, \tag{4.16}$$

which is plotted in figure 4.5.5. This equation holds for $0 < c \le w(0)$ (recall, $w(0) = 1 - e^{-R_0 w(0)}$). Cases 2 and 3 above cover the alternatives.

SOCIAL OPTIMUM

The population's optimal vaccination coverage can be obtained by minimizing the total expected cost from both vaccination and infection. If a fraction x of the population is vaccinated, the expected cost is

$$\begin{aligned}
E(x) &= N\left\{xc + (1 - x)\left\{0 \cdot [1 - w(x)] + w(x) \cdot 1\right\}\right\}, \tag{4.17}\\
&= N\left[xc + R(\infty)\right], \tag{4.18}\\
&= N\left[xc + (1 - x)(1 - e^{-R_0 R(\infty)})\right].
\end{aligned}$$

We show that the social optimum is exactly the herd immunity threshold,

139

$x_h = 1 - 1/R_o$.

For x above x_h, the final epidemic size $R(\infty)$ is zero. $E(x)$ therefore increases as x rises above x_h.

For x below x_h, $\frac{dE(x)}{dx} = N\left[c + \frac{dR(\infty)}{dx}\right]$. It is easy to show that $\frac{dR(\infty)}{dx} < 0$; that is, that final infection size decreases with vaccination coverage. Furthermore, differentiating both sides of Eq. (4.10) with respect to x, we obtain:

$$\frac{dR(\infty)}{dx} = -\frac{e^{R_o R(\infty)} - 1}{e^{R_o R(\infty)} - (1 - x)R_o}, \qquad (4.19)$$

which is guaranteed to be less than -1 for $(1 - x)R_o > 1$; i.e., $x < x_h$. Hence we know that $\frac{dE(x)}{dx} = N\left[c + \frac{dR(\infty)}{dx}\right]$ is negative for $c < 1$ and $x < x_h$.

Since $E(x)$ decreases for $x < x_h$ and increases for $x > x_h$, the socially optimal vaccination level is precisely x_h.

For any $c \in (0, 1)$, the Nash equilibrium falls short of the social optimum, leading to the well-known dilemma of voluntary vaccination in a population of selfish, rational individuals.

EFFECTS OF R_o ON VOLUNTARY VACCINATION

For a fixed relative cost of vaccination, the Nash equilibrium increases with rising R_o (figure 4.5.5): given a higher risk of infection, rational individuals are more likely to vaccinate. In the limiting case $R_o \to \infty$, unvaccinated individuals cannot free-ride on the immunity generated by others, and so they eventually get infected. In this case, the Nash equilibrium and social optimum converge to 100% vaccination. For the opposite limiting case $R_o \leq 1$, individuals have zero
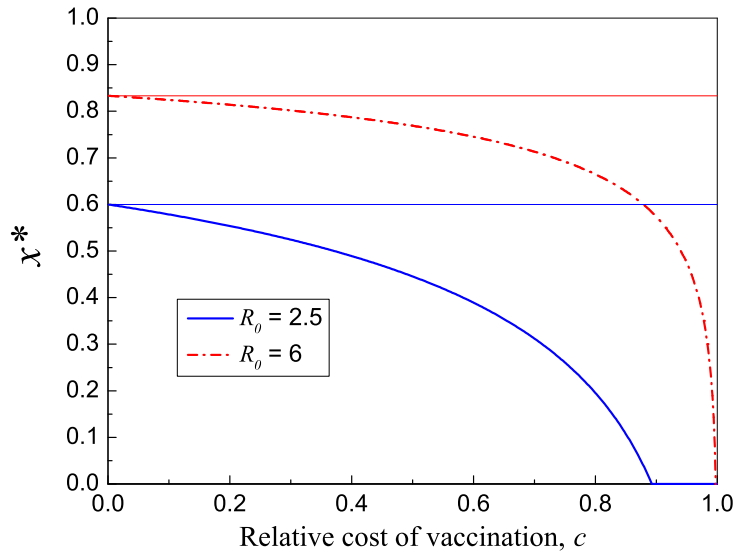
Figure 4.5.5: Nash equilibrium as a function of relative cost of vaccination cost c with different disease transmissibility R_0. The horizontal lines correspond to socially optimal vaccination levels.

risk of infection, so that the Nash equilibrium and social optimum again agree — no one is vaccinated.

4.5.4 EVOLUTION OF VACCINATING BEHAVIOR IN WELL-MIXED POPULATIONS

Above, we analyzed the vaccination dilemma from the perspective of classical game theory. Here we consider vaccination dynamics from an evolutionary game perspective. We derive a diffusion approximation for large populations of size N. Let m/N be the fraction of vaccinated individuals, who are immune from the seasonal infectious disease. Individuals imitate others based on the *pairwise comparison rule*, which preferentially copies others with higher payoffs [311]. Each round, a randomly chosen individual i selects another random individual j

as role model, and compares her own payoff to that of the role model. Individual i adopts the strategy of individual j with the probability given by the Fermi function

$$\varphi(s_i \leftarrow s_j) = f(P_j - P_i) = \frac{1}{1 + \exp[-\beta(P_j - P_i)]},\qquad(4.20)$$

where β represents the intensity of selection. The population can change only if individuals i and j have different strategies. Hence, the probability that the number of vaccinated individuals increases from m to $m + 1$ (denoted T_m^+) and the probability that the number decreases from m to $m - 1$ (denoted T_m^-) are

$$T_m^\pm = \frac{m}{N}\frac{N-m}{N}\left\{[1 - w(m/N)]\frac{1}{1 + e^{\mp\beta(P_A - P_{B_0})}} + w(m/N)\frac{1}{1 + e^{\mp\beta(P_A - P_{B_1})}}\right\},$$

$$(4.21)$$

where P_A is the payoff of vaccinated individuals, P_{B_0} the payoff of unvaccinated (and healthy) individuals, and P_{B_1} the payoff of unvaccinated (and infected) individuals (see Table 4.5.1).

For large populations [310], this process can be approximated by a stochastic differential equation with drift $T_m^+ - T_m^-$ and diffusion $\sqrt{(T_m^+ + T_m^-)/N}$, yielding

$$\dot{x} = x(1-x)\left\{[1 - w(x)]\tanh\left[\frac{\beta}{2}(P_A - P_{B_0})\right] + w(x)\tanh\left[\frac{\beta}{2}(P_A - P_{B_1})\right]\right\} + \sqrt{\frac{x(1-x)}{N}}\xi$$

$$(4.22)$$

where $x = m/N$ is the fraction of vaccinated individuals and ξ is Gaussian white noise with variance one. For $N \to \infty$, the stochastic term vanishes. As a result,

Table 4.5.1: The fraction of individuals with different states and their corresponding payoffs.

	vaccinated	unvaccinated and infected	unvaccinated and healthy
fraction	x	$R(\infty)$	$1 - x - R(\infty)$
payoff	$P_A = -c$	$P_{B_1} = -1$	$P_{B_0} = 0$

for large populations, we can use the deterministic approximation

$$\dot{x} = x(1-x)\left\{[1-w(x)]\tanh\left[\frac{\beta}{2}(-c-0)\right] + w(x)\tanh\left[\frac{\beta}{2}(-c+1)\right]\right\}.$$

$$(4.23)$$

At equilibrium where a fraction x is vaccinated, the fraction infected is expected to be $R(\infty)$ (as given in Eq. (4.10)), and the fraction that are successful free-riders (unvaccinated and healthy) is expected to be $1 - x - R(\infty)$ (see Table 4.5.1).

A. FOR SMALL β, we have $\tanh(\beta x) \sim \beta x$. Thus Eq. (4.23) simplifies to

$$\begin{aligned}
\dot{x} &= x(1-x)\left[-c[1-w(x)]\frac{\beta}{2} + w(x)\frac{\beta}{2}(1-c)\right] \\
&= \frac{\beta}{2}x(1-x)[w(x)-c].
\end{aligned}$$

$$(4.24)$$

The replicator dynamics is recovered in this limit, $\beta \ll 1$, but with the time scale adjusted by a factor $\frac{\beta}{2}$. For any vaccination cost $0 < c < w(0)$, the system converges to the interior equilibrium $x^* = 1 + \ln(1-c)/(cR_0)$, which is evolutionarily stable as remarked above.

B. For large β:

B1. For $c \to 0$ $(c < 1/\beta)$, $-\frac{\beta c}{2} \to 0$,   Eq. (4.23) becomes

$$\dot{x} = x(1-x)\left[-\frac{\beta}{2}c[1-w(x)] + w(x)\right],\tag{4.25}$$

which has a stable interior equilibrium $x^* = 1 - (1 + \frac{\beta c}{2})\ln(1 + \frac{\beta c}{2})/(\frac{\beta c}{2}R_0)$. For small c, the first-order approximation of this expression is $1 - \frac{1+\frac{\beta c}{4}}{R_0}$. Comparing this approximation to the Nash

equilibrium $(x^* = 1 + \ln(1-c)/(cR_0) \approx 1 - \frac{1+c/2}{R_0})$, we note that the effect of large β can be described as rescaling small values of c by a factor of $\frac{\beta}{2}$.

B2. $c \to 1$ $(c > 1 - 1/\beta)$, $-\frac{\beta(1-c)}{2} \to 0$,   Eq. (4.23) becomes

$$\dot{x} = x(1-x)\left\{-[1-w(x)] + \frac{\beta}{2}(1-c)w(x)\right\}.\tag{4.26}$$

The third factor in Eq. (4.26) equals zero for vaccination level $\tilde{x} := 1 - \frac{(2+\beta(1-c))\ln\left(\frac{2+\beta(1-c)}{\beta(1-c)}\right)}{2R_0}$. If this value is positive, then the stable interior equilibrium is $x^* = \tilde{x}$; otherwise, $x^* = 0$.

B3. For intermediate c $(1/\beta < c < 1 - 1/\beta)$,   the vaccination level Eq. (4.23) depends little on β and can be approximated as

$$\dot{x} = x(1-x)[2w(x) - 1],\tag{4.27}$$

144

which has a stable interior equilibrium $x^* = 1 - \frac{2\ln 2}{R_0}$. Therefore the vaccination level has a plateau at $1 - \frac{2\ln 2}{R_0}$ for large β and intermediate c values.

Notice that at $c = 0.5$, $x^* = 1 - \frac{2\ln 2}{R_0}$ is an equilibrium for any β value (figure 4.5.6).

Figure 4.5.6 shows the effects of selection strength β and R_0 on equilibrium vaccination coverage (figures 4.5.6a versus 4.5.6c) and final epidemic size (figures 4.5.6b versus 4.5.6d). We can see that for low β the equilibrium vaccination level under imitation dynamics converges to the Nash Equilibrium. Strong selection (large β values) causes the vaccination level to drop below the Nash equilibrium when vaccination cost is low. Furthermore, greater risk of infection (higher R_0) does prompt higher levels of vaccination among imitating individuals, shrinking the gap between the utilitarian optimum and the voluntary outcome (cf. figures 4.5.6a and 4.5.6c).

4.5.5 Vaccination dynamics on scale-free networks

In addition to random graphs reported in the main text, we consider vaccination dynamics on scale-free networks. The degree distribution of real-life social networks follows a power law, which can be represented using a Barabási-Albert scale-free network model [20]. Scale-free networks generally possess larger degree heterogeneity than random graphs, leading to more severe persistence of epidemic outbreaks, making herd immunity more difficult to achieve (cf. figures 4.5.3 and 4.5.4). As a consequence, network heterogeneity further promotes individuals' vaccination on scale-free networks (figure 4.5.7). The range of c that

Figure 4.5.6: Imitation dynamics in a large, well-mixed population (diffusion approximation). Upper panels (a), (c) show the equilibrium vaccination level, as a function of relative cost of vaccination c with different intensities of selection β. Lower panels (b), (d) show the final epidemic size, as a function of relative cost of vaccination c with different intensities of selection β. Parameters: (a)(b) $R_o = 2.5$, (c)(d) $R_o = 6$.

promotes 100% vaccination is larger than in the case of random graphs (for $\beta = 1$: $c < 0.3$ for scale-free versus $c < 0.2$ for random; for $\beta = 10$: $c < 0.1$ versus $c < 0.05$).

Despite this difference, the overall pattern of equilibrium vaccination coverage in scale-free networks is similar to that of random networks: "hubs" with many neighbors tend to vaccinate more often than small-degree individuals do (figure 4.5.7d). The particular structural characteristics of the Barabási-Albert scale-free network model seem to complicate this pattern slightly, in that the most likely free-riders actually have intermediate degree ($k = 3, 4, 5$ in figure 4.5.7d) rather than lowest degree ($k = 2$). Since many degree-two nodes are connected only to large hubs, their vaccination decisions are determined by imitation of these hubs. This peer influence appears to outweigh the fact that they can easily free-ride on the hubs' immunity, increasing the vaccination frequency of degree-two nodes above that of slightly better-connected individuals.

Although degree heterogeneity promotes vaccination, the equilibrium vaccination coverage is still sensitive to the cost of vaccination. Above a critical cost, the vaccination coverage rapidly falls below both the herd immunity threshold (figure 4.5.7a) and the final size of the epidemic grows (figure 4.5.7b). For influenza, the estimated relative cost of vaccination to infection is less than 0.1, which is approximately the threshold found in the $\beta = 10$ case. Misperceived vaccination risks and individual variation in attitudes towards vaccination may, however, tip the effective value of c above this threshold.
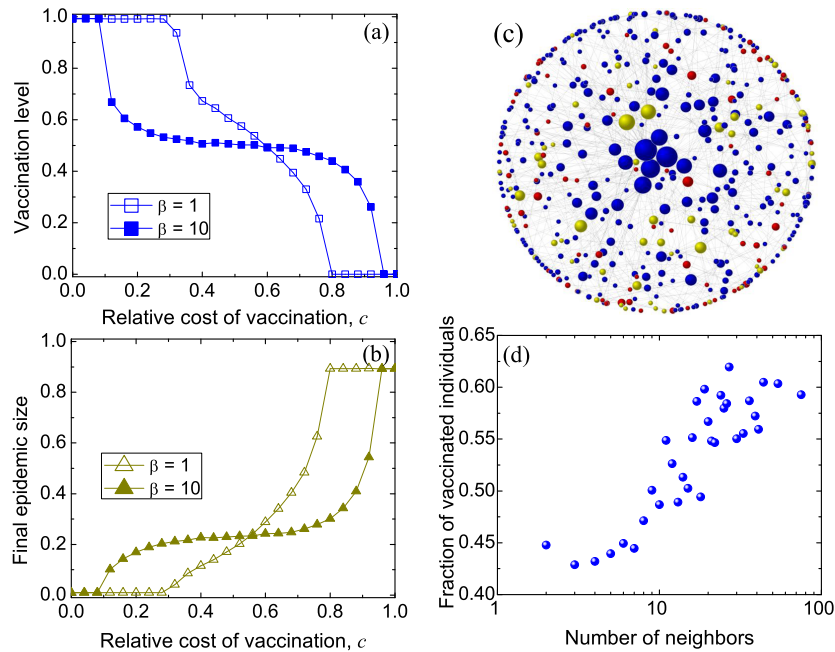
Figure 4.5.7: Vaccination dynamics on scale-free networks. Left panels show the fractions of (a) vaccinated and (b) infected individuals as a function of relative cost of vaccination c with the intensity of selection $\beta = 1$ and 10. Right panels: (c) Snapshot of a single simulation on a scale-free network. The size of a node corresponds to its degree. Blue nodes are vaccinated, yellow are infected, and red are successful free-riders. (d) The frequency of vaccination on a scale-free network, as a function of the number of social contacts an individual has (node degree). Parameters: (a)–(d)average degree $\bar{k} = 4$, disease transmission rate $r = 0.55$, recovery rate $g = 1/3$, $I_0 = 10$; (a)(b)(d) $N = 1000$, (c) $N = 500$; (c)(d) $c = 0.2$, $\beta = 10$. Results in panels (a), (b), and (d) are averaged over 100 runs. The lines in (a) and (b) are visual guides.

4.6 Acknowledgements

4.7 Manuscript information

4.7.1 Previously published as

This manuscript appeared in [96]:

Fu F, Rosenbloom DI, Wang L, Nowak MA (2011). Imitation dynamics of vaccination behaviour on social networks. *Proc. R. Soc. B*, 278(1702): 42 – 49.

5

Frequency-dependent selection can lead to evolution of high mutation rates

5.1   INTRODUCTION

IN POPULATIONS THAT ARE WELL-ADAPTED TO A STATIC ENVIRONMENT, most mutations are neutral or deleterious, and we may generally expect mutation rates to evolve to levels as low as feasible, given physical constraints and costs associated with faithful DNA replication [163, 174, 175, 295]. However, theoretical [69, 85, 141, 158, 174, 226, 299, 313, 326] and

empirical [107, 260, 294, 321] findings indicate that in novel or rapidly changing environments, evolution may select for higher mutation rates. While strains with higher mutation rates will experience increased mutational load, they also stand a better chance of giving rise to beneficial alleles, upon which the linked genes coding for increased mutation may "hitchhike" to fixation. High-rate mutator alleles can have a competitive advantage if beneficial mutations are strong and frequent enough to outweigh deleterious load [7, 69, 105, 307, 326].

Most studies that examine the role of environmental change in mutation rate evolution have considered exogenous fluctuations in the environment [141, 163, 174]. More recently, researchers have studied how host-parasite coevolution may generate an environmental feedback loop that drives mutation rates upward [150, 201, 227]. Yet the possibility that mutation rate evolution may be influenced by evolutionary dynamics within a single species remains largely unexplored [6]. Frequency-dependent competition, by which the composition of the population determines the fitness of phenotypes [120], can generate complex, unstable trait dynamics [132, 134, 192, 194, 218], which may in turn generate selection pressures on the mutation rate. Frequency dependence is common in nature [288] and occurs in a wide range of interactions: predator-prey systems [60, 169, 192], host-parasite systems [193, 274], niche competition [42, 177], cooperative dilemmas [73, 108, 109, 212, 251], and nontransitive competition dynamics [157, 165, 289, 290].

The evolution of high mutation rates may be an important force determining

the fate of obligately asexual lineages. When alleles increasing the genomic mutation rate ("global mutator alleles") are fully linked to the beneficial mutations that they promote (as in an ideal asexual population), the mutation rate may quickly rise beyond a maximum tolerable level, dooming the population to extinction within as few as several thousand generations [105]. This catastrophe occurs due to the myopia of natural selection, which can favor a mutator for its short-term beneficial consequences, despite the increased mutational load that eventually results as a strain bearing a mutator allele converges to mutation-selection equilibrium [7, 105]. Unlike Muller's ratchet and the related "mutational meltdown" phenomenon [115, 184], this effect can burden even a large population, with many available beneficial mutations, with a maximum possible "lifespan." The dynamics of mutation rate evolution in large populations with little or no recombination should therefore have broad implications for early microbial evolution, particularly regarding the evolution of sex.

Here we show how frequency-dependent dynamics can promote evolution of high mutation rates, both in the presence and absence of recombination. Our framework combines methods of evolutionary game theory [61, 132, 134, 195, 196, 218, 266, 320] and adaptive dynamics [70, 103, 131, 199, 216] to find and understand *evolutionarily stable mutation rates* (ESMRs, 141). We previously introduced our mathematical framework and used it to show that if trait frequencies converge to a stable equilibrium, then mutation rates evolve downward [6]. This result extends the classical observation that mutation rates evolve to zero in constant fitness

landscapes, if no new beneficial mutations are available [163, 174, 175]. Outside this special case, it is not known how mutation rates evolve under frequency-dependence, or even if they tend toward a single value. We therefore analyze and numerically simulate competition between strains of different mutation rates, in which the traits subject to mutation follow "rock-paper-scissors" dynamics. In most cases, the evolutionary dynamics converge upon a unique positive ESMR that depends on the rate of trait substitutions. The emergence of a unique ESMR is perhaps surprising against the backdrop of trait frequencies that are continually in flux. If mutation rates evolve in large steps, as can occur through damage to the mismatch-repair system [65, 191], we find that bistable outcomes may result: a high and a low rate may each resist invasion by the other. Recombination lowers the ESMR, but it also allows for stable mutation rate polymorphism, occurring between mutation rates above and below the ESMR.

The sort of cyclical dynamics that promote evolution of high mutation rates may be common in many species [157, 165, 288–292]. These dynamics involve periodic or chaotic transitions between population trait distributions, causing the favored phenotype(s) to change over time. Each transition provides another opportunity for a strain with high mutation rate to seize upon a beneficial mutation, thus giving it a fitness advantage over other strains. This scenario, in which a steady supply of beneficial substitutions can select for high mutation rates, is similar to those contemplated by other authors investigating constant selection [105, 307, 326]. We primarily study this scenario in an idealized infinite

population, using replicator-mutator equations to track the expected population trajectory. Strains with relatively low mutation rates face a disadvantage during each substitution event, as it takes longer for them to produce the optimal trait at high frequency. This idealized framework contrasts with finite-population models in which each strain must wait for a beneficial mutation to arrive stochastically, after which the mutator allele associated with that beneficial mutation may hitchhike to fixation [307, 326]. To highlight the relevance of our infinite-population model to finite populations, we demonstrate that both the ESMR and the invasion dynamics of mutator lineages in the infinite-population model coincide with a model of a small population ($N = 100$), in the regime where the ESMR greatly exceeds $1/N$ (implying that population waiting times for beneficial mutations are minimal).

Recombination inhibits the evolution of high mutation rates by separating the genetic determinants of the mutator phenotype from the beneficial mutations it produces [147, 175, 295, 308]. To analyze the effect of recombination under frequency-dependent selection, we introduce the *replicator-mutator-recombinator equations*. For a wide range of plausible recombination rates and frequencies of unconditionally deleterious mutations, the model supports local mutator alleles with ESMR values far above rates in the $10^{-10} - 10^{-8}$ range that is typical for cellular division in most species. In the Discussion, we outline a diverse set of genetic mechanisms that may meet our model's criteria for producing favored local hypermutation. We also find that the ESMR concept supports global hypermutation in asexual populations, as mutator hitchhiking is not disrupted by

recombination. More generally, while sufficiently strong recombination can drive the ESMR to zero, there still exist, for any frequency of recombination less than unity, evolutionary games that select for positive mutation rates. This finding bears similarity to the result of Ishii, et al. (1989) that for strongly beneficial mutations, recombination decreases the ESMR, but does not drive it to zero. In the rock-paper-scissors game we discover an elegant rule for recombination: evolution favors positive mutation rates if the selective advantage of beneficial mutants exceeds the ratio of recombining to non-recombining offspring. This rule holds regardless of the frequency of unconditionally deleterious mutation.

5.2   Model

We explore the evolution of mutation rates with a two-locus, many-allele model: the *trait locus* encodes the phenotype under frequency-dependent selection, while the *modifier locus* controls the rate at which the trait locus mutates. For simplicity, we assume that mutation on the modifier locus occurs on a slower time-scale than trait mutation, which puts aside the extinction risk presented by rapid evolution of mutation rates [78, 105]. Appealing to this separation of time-scales, we focus on competition between just two mutation rates at a time, following the approach of adaptive dynamics [70, 131, 199, 216]. We call a collection of individuals with the same modifier allele a *strain*. When a small strain of mutation rate u' (the "invader") appears in a population of mutation rate u (the "resident"), these two rates act as "meta-traits," as each strain can consist of different trait distributions. Competition between these mutation rates emerges

from the dynamics of the resident and invader trait distributions.

5.2.1   SHORT TIME-SCALE: TRAIT FREQUENCY DYNAMICS

Frequency-dependent competition is studied using evolutionary game theory [132, 134, 195, 196, 218], in which the reproductive fitness of an individual depends on the average of fitness payoffs obtained from interacting with other individuals according to a game. The distribution of traits found in the population evolves over time through the processes of interaction, reproduction, mutation, and death.

A game is characterized by a nonnegative payoff matrix A, where A_{ij} is the fitness that an individual of trait i obtains by interacting with an individual of trait j. Given a randomly interacting population with trait frequencies described by the nonnegative column vector $\mathbf{x} = (x_1, \ldots, x_n)^{\mathrm{T}}$, the fitness of trait i is $(A\mathbf{x})_i$. This definition of fitness ignores random deviations from the total population trait distribution in the subset that an individual meets for interaction. Note that the special case where each trait has a constant (frequency-independent) fitness is recovered by setting all columns of the matrix A equal. We assume reproduction is proportional to fitness and death occurs at the same rate for all traits.

The replicator equations describe the deterministic, infinite-population model without mutation, approximating the situation where all phenotypes are typically present in the population, though possibly at very low frequency [270, 305]:

$$\dot{r}_i = r_i(A\mathbf{r})_i - \varphi r_i, \tag{5.1}$$

Fig. 5.2.1(A) depicts one possible trajectory of the replicator equations where payoff matrix A describes a rock-paper-scissors game (formally defined below, Eq. (5.8)). This is the simplest model of cyclical population dynamics.

The replicator-mutator equations modify this basic model by assuming that each offspring adopts the trait of its parent with probability $1 - u$ and otherwise chooses from among all traits (including that of its parent) with equal probability [114, 128, 221]:

$$\dot{x}_i = (1 - u)x_i(A\mathbf{x})_i + \frac{u}{n}\varphi - \varphi x_i, \qquad (5.2)$$

Fig. 5.2.1(B) depicts one possible trajectory of the replicator-mutator equations with the same cyclical payoff matrix. Compared to the case without mutation, the population in (B) is always highly polymorphic. Mutation is only one possible mechanism that can preserve trait coexistence in rock-paper-scissors interactions; spatial population structure [83, 157] and alterations in the payoff matrix [132] also promote coexistence.

5.2.2   COMPETITION BETWEEN MUTATION RATES

To extend this model to competition between two mutation rates, we describe the system state by two nonnegative vectors $\mathbf{r} = (r_1, \ldots, r_n)$ and $\mathbf{z} = (z_1, \ldots, z_n)$, giving the relative abundances (with respect to the whole population) of residents and invaders of each trait. We denote the total relative

Figure 5.2.1: Example trajectories depicting evolution of the population according to (A) the replicator equations (5.1) and (B) the replicator-mutator equations (5.2) with $n = 3$ traits. Payoff matrix A determines a rock-paper-scissors game (described in Eq. (5.8)). Each trajectory is drawn on the "population simplex," where each point represents the trait distribution of the population at a particular time: points close to a vertex represent populations consisting mostly of one trait, while points close to the center of the simplex represent populations with a nearly equal mixture of traits. The trajectory in panel (A) starts near the center of the simplex and proceeds for 50 timesteps, spiraling outward and converging to the simplex boundaries. Once near the simplex boundaries, the population spends most of the time consisting mostly of one trait, and it transitions infrequently to the next trait in the cycle. The trajectory in panel (B) starts at the lower-left vertex, representing a population of 100% rock, and proceeds for 100 timesteps, spiraling inward and converging to a limit cycle. Once on this cycle, the population is always highly polymorphic, and the trait frequencies oscillate periodically. Comparing this limiting behavior to that of panel (A) shows that mutation acts to maintain trait diversity. The light blue vectors in both panels show the action of selection, directing the population counterclockwise around the simplex as determined by the game. The short red vector in panel (B) shows the action of mutation, pointing toward the center of the simplex, generating increased trait diversity. The trajectory determined by the replicator equations follows the selection vectors (A). The trajectory determined by the replicator-mutator equations follows the sum of the selection and mutation vectors — the resultant vector, in purple (B). Thin gray lines are axis guides meeting at the center of the simplex. Parameters: Fitness benefit for winning $a = 0.5$ (both panels); mutation rate $u = 0.07$ (panel B only).

abundances of residents and invaders by R and Z respectively,

$$R = \sum_{i=1}^{n} r_i, \qquad Z = \sum_{i=1}^{n} z_i.$$

We note that at every point in time, $R + Z = 1$. These assumptions lead to the dynamical equations

$$
\begin{aligned}
\dot{r}_i &= (1 - u)r_i(A(\mathbf{r} + \mathbf{z}))_i + \frac{u}{n}R\varphi_r - \varphi r_i, \\
\dot{z}_i &= (1 - u')z_i(A(\mathbf{r} + \mathbf{z}))_i + \frac{u'}{n}Z\varphi_z - \varphi z_i,
\end{aligned}
\tag{5.3}
$$

where

$$
\begin{aligned}
\varphi &= (\mathbf{r} + \mathbf{z})^T A(\mathbf{r} + \mathbf{z}), \\
\varphi_r &= \frac{1}{R}\mathbf{r}^T A(\mathbf{r} + \mathbf{z}), \\
\varphi_z &= \frac{1}{Z}\mathbf{z}^T A(\mathbf{r} + \mathbf{z})
\end{aligned}
\tag{5.4}
$$

are the average fitnesses of the whole population and of the two strains, respectively.

Initial success of the invading strain when it first appears in the population is measured by its *invasion fitness* $s_u(u')$. This value is defined as the time-averaged exponential growth rate of the frequency of the invading strain when rare [6]. Taking the time-average is required in this setting, as growth rate at each point in time depends on the current trait distributions, which themselves evolve. This

concept is formalized as

$$s_u(u') = \lim_{T \to \infty} \frac{1}{T} \int_0^T \frac{\dot{Z}(t)}{Z(t)} \, dt, \tag{5.5}$$

or equivalently,

$$s_u(u') = \lim_{T \to \infty} \frac{1}{T} \int_0^T \left(\varphi_z - \varphi \right) \, dt. \tag{5.6}$$

These definitions coincide with other notions of fitness as time-averaged exponential growth rate [198, 252]. As a practical matter, the dynamics (5.3) can be used to compute positive invasion fitness values only if the limiting value of T used in equations (5.5), (5.6) is not so large as to allow Z to reach a substantial frequency. In Eq. (5.20) in Methods we describe modified dynamics that keep the invader forever rare, avoiding this complication.

Replicator-mutator dynamics can yield a wide variety of behaviors, including multiple attractors, limit cycles, and chaos [132, 194, 268]. Invasion fitness is therefore difficult to compute in general and may depend on initial trait distributions. In certain cases, the limits taken in Eqs. (5.5), (5.6) do not even exist [282]. To minimize these complications, we focus on the simplest type of payoff matrix A that supports the evolution of high mutation rates: n-trait rock-paper-scissors interactions where no stable equilibrium trait distribution exists. This form of competition is simple enough that the dynamics have global, nonchaotic attractors; invasion fitness is readily computed; and ESMRs can be approximated analytically.

Many researchers have observed that the evolution of high mutation rates is particularly sensitive to population size effects, as mutator alleles must often wait long periods before giving rise to a beneficial mutation on which they can hitchhike, and they may drift to extinction before that time comes [7, 69, 299, 307, 321, 326]. Although invasion fitness and the ESMR are here defined with respect to a deterministic model in which drift plays no role, the same concepts can be recast in a stochastic model; to do so, we impute an effective invasion fitness from the fixation probability in simulation of a small population ($N = 100$) (see Methods).

5.2.3 Long time-scale: adaptive evolution of mutation rates

Strain u' is favored to invade strain u if $s_u(u') > 0$. In all cases examined in this article, if the invading rate u' is sufficiently close to u, an initially successful invasion leads to the eventual fixation of u'. This principle, central to adaptive dynamics theory, is known as "invasion implies substitution" and has been proven for simpler models of evolving quantitative characters [67, Appendix B]. In order to understand the long-term behavior of a gradually evolving trait that adheres to this principle, it suffices to analyze invasion fitness alone. Curiously, this principle *fails* for mutation rate evolution in general, and even a simple two-trait frequency-dependent interaction can support the coexistence of two arbitrarily close mutation rates [6].

An *evolutionarily stable mutation rate* (ESMR) $\tilde{u}$ is defined as one that is favored against invasion by any other mutation rate; that is, for all $u' \neq \tilde{u}$, it is the case

that $s_{\tilde{u}}(u') < 0$. While it is possible for multiple ESMRs to exist, all examples that we explore have no more than one ESMR. When it exists, the ESMR can be thought of as the unique mutation rate favored by evolution [6].

5.2.4 UNCONDITIONALLY DELETERIOUS MUTATION

The two-strain replicator-mutator equations (5.3) treat the modifier locus as determining the mutation rate on a single trait locus. In an actual genetic setting, a mutator allele may target other genomic regions to cause unconditionally deleterious mutations. This effect is known to weaken — or eliminate — selection for high mutation rates [85, 295, 313]. To capture the aggregate effect of a modifier determining a mutation rate u on other parts of the genome, we subtract a mutational load $u\ell$ from the fitness of each individual in the strain, representing the cost that the strain bears at mutation-selection equilibrium [162]. To ensure nonnegative fitness, we also add a sufficiently large baseline value f_0 to all fitnesses. Here, ℓ represents the size of the region targeted by the modifier relative to the size of the frequency-dependent trait locus. We assume that no beneficial mutations occur at this target in the time-scale under consideration. We also assume that deleterious mutations are strong enough so that the population's approach to mutation-selection equilibrium occurs more rapidly than trait substitution does; otherwise, the full cost $u\ell$ would not be realized [69, 105, 146].

5.2.5 RECOMBINATION

Recombination between the mutation-controlling locus and the trait locus is also known to weaken or eliminate selection for high mutation rates [85, 141, 295, 313]. Recombination can be modeled using the following simple scheme: After adults interact and reproduce proportional to fitness, the (possibly mutant) offspring pair randomly with one another. Each pair exchanges traits with probability c. The offspring then disperse and join the adult population. This lifecycle implies the following dynamics, which we dub the *replicator-mutator-recombinator equations*:

$$
\begin{aligned}
\dot{r}_i &= \left(1 - c + c\frac{R\varphi_r}{\varphi}\right)\left[(1-u)f_{i,u}r_i + \left(\frac{u}{n}\right)R\varphi_r\right] \\
&\quad + \left(c\frac{R\varphi_r}{\varphi}\right)\left[(1-u')f_{i,u'}z_i + \left(\frac{u'}{n}\right)Z\varphi_z\right] \\
&\quad - \varphi r_i,
\end{aligned}
$$

$$ (5.7) $$

$$
\begin{aligned}
\dot{z}_i &= \left(1 - c + c\frac{Z\varphi_z}{\varphi}\right)\left[(1-u')f_{i,u'}z_i + \left(\frac{u'}{n}\right)Z\varphi_z\right] \\
&\quad + \left(c\frac{Z\varphi_z}{\varphi}\right)\left[(1-u)f_{i,u}r_i + \left(\frac{u}{n}\right)R\varphi_r\right] \\
&\quad - \varphi z_i,
\end{aligned}
$$

where $f_{i,u}$ is the fitness of trait i subject to mutation rate u, $f_{i,u} = A(\mathbf{r} + \mathbf{z})_i + f_o - u\ell$. The first two lines of each equation provide the contribution from non-recombinant and recombinant individuals, respectively.

5.3 Results

5.3.1 Cyclical population dynamics can favor high mutation

Rock-paper-scissors games are a simple class of interactions that can exhibit cyclical replicator-mutator dynamics [132]. We study an n-trait generalization of the game $(n \geq 3)$, with $n \times n$ payoff matrix A given by

$$
\begin{pmatrix}
1 & 0 & 0 & \cdots & 0 & 1+a \\
1+a & 1 & 0 & \cdots & 0 & 0 \\
0 & 1+a & 1 & \cdots & 0 & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & 1 & 0 \\
0 & 0 & 0 & \cdots & 1+a & 1
\end{pmatrix} .
\tag{5.8}
$$

For each trait, there is another trait that gets a fitness benefit $a > 0$ from it (that is, there is a "scissors" for every "paper"), while all other traits do worse. The traditional rock-paper-scissors game corresponds to $n = 3$. Selection causes the population to cycle through successive traits, while mutation increases polymorphism in the population (Fig. 5.2.1). If all traits are present in equal proportion (that is, at the frequency vector $\left(\frac{1}{n}, \ldots, \frac{1}{n}\right)$), neither selection nor mutation causes the trait distribution to change. In the absence of mutation (using Eq. (5.1)), this fixed point is stable if $a > n - 2$. We focus on the case $a < n - 2$, which guarantees that a nonmutating population will produce unstable cyclical dynamics: any population that starts as a mix of all traits (but

not at the fixed point), converges to a heteroclinic orbit [129], meaning that it cycles between nearly-monomorphic states in order, as each trait is displaced by the one that defeats it [132, 136]. Long-term average population fitness approaches 1 in this case since, as the system converges to the heteroclinic orbit, it spends progressively more time near each monomorphic state, with transitions between them taking relatively little time. Increasing mutation draws this cyclical trajectory closer to the fixed point (i.e., towards greater polymorphism), causing long-term average fitness to fall. As the mutation rate crosses a threshold value of

$$u_\Omega = \frac{1 + (1 + a) \cos\left(\frac{2\pi}{n}\right)}{3 + a + (1 + a) \cos\left(\frac{2\pi}{n}\right)}, \tag{5.9}$$

the interior fixed point becomes a stable equilibrium (calculation in Online Appendix B), at which each trait has fitness $\frac{2+a}{n}$.

Mutator invasion of nonmutators

Before carrying out any numerical simulations of mutation rate competition, we show that the zero mutation rate is evolutionarily unstable in the absence of recombination and unconditionally deleterious mutations when $a < n - 2$; that is, there always exists some positive mutation rate u' that can invade a resident strain of nonmutators.

As noted above, the resident strain exists for long periods of time effectively consisting of a single trait, which has fitness $\varphi_r = 1$, given by the diagonal entries of payoff matrix A. While the invading strain is of negligible size, trait fitnesses are therefore constant: 1 for the resident trait, $1 + a$ for the favored trait, and 0 for all

other traits. The dynamics of the invading strain, given by (5.3), can then be rewritten in matrix form as

$$\dot{\mathbf{z}} = \left(Q_{u'}F - I\right)\mathbf{z},\tag{5.10}$$

where I is the identity matrix, $Q_{u'}$ is the mutation matrix

$$Q_{u'} = \begin{pmatrix} 1 - \frac{n-1}{n}u' & \frac{u'}{n} & \cdots & \frac{u'}{n} \\ \frac{u'}{n} & 1 - \frac{n-1}{n}u' & \cdots & \frac{u'}{n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{u'}{n} & \frac{u'}{n} & \cdots & 1 - \frac{n-1}{n}u' \end{pmatrix},\tag{5.11}$$

and F is a diagonal matrix with f_i, the fitness of the ith trait, as the ith diagonal entry.

Based on this dynamic, invasion fitness $s_0(u')$, defined by (5.5), is equal to the largest eigenvalue of the product matrix $Q_{u'}F$, minus 1. Since this matrix has only two nonzero columns, the condition for invader success, $s_0(u') > 0$, can be computed directly as

$$0 < u' < \frac{a\left(n-1\right)}{\left(1+a\right)\left(n-2\right)}.\tag{5.12}$$

In the limit of large n, most mutations are deleterious. Mutant offspring face a near-certain chance of mutating to a trait of zero fitness, while the nonmutator strain faces no such risk. Despite this large cost of mutation, mutators with

mutation rate less than $\frac{a}{1+a}$ can invade nonmutators for any combination of a and n. This condition can be derived by taking the limit as $n \to \infty$ of Eq. (5.12), or alternatively, by recognizing that the mutating strain incurs a fitness cost equal to the mutation rate u', but attains fitness advantage $\frac{a}{1+a}$ versus the resident (relative to the fitness of the optimal trait). The mutator can invade whenever this benefit exceeds the cost.

The analysis is similar if recombination c and load ℓ are introduced, again assuming $a < n - 2$. The equation describing invader dynamics, analogous to (5.10), is now

$$\dot{\mathbf{z}} = \left(CQ_{u'}F_{f_o} - I\left(1 + f_o + u'\ell\right) \right) \mathbf{z}, \qquad (5.13)$$

where $Q_{u'}$ is as above, F_{f_o} is a diagonal fitness matrix where the baseline fitness f_o is included, and C is a matrix whose ijth entry gives the probability that an invader of trait j will have trait i after recombination,

$$c_{ij} = \begin{cases} 1 - c + cr_i & \text{for } i = j, \\ \\ cr_i & \text{otherwise.} \end{cases}$$

In the present scenario, recall that $r_1 = 1$ and $r_i = 0$ for all $i > 1$. Invasion now succeeds if and only if the largest eigenvalue of $CQ_{u'}F_{f_o}$ exceeds $1 + f_o + u'\ell$. In

the case $f_0 = 0$, this condition simplifies to

$$0 < u' < \frac{a(1-c)(\mathcal{M}+1) - \mathcal{M}c}{a(1-c)(\mathcal{M}-\ell) + \mathcal{M}(1-c+\ell)}, \qquad (5.14)$$

where $\mathcal{M} = n\ell + (1-c)(n-2)$. Selection favors positive mutation rates if this threshold value of u' is positive, which occurs when $\frac{a}{\frac{\mathcal{M}}{\mathcal{M}+1}+a} > c$. For arbitrarily high n and ℓ, this bound can be no worse than $\frac{a}{1+a} > c$. If the baseline fitness f_0 is positive, however, the expression for the maximum invading u' admits no simple formula. Selection is guaranteed to favor positive mutation rates if the largest eigenvalue of $CQ_0 F_{f_0}$ exceeds $1 + f_0$, which is equivalent to

$$\frac{a}{1 + f_0 + a} > c. \qquad (5.15)$$

That is, the selective advantage — as a fraction of maximum possible fitness — must exceed the recombination rate in order for positive mutation to be favored.

The above calculations can be carried out for any game in which competition is nontransitive (each pure strategy is bested by another) and the dynamics cycle between nearly-monomorphic states (formally, the monomorphic states constitute an attracting heteroclinic cycle). In all such games, there exist positive mutation rates that can invade nonmutating residents [6].

Interaction between two strains of nonzero mutation rates $u_1 < u_2$ introduces new features into the evolutionary dynamics. Our numerical analysis of (5.3) (see Methods) has shown three possible outcomes:

(i) The higher-rate strain can invade and replace the lower, and not vice-versa,

$$s_{u_1}(u_2) > 0 > s_{u_2}(u_1);$$

(ii) The lower-rate strain can invade and replace the higher, and not vice-versa,

$$s_{u_2}(u_1) > 0 > s_{u_1}(u_2);$$

(iii) Bistability (neither can invade the other), $s_{u_1}(u_2) < 0$ and $s_{u_2}(u_1) < 0$.

Absent recombination, we have not observed any cases of coexistence between strains of different mutation rate in the rock-paper-scissors game.

To show a representative example of outcome (i), we consider competition between strains with mutation rates $u = 0.02$ and $u' = 0.15$, using $n = 3$ and $a = 0.1$. The faster-mutating strain succeeds because it is better able to adapt to the changing fitness landscape created by the population dynamics, shown in Fig. 5.3.1. Fig. 5.3.2 shows that the rise in abundance of an invading strain is not necessarily monotonic: there may be an initial decrease if the invader is initially in a disadvantageous state, and there may be periodic decreases as the fitnesses of traits shift. Each episode of trait substitution (e.g., scissors replacing paper) is accompanied by a rise in invader abundance; there are 13 such episodes between

timesteps $\sim$ 10 and $\sim$ 190 in the scenario depicted, during which time the invader reaches 1:1 parity with the resident. At the end of each episode, the substitution has occurred in the resident strain as well, causing mutation rate evolution to pause, but also setting the stage for frequency-dependent selection to favor a new substitution. The average slope of invader growth in Fig. 5.3.2 in this initial growth phase equals the invasion fitness $s_u(u')$. After timestep $\sim$ 230, the invader is a large majority of the population, and its growth rate relative to the resident is nearly constant: instead of discrete episodes of substitution and stasis as before, the prevailing mutation rate in the population is so high that the slower-mutating strain never manages to "catch up" to the favored trait.

This example scenario illustrates an important possibility: competition between mutation rates often takes the form of a cooperative dilemma [13], with the strain of higher mutation rate playing the role of defector. When such a strain invades and grows, it increases polymorphism in the population, decreasing overall population fitness. Fig. 5.3.3 shows how fitness depends on the frequencies of A and B. Specifically, when A is abundant and B is rare, the two strains have time-averaged fitnesses $\varphi_A = 0.87$, $\varphi_B = 0.91$. When B is abundant and A is rare, the fitnesses are $\varphi_A = 0.7$, $\varphi_B = 0.72$. Written in matrix form, this "Mutator's Dilemma" is

Figure 5.3.1: Dynamics of a faster-mutating strain successfully invading a slower-mutating resident. Trait distributions of residents (blue trajectories) and invaders (red trajectories) are shown at time intervals (A) 0 to 25, (B) 135 to 175, (C) 200 to 230, and (D) 280 to 300. For each interval, the trajectories proceed counter-clockwise, indicated by shading from darker to lighter. Thickness of each curve indicates relative abundance of the corresponding strain; the thicker trajectories of the invaders in (C) and (D) partially obscure the resident trajectories. At $t = 0$, the resident strain is at the point on its stable limit cycle where frequency of Rock is greatest, and the invading strain is at 100% Scissors with a total frequency of 0.01. Invaders initially decline due to their disadvantageous trait distribution, but eventually sweep to fixation. Thin gray lines are axis guides meeting at the center of the simplex. Parameters: Resident mutation rate $u = 0.02$, invader mutation rate $u' = 0.15$, fitness benefit for winning $a = 0.1$, and $n = 3$ traits.

Figure 5.3.2: Ratio of faster-mutating invader frequency to slower-mutating resident frequency (log-scaled), showing the indirect route to fixation of a mutator allele. The slope equals the relative fitness of the invading strain. The shaded regions correspond to the panels of Fig. 5.3.1, highlighting the four phases of a successful mutator invasion: (A) a possible initial decline due to disadvantageous invader trait distribution, (B) a first growth phase in which fitness fluctuations result from trait dynamics of the slower-mutating residents, (C) a rapid transition as invader abundance surpasses resident abundance, and (D) a second growth phase in which fluctuations are due to the now-abundant invaders. The top-left inset shows one fluctuation during (B), corresponding to one trait substitution; this fluctuation includes a small dip and a rapid rise in invader frequency. The bottom-right inset shows one fluctuation during (D), also corresponding to a substitution of the predominant trait; here frequency increases smoothly since trait substitutions are less dramatic. Both insets use a linear scale. Parameters (same as Fig. 5.3.1): Resident mutation rate $u = 0.02$, invader mutation rate $u' = 0.15$, fitness benefit for winning $a = 0.1$, and $n = 3$ traits.

**Figure 5.3.3:** Time-averaged fitnesses of fast-mutating invaders (upper curve) and slow-mutating residents (lower curve), as functions of invader frequency. To calculate the strain fitnesses at a particular invader frequency, equation (5.3) was modified so that the resident / invader frequency remains constant, permitting only trait fluctuations within a strain (see equation (5.21) in Methods). Evolution always favors the invading strain, but each strain's fitness decreases monotonically with invader frequency; mutator competition is therefore a cooperative dilemma. Inset: Average fitness of the population oscillates and falls over time, as invader abundance increases. Parameters (same as Figs. 5.3.1, 5.3.2): Resident mutation rate $u = 0.02$, invader mutation rate $u' = 0.15$, fitness benefit for winning $a = 0.1$, and $n = 3$ traits.

	A	B
A	0.87	0.70
B	0.91	0.72.

This interaction differs from matrix games as typically considered in evolutionary game theory, since the fitnesses of the two strains depend nonlinearly on their frequencies and fluctuate in time. It does, however, exhibit a key feature of the Prisoner's Dilemma: competition unequivocally favors a strategy that reduces average population fitness.

Considering the entire range of possible mutation rates, the pairwise invasibility plot Fig. 5.3.4(A) shows invasion fitness $s_u(u')$ as a function of u and u' for a representative example with $n = 4$, $a = 0.4$ (see Methods). This plot shows that, in general, residents of very low mutation rate can be invaded by slightly higher mutation rates, and residents of very high mutation rate can be invaded by slightly lower rates. Bistability can occur between strains of low and high rates. It appears in these cases that the lower-rate strain is not able to adapt to the fluctuating fitness landscape created by the higher-rate strain, while the higher-rate strain mutates too rapidly away from advantageous traits when the fitness landscape is more stable.

For the traditional rock-paper-scissors game $(n = 3)$ with $0 < a < 1$, our simulations show that any resident strain of mutation rate $u < u_\Omega = \frac{1+a}{5+a}$ can be invaded by a strain of rate u' slightly larger than u. For $u \geq u_\Omega$, the resident distribution converges to the equilibrium frequencies $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$, at which all
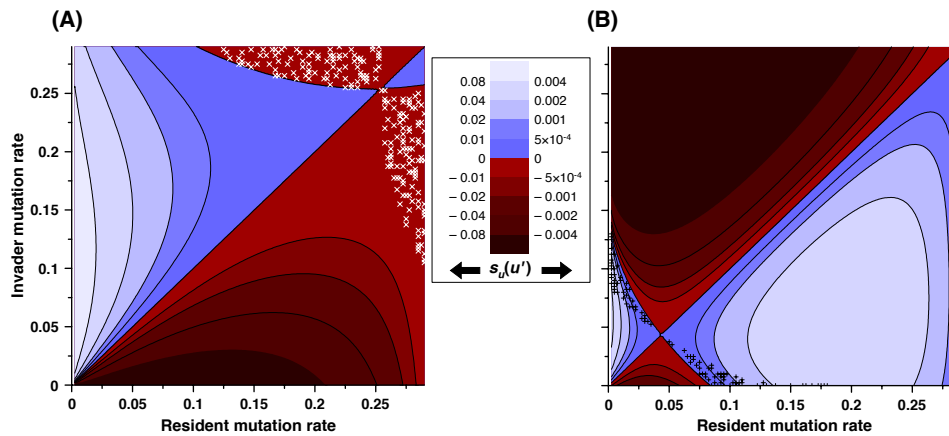
Figure 5.3.4: Pairwise invasibility plots describe the adaptive dynamics of mutator competition (A) without recombination and (B) with recombination. Panels show contour plots of invasion fitness $s_u(u')$, computed using Eq. (5.20) in Methods, as a function of resident and invader mutation rates u and u'. The white x's in (A) mark the region of bistability, where neither mutation rate in a pair can successfully invade the other (negative invasion fitness at a point and at its reflection across the diagonal). The black +'s in (B) mark the region of coexistence, where both mutation rates in a pair can successfully invade the other (positive invasion fitness at a point and at its reflection). Recombination reduces the ESMR from $\tilde{u} \approx 0.25$ to $\tilde{u} \approx 0.04$. Since dynamics are discontinuous at resident mutation rate $u = 0$, invasion fitness against resident nonmutators is plotted separately, in Suppl. Fig. 7.4.1. Parameters: $n = 4$, $a = 0.4$, $c = 0$ (A) or 0.3 (B).

traits are equally fit and the invasion fitness of any other strain is zero. So u_Ω is stable against invaders from below, and it is neither favored nor disfavored against invaders from above. Strictly speaking, there is no ESMR in this instance, though perturbation of the system by introducing rare deleterious mutations would make u_Ω the unique evolutionary endpoint.

For the generalized game with $n > 3$, our simulations consistently show a single ESMR $\tilde{u}$ between 0 and u_Ω, which is also an evolutionary endpoint [6], as plotted in Fig. 5.3.5. The value of $\tilde{u}$ increases with a (up to the point where $\tilde{u} = u_\Omega$), as the potential benefit of a change in trait increases. It decreases with n, as the probability that a mutation is disadvantageous increases.

If mutations are common, the qualitative invasion dynamics appear insensitive to population size, as an analogous finite-population model returns similar results. Fig. 5.3.6 plots the effective invasion fitness implied by the fixation probability of the invading strain, in a population of size $N = 100$, where the invader starts as a single individual. This effective fitness parameter is defined as the constant selective advantage in a Moran model that would produce the same fixation probability as observed in the simulation (see Methods). In a finite population, although an invader may replace the ESMR, it does so only with probability less than the neutral value $1/N$. The ESMR is approximately the same in both the infinite-population and finite-population scenarios ($\tilde{u} \approx 0.25$, see Figs. 5.3.4 and 5.3.6). The magnitudes of the effective invasion fitnesses are overall smaller in the finite case, likely owing to additional stochastic effects introduced by mutator dynamics; for instance, the time that a lone invader must wait before it produces
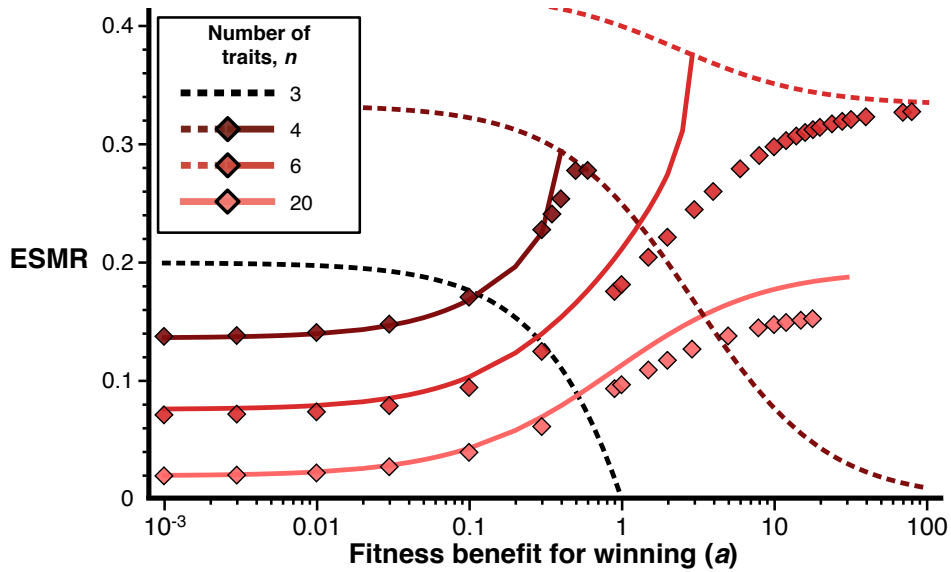
Figure 5.3.5: The ESMR increases with fitness benefit a (plotted on a log scale) and decreases with number of traits n; no recombination is shown in this figure. Points show simulated ESMR values (see Methods), and solid lines show approximations using Eq. (5.17). Dashed lines plot u_Ω, which is "quasi-stable" for cases where an actual ESMR is not defined (see text). For the case $n = 3$ (black), there is no ESMR and only u_Ω is shown. For $n = 4$ (dark red) and $n = 5$ (not shown), the ESMR increases with a until it reaches u_Ω; for higher values of a there is no ESMR. For $n = 6$ (medium red), the ESMR appears to converge to u_Ω as a increases. For $n > 6$ ($n = 20$ in pink), the ESMR is bounded below u_Ω, which is greater than 0.4.

a beneficial mutation varies between runs. In general, this factor substantially influences survival and fixation of an invading mutator strain [302, 307].

Recombination lowers the ESMR

Fig. 5.3.4(B) shows the effect of recombination on an invasion fitness landscape, using the replicator-mutator-recombinator equations (5.7). In the case illustrated, the value of the ESMR decreases but remains positive when recombination is introduced. One important effect of recombination is that coexistence between mutation rates above the ESMR and below the ESMR becomes possible. In this case, the two rates appear to serve complementary functions: beneficial mutations acquired by the higher-rate mutators are transferred by recombination to lower-rate mutators, which faithfully replicate them. Numerical simulations show, however, that a third, intermediate mutation rate can invade this polymorphic state. If the mutation rate evolves in small steps (such that an intermediate rate is attainable), then the ESMR is still the unique endpoint of evolution in this game [103]. Moreover, if two competing strains have very similar mutation rates, coexistence is never observed. In other words, even with recombination, the rock-paper-scissors game appears to behave according to the "invasion implies substitution" principle of adaptive dynamics [67].
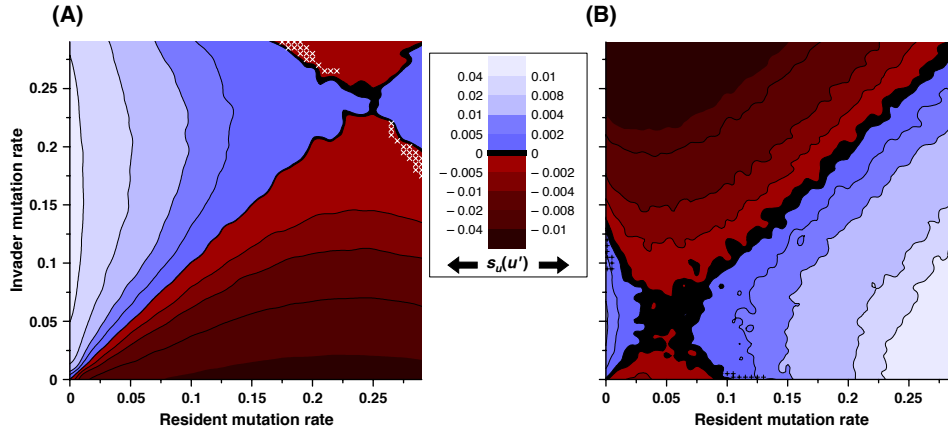
Figure 5.3.6: Stochastic simulation confirms that deterministic replicator-mutator-recombinator equations can be used to determine the structure of mutator competition in finite populations without (A) and with (B) recombination, as long as $\tilde{u}N \gg 1$. Panels show contour plots of effective invasion fitness $s_{e,u}(u')$, computed using Eq. (5.23) in Methods, as a function of resident and invader mutation rates u and u'. If $s_{e,u}(u') > 0$, then the invader has fixation probability greater than neutral probability $1/N$; the opposite is true for $s_{e,u}(u') < 0$. The white x's in (A) mark the region of bistability, where neither mutation rate in a pair is favored to invade the other (negative effective invasion fitness at a point and at its reflection across the diagonal). The black +'s in (B) mark the region of coexistence, where both mutation rates in a pair are favored to invade the other (positive effective invasion fitness at a point and at its reflection). The black region between positive and negative fitness shows nearly neutral competition (effective invasion fitnesses within (A) 5×10^{-5} or (B) 4×10^{-4} of zero). Population size $N = 100$, other parameters (same as Fig. 5.3.4): $n = 4$, $a = 0.4$, $c = 0$ (A) or 0.3 (B).

Unconditionally deleterious mutation lowers the ESMR

Introduction of unconditionally deleterious mutation with a target size ℓ further decreases the ESMR. In the absence of recombination, the ESMR scales roughly as $1/\ell$. With recombination, it decreases by a larger negative power of ℓ (Fig. 5.3.7).

Estimating the ESMR in the absence of recombination

To gain insight into how the ESMR depends on the costs and benefits of mutation, we develop an analytical approximation of the ESMR as a function of fitness benefit for winning a and deleterious target size ℓ. We adapt an argument made in a frequency-independent setting [174] that itself has its roots in an earlier model of trade-offs between time required for adaptive substitutions (minimized by a high mutation rate) and mutational load (minimized by a low mutation rate) [162, 163].

We first consider the case without unconditionally deleterious mutation, $\ell = 0$. For values of a and u that are not too large, the trait frequencies transition between consecutive nearly-monomorphic states, each of which is a temporary balance between mutation and selection. For each state, there is a newly favored trait that starts to grow. The long-term performance of a strain therefore depends on how quickly its frequency within a newly favored trait increases.

At one of these temporary mutation-selection equilibria, each minority trait is present essentially only due to the action of mutation; we suppose that the next

Figure 5.3.7: The ESMR declines with the target size of unconditionally deleterious mutation (ℓ) and the recombination rate (c). Solid trend lines (shown for all series except $c = 0.016$) on the log-log axes track the decline of ESMR with ℓ; the slope for the series $c = 0$ is approximately -1, as predicted by Eq. (5.17). Series with higher c have steeper slopes and lower intercepts. The horizontal dashed line at $u = 5 \times 10^{-10}$ represents a typical per-site, per-replication mutation rate for eukaryotes; below this level, simulation was not reliable due to numerical error in the stiff system. The dotted line for the series $c = 0.016$ is a rough estimate; only one reliable ESMR value was computed. Though the ESMR falls rapidly with increasing c and ℓ, Eq. (5.14) guarantees a positive ESMR for any $c < \frac{a}{1+a} = \frac{1}{3}$. Parameters: $n = 3$ traits, fitness benefit for winning $a = 0.5$, baseline fitness $f_0 = 30$.

round of frequency-dependent selection has not yet made an appreciable contribution to the abundance of the newly favored trait. The frequency of the newly favored trait within a strain is therefore proportional to that strain's mutation rate. If trait i has just been selected, then the ratio of invader to resident within trait $i + 1$ is $\frac{u'}{u}$. This ratio represent's the invader's relative advantage due to mutation; if $u' > u$, then the invader has a "head start" in the race to increase its frequency within the soon-to-be-majority trait.

As trait $i + 1$ grows, its representatives in each strain experience a mutational load per generation equal to their respective deleterious mutation rates — $u\frac{n-1}{n}$ for the resident, $u'\frac{n-1}{n}$ for the invader [162, 163]. Since nearly all mutations away from the optimal trait are lethal, this cost is fully realized, with negligible delay (c.f. 105). If the entire limit cycle has period T, then one substitution in the cycle takes time $\frac{T}{n}$. The total multiplicative effect of mutational load on the invader population (i.e., the factor by which it shrinks, relative to the ideal case without load) during this leg is $e^{-u'\bar{f}T\frac{n-1}{n^2}}$, where $\bar{f}$ is the geometric mean fitness of the favored trait. Note that both $\bar{f}$ and T depend only on the resident's mutation rate, as the invader has negligible effect on the fitness landscape and dynamics.

Combining both the advantageous and disadvantageous effects of mutation, the relative performance of the invading strain is proportional to $u'e^{-u'\bar{f}(u)T(u)\frac{n-1}{n^2}}$. This expression attains its maximum at

$$u' = \frac{n^2}{T(u)\bar{f}(u)(n-1)}. \tag{5.16}$$

If the rate $\tilde{u}$ is an ESMR, then the mutation rate best equipped to invade $\tilde{u}$ is $\tilde{u}$

itself. Thus $\tilde{u}$ satisfies

$$\tilde{u}\bar{f}(\tilde{u}) \approx \frac{n^2}{T(\tilde{u})(n-1)}. \qquad (5.17)$$

Both $\bar{f}(u)$ (calculated as the geometric mean of the optimal trait's fitness) and $T(u)$ can be obtained numerically as functions of u through simulation of the resident replicator-mutator equations (5.2) (see Methods). This method yields an accurate prediction of ESMR values, particularly for small a (Fig. 5.3.5). The ESMR increases with a, unless it reaches u_Ω, after which there is no ESMR, as invaders $u' > u_\Omega$ have zero invasion fitness against the resident u_Ω. For $n \geq 6$, the ESMR never reaches u_Ω, and so it always exists and increases with a.

As an alternative to numerical simulation, the fitness $\bar{f}(u)$ can be approximated by

$$\bar{f}(u) \approx (1+a)\left(1 - u\frac{n-1}{n}\right) + \frac{u}{n}. \qquad (5.18)$$

This approximation is based on the notion that the transient equilibrium state describes the typical trait distribution and represents a mutation-selection balance, at which one trait has frequency $\approx 1 - u + \frac{u}{n}$ and the remaining population is divided roughly equally among the other strategies. Using this approximation in Eq. (5.17) gives an ESMR estimator that only requires measurement of the period. This approach is appropriate for $n \geq 6$ and $a \ll 1$, but Eq. (5.18) substantially overestimates fitness for larger a or smaller n, producing an ESMR prediction well below the actual value (Suppl. Fig. 7.4.2).

If the deleterious target size ℓ is positive, then the above argument is modified slightly. The mutational load per generation is instead $u\left(\frac{n-1}{n}+\ell\right)$ for the resident, with u replaced by u' for the invader. The ESMR $\tilde{u}$ then satisfies

$$\tilde{u}\overline{f}(\tilde{u}) \approx \frac{n^2}{T(\tilde{u})(n(1+\ell)-1)},\qquad (5.19)$$

which scales as $1/\ell$ in the limit of large ℓ. Numerical simulation confirms this asymptotic relationship between $\tilde{u}$ and ℓ (Fig. 5.3.7, $c = 0$).

5.4   Discussion

5.4.1   Overview

Sustained adaptive evolution of high mutation rates requires neither exogenous environmental variation nor an inexhaustible pool of possible adaptive mutations; rather, it may result from cyclical competition dynamics taking place within a population. The fitness benefit for prevailing in competition, the number of possible deleterious mutations, and the recombination rate together determine the evolutionarily stable mutation rate, or ESMR. This mutation rate can exceed typical per-generation mutation rates by several orders of magnitude. Strong cyclical competition such as that considered in Fig. 5.3.7 (fitness bonus $a = 0.5$ as a fraction of maximum fitness $f_0 + 1 + a = 31.5$ corresponding to a selection coefficient of 1.6%) selects for particularly high mutation rates in the absence of recombination, even for lethal mutation target sizes of $10^4 - 10^6$ sites (Fig. 5.3.7). This mechanism is one means by which global hypermutators may be selected

for, and it may help explain the prevalence in the wild of mismatch repair-deficient bacterial strains with genome-wide mutation rates 100- to 1,000-fold that of wild type [65, 191]. Recent experimental and theoretical work suggests that host-pathogen competition, which can be thought of as a type of cyclical competition with two separate populations, also supports evolution of high mutation rates [150, 201, 227].

Whether a mutator strain can invade a resident with low mutation rate depends on the period of the resident's trait substitutions; a shorter period favors faster-mutating invaders (Eq. (5.16)). If the mutator achieves fixation, it generally does so in a succession of "mini-sweeps," each one corresponding to a single trait substitution (Fig. 5.3.2). Patterns of episodic partial sweeps have also been observed in stochastic models of mutator evolution [302, 307]. Given the infinite population in our model, the limiting factor on sweep timing is not the appearance of mutations, but rather the time it takes for the frequency-dependent fitness landscape to shift, favoring the next trait in the cycle. This timing depends on the strength of selection itself. One aspect of frequency-dependent competition that we did not explore is that an ESMR need not be unique. Since the favored mutation rate depends on the period of trait cycling, we expect that systems with multiple cyclical attractors of different periods will have multiple ESMRs, such that long-term mutation rate evolution would depend on initial trait distributions.

Recombination is known to temper selection for global mutator alleles [147, 175, 295, 308], and our model agrees with others in this respect. We

introduced and analyzed replicator-mutator-recombinator equations (5.7) to conclude that the selective advantage from winning the cyclical competition must exceed the recombination frequency between the trait and mutator loci in order for a mutating strain to be favored over nonmutators (Eq. (5.15)). This conclusion holds regardless of the deleterious load incurred. The favored positive mutation rates may, however, become very small as load and recombination increase: For a selection coefficient of 1.6% and deleterious target size of 10^5, introducing recombination at a distance of 0.4 centimorgans between the two loci causes the ESMR to drop from 3×10^{-7} to 10^{-8}. Doubling the distance to 0.8 centimorgans brings the ESMR below 10^{-10} (Fig. 5.3.7). At this value, the infeasibility of achieving perfectly faithful replication would enforce a floor on the mutation rate, rather than any adaptive benefit of the ESMR. Recombination also permits stable coexistence between multiple mutation rates (Figs. 5.3.4(B), 5.3.6(B)) by allowing a "division of labor" in adaptation: beneficial traits arise more frequently within the high-rate strain, but they are more likely to persist after crossing into the low-rate strain. If fine-tuning of the mutation rate to the ESMR is not possible, the next-best state of affairs may involve such coexistence. A previous modeling study demonstrated that even the low rates of recombination found in bacterial species are sufficient for non-mutating lineages to "steal" adaptive mutations generated by a mutator strain, increasing the population's rate of adaptation while severely limiting the mutator's chance of fixation [308]. Stable coexistence was not observed in that study, however, as the simulated genome contained few sites at which beneficial

mutations could appear, and simulation ended once most of the population carried all favorable alleles at these sites. The adaptive complementarity of mutator and non-mutator phenotypes is also suggested by evidence that mismatch repair genes in *Escherichia coli* have undergone frequent horizontal transfer [65].

5.4.2 Limitations of the model

Our model is situated in an adaptive dynamics framework that we introduced for studying evolution of mutation rates under general cases of frequency-dependent competition [6]. Within this framework, adaptive evolution of high mutation rates may occur only in the absence of a stable equilibrium trait distribution — hence our choice to analyze "rock-paper-scissors" competition, the simplest unstable case. Our framework relies on two simplifying assumptions: competition between two mutation rates follows deterministic population dynamics and is fully resolved before a third mutation rate is introduced. These two formal assumptions constrain the biological scenarios to which the framework applies. First, population size must be sufficiently large that both beneficial and deleterious mutations appear with some regularity (as in Fig. 5.3.6); if there are many generations between mutations (even those caused by the high-rate strain), then fitness differences between the two strains will be negligible for long periods of time, during which mutation rate competition would be governed by neutral drift [307]; this scenario is not contemplated by our model. Second, the framework rules out cases of runaway selection for high

mutation rates: if the mutation rate is itself labile and most deleterious mutations are mild, it is possible for rapid increases in the mutation rate to occur, subjecting the population to rapid fitness declines (or error catastrophe) when mutational load equilibrates [7, 105]. Since our framework assumes that deleterious mutations are strong, all mutational load costs are fully realized by the prevailing mutation rate. Though our framework makes strong assumptions, we anticipate selection pressures for increased mutation rates to be a general consequence of cyclical and other unstable frequency-dependent interactions.

5.4.3   The mutator's dilemma

Though mutator alleles are commonly viewed as mechanisms for increasing the adaptation rate and fitness of populations, our model postulates a mechanism by which a high mutation rate, though favored by evolution, can cause average population fitness to decrease over time (Fig. 5.3.3). If, as in the payoff matrix considered, population fitness declines with trait diversity, then a successful mutator allele essentially behaves as a selfish variant in a cooperative dilemma. Prior findings that mutator lineages generate "social cheats" in microbial populations demonstrated that mutator evolution can lead to a breakdown in cooperative scavenging behaviors [39]. Our model broadens the concept of mutator-mediated selfishness by showing how a cooperative dilemma may emerge from the dynamics of mutator competition, even though the underlying interaction does not involve cooperative or selfish traits. Together with the case of runaway selection described above, our model represents a third possible

mechanism by which mutator alleles may prevail but subvert natural selection.

5.4.4 MECHANISMS FOR MUTATION RATE EVOLUTION

Our model can represent a range of different biological mechanisms controlling mutation rates. Mechanisms that modulate the genome-wide mutation rate, such as altered DNA polymerases, would have large associated rates of deleterious mutation and recombination (large ℓ and c values), suggesting that their role in adaptive mutation may be limited in sexual species [204, 278, 295]. On the other hand, mechanisms affecting local mutation rates correspond to lower ℓ and c values, and should be more responsive to selection for increased mutation rate on a trait locus. A number of such mechanisms have been identified, typically involving features of the immediately surrounding sequence.

First, the rates of different types of nucleotide substitutions appear to be modulated by the frequency of certain short (3 bp) sequences up to 100 bp away [82, 127] and by GC content up to 1500 bp away [87, 95, 183, 330]. In the latter case, each individual site contributing to local GC content may act as a "fine-tuning" local mutator allele. One consequence not explored in our model is that certain features have a directional effect; e.g., low GC content promotes methylation-dependent mutation of C to T (G with A), but not the reverse [87, 95, 183, 330].

Second, short DNA repeats, which can experience frequent duplication and deletion events caused by slipped-strand mispairing, may be potent local mutators. If a repeat originates within an open reading frame and its subsequent

duplication or deletion would result in a downstream frame shift, the increased mutation rate can cause rapid switching between alternate gene products [37, 164, 178, 179, 205, 278]. Even if the repeat itself does not appear in a coding region, the sequence length heterozygosity that it generates can increase nucleotide substitution rates by an order of magnitude within a 100 bp radius, with smaller noticeable effects out to 600 bp [46, 82, 309].

Since a local mutator tends to remain linked to the nearby mutations that it causes, recombination is unlikely to substantially impede selection for increased mutation. Consider the 3-trait competition in Fig. 5.3.7 with a trait-controlling site at the end of a 500 bp effect radius of a local mutator. The deleterious target size ℓ would be on the order of 10^3, the sequence length afflicted by increased deleterious mutation. The ESMR in this scenario is $\tilde{u} = 6 \times 10^{-5}$. To obtain five-fold reduction in the ESMR, the recombination probability would need to be $c = 4 \times 10^{-3}$; over 500 bp, this value requires a recombination rate of 800 cM/Mb, extraordinarily high among multicellular eukaryotes, even at hotspots [5, 143, 171, 172, 200, 322].

Our model can also describe migration of a trait-controlling gene to a region with a different mutation rate. Mutation rates vary across regions of the genome due to features such as DNA structure, conformation, and replication timing [14, 35, 173, 246, 286]. In the asexual ($c = 0$) case, our model can treat individuals affected by the gene movement as the invading strain. Parameter ℓ then describes the deleterious target size of the region that is transferred to the new mutational milieu, including the trait-controlling gene. Recombination

would introduce genetic incompatibilities, requiring modifications to the model.

5.4.5 Outlook

The predictions of our model apply broadly, as cyclical competition is present in diverse biological scenarios. The strategic dynamics among male mating strategies in the side-blotched lizard, *Uta stansburiana* [289, 290], is a well-documented example. The three male morphotypes of this species vary in aggressiveness, territoriality, and cooperation in territory defense, generating a rock-paper-scissors game and cyclical population dynamics among the types [289–291]. Similar dynamics have also been observed in the viviparous lizard, *Lacerta vivipara* [292]. *E. coli* also exhibits rock-paper-scissors dynamics with regard to production of and resistance to colicins, a form of bacteriotoxin [157, 165]. More generally, cyclical competition should be common wherever individuals face fitness tradeoffs on multiple interacting dimensions, such as life history characters, sexual behaviors, and social behaviors [288]. Since selection coefficients for social or sexual traits tend to be large, we believe that the 1.6% advantage considered in Fig. 5.3.7 is quite reasonable. Our model predicts elevated levels of mutation on the loci contributing to behavior in these cyclical interactions, as compared with the genome-wide average. Moreover, global mutation rates are generally predicted to be higher in asexual species where cyclical dynamics are common.

5.5 Methods

5.5.1 Numerical simulation of ODE and invasion fitness calculation

Numerical analysis of the replicator-mutator-recombinator equations (5.7) was done using Mathematica version 7.0 for Linux x86 (64-bit), using default options unless otherwise specified. To compute invasion fitness, the following modified dynamics were used to treat the invader strain as forever rare:

$$\dot{r}_i = (1-u)f_{i,u}r_i + \left(\frac{u}{n}\right)\varphi_r - \varphi_r r_i,$$

$$\dot{z}_i = (1-c)\left[(1-u')f_{i,u'}z_i + \left(\frac{u'}{n}\right)\varphi_z\right] \qquad (5.20)$$
$$+ c\left[(1-u)f_{i,u}r_i + \left(\frac{u}{n}\right)\varphi_r\right]$$
$$- \varphi_z z_i.$$

Here, r_i and z_i are interpreted as the proportion of trait i within the resident strain and invader strain, respectively, not within the total population. Accordingly, the final term of each equation uses φ_r or φ_z, not φ, ensuring that $\sum_i r_i = \sum_i z_i = 1$ at all times. Fitness computation ignores the rare invader, and so $f_{i,u} = A\mathbf{r} - u\ell$. From the resident perspective, recombination with an invader never occurs, and so the equation for $\dot{r}_i$ recapitulates the form of the replicator-mutator equation. From the invader perspective, all of the fraction c offspring that recombine do so with a resident. The dynamics were simulated using the NDSolve option, "MaxSteps -> 100000", and numerical integration of equation (5.6) was done

using the NIntegrate option, "Method -> LocalAdaptive". The population started

with only trait 1, and the first 1500 timesteps were discarded to ensure

convergence to the limiting trajectory. The time limits of integration in (5.6)

were then chosen to contain an integer number of oscillations (at least 10 as

counted by the method below, and not more than 1500 timesteps). In the case

where the resident was a nonmutator, there are no oscillations, and the limits of

integration were $[1500, 3000]$.

For the time-averaged fitness calculation in Fig. 5.3.3, a modified version of

(5.3) was used to keep the invader/resident ratio constant:

$$
\begin{aligned}
\dot{r}_i &= (1 - u)r_i \left(A\left((1 - Z)\,\mathbf{r} + Z\mathbf{z}\right)\right)_i + \frac{u}{n}\varphi_r - \varphi_r r_i \\
\dot{z}_i &= (1 - u')z_i \left(A\left((1 - Z)\,\mathbf{r} + Z\mathbf{z}\right)\right)_i + \frac{u'}{n}\varphi_z - \varphi_z z_i.
\end{aligned}
\tag{5.21}
$$

As above, r_i and z_i are interpreted as trait frequencies within the respective strain.

The value of Z determines the invader/resident ratio and is constant. Numerical

integration proceeded as above to obtain the time-average values of φ_r and φ_z.

5.5.2   MEASURING THE PERIOD OF OSCILLATIONS

To measure the period of oscillations when just the resident strain is present, a

piecewise projection function π was devised to track the resident population

vector's progress around the cycle:

$$
\pi(\mathbf{r}) = \begin{cases} m + \dfrac{r_{m+1} - r_{m-1}}{r_m + r_{m+1} - 2r_{m-1}} & \text{if } r_{m-1} < r_{m+1}, \\[2ex] m - \dfrac{r_{m-1} - r_{m+1}}{r_m + r_{m-1} - 2r_{m+1}} & \text{if } r_{m-1} > r_{m+1}, \\[2ex] m & \text{if } r_{m-1} = r_{m+1}, \end{cases}
\tag{5.22}
$$

where $m = \mathrm{argmax}(\mathbf{r})$, $m + 1$ is replaced with 1 if $m = n$, and $m - 1$ is replaced with n if $m = 1$. This quantity progresses from 0 to $n + 1$, though not monotonically or continuously, as the vector $\mathbf{r}$ traverses the cycle. For the purpose of counting the number of oscillations, the small discontinuities and declines can be handled by the following algorithm: Data from the first 1500 timesteps were not used, to allow the period to stabilize, and then the value of $\pi(\mathbf{r})$ was calculated from timestep 1500 to 3000, in increments of 0.1. Each timestep where $\pi(\mathbf{r}) \in (0.1, 0.9)$ was marked with an **A**, each timestep where $\pi(\mathbf{r}) \in (1.1, 1.9)$ was marked with a **B**. Each yet-unmarked timestep was then marked with the symbol of the most recent marked timestep. A new oscillation is then deemed to start at each transition from **A** to **B**. The average period of all oscillations was then used.

5.5.3 ESMR CALCULATION

A binary search algorithm was used to find the ESMR via simulation of (5.7). Mutation rates strictly between 0 and u_Ω were explored to find a value $\tilde{u}$ that could resist invasion both by $1.01 \times \tilde{u}$ and by $\tilde{u}/1.01$. This analysis is justified for the generalized rock-paper-scissors game based on the observations that (1) if u

can be invaded from slightly above, then $\tilde{u} > u$, (2) if u can be invaded from slightly below, then $\tilde{u} < u$, (3) if u resists nearby invaders, then u resists all invaders. These observations need not hold for more complicated games. The rate $u_\Omega/(50 + \ell)$ was used as an initial guess for $\tilde{u}$. At time 0, all individuals were of trait 1 and the invading strain frequency was 10^{-8}. Invasion was deemed to fail if the invader frequency after 3000 timesteps was less than this initial value.

5.5.4 FINITE POPULATION SIMULATION

Finite population dynamics were simulated as a frequency-dependent Moran process with constant population size 100. Each timestep, every individual interacted with every individual (including itself), receiving fitness 1 from interaction with another of the same trait, fitness $1 + a$ from interaction with one of the trait that it defeats, and 0 from all other interactions. One individual was chosen proportional to fitness to reproduce, its offspring replacing one randomly chosen (regardless of fitness) individual (possibly the parent). At each reproduction event, the offspring inherited its parent's trait with probability $1 - u$ and selected a random trait with probability u, where u is the parent's mutation rate. The offspring's mutation rate always equaled that of the parent. The offspring then replaced a randomly chosen individual in the population (possibly its parent).

For simulations with recombination, with probability c each timestep, two (possibly mutated) offspring are generated and produce a recombinant, which has the mutation rate of one and the trait of the other. This recombinant then

replaces a randomly chosen individual in the population (possibly one of its parents).

The Moran simulation was first run with only the resident strain, to create representative initial trait distributions in which the invader would appear: starting from a population with 25 of each trait, the resident population evolved for 5000 timesteps, after which the population frequencies were sampled every 1000 timesteps thereafter. 1000 frequency distributions were sampled in all. In each of the sampled populations, one individual was randomly designated to be the initial invader, with mutation rate u'.

Next, to determine the fixation probability of an invading strain for a particular pair of values (u, u'), 5,000,000 separate invasion attempts were simulated, each starting at a trait frequency distribution randomly selected from the 1000 distributions created previously. Each simulation progresseed until one of the two strains reached fixation.

Effective invasion fitness $s_{e,u}(u')$ was computed by solving for s in the equation

$$\rho = \frac{1 - \frac{1}{1+s}}{1 - \frac{1}{(1+s)^N}}, \tag{5.23}$$

(see [90], Eq. 3.66) where $N = 100$ is the population size and ρ is the fraction of the 5,000,000 simulations for which the invader went to fixation.

While it may be more plausible to suppose that the initial invader is a randomly chosen *offspring* that had just mutated to the new mutation rate u', rather than a randomly chosen individual, this method would unduly favor the invading strain, as it would be more likely to start as the high-fitness offspring of a

high-fitness parent. In that case, a strain would be favored to invade against a resident of the very same mutation rate, invalidating the ESMR concept.

All finite population simulations were carried out using Matlab R2010b.

5.6 Acknowledgements

5.7 Manuscript information

5.7.1 In submission as

This manuscript is being submitted as:

Rosenbloom DIS, Allen B. Frequency-dependent selection can lead to evolution of high mutation rates.

6

Mutation rate evolution in replicator dynamics

6.1 INTRODUCTION

THE MUTATION RATE OF AN ORGANISM IS SUBJECT TO HERITABLE VARIATION, and therefore evolves [4, 85, 166, 244, 250, 295, 317, 319]. In populations that are well-adapted to a static environment, most mutations are neutral or deleterious, and mutation rates are therefore expected to evolve to levels as low as feasible, given the constraints and costs associated with faithful

replication [163, 174, 175, 295]. However,

theoretical [85, 141, 158, 174, 226, 299, 307, 313, 326] and

empirical [107, 260, 294, 324] evidence is growing that in novel or rapidly

changing environments, evolution may select for higher mutation rates. While

strains with higher mutation rates are susceptible to increased mutational load,

they are also more likely to produce beneficial mutations, upon which the linked

genes coding for increased mutation may "hitchhike" to fixation. The success of

mutator strains depends on whether beneficial mutations occur with sufficient

frequency and strength to outweigh load [326].

Theoretical study of mutation rate evolution has focused principally on cases

where the environment is either static or fluctuates independently of the evolving

population. Cases of fitness fluctuations generated by the evolving population

itself remain unexplored. Such fluctuations may arise through

frequency-dependent selection [120, 288] in scenarios such as predator-prey

systems [60, 169, 192] (including Red Queen dynamics [68, 71, 159, 190])

host-parasite systems [193, 274], cooperative dilemmas [108, 203, 212, 251], and

nontransitive competition systems [40, 157, 165, 231, 289, 290].

This work introduces a theoretical framework for studying the evolution of

mutation rates in a large population subject to frequency-dependent selection.

Our framework combines two major approaches to evolutionary modeling:

replicator dynamics [132, 134, 135, 218, 270, 305], which studies the

demographic dynamics of trait frequencies in a large population, and adaptive

dynamics [70, 103, 131, 199], which studies long-term evolution through

sequences of trait substitutions.

The model we describe considers two loci, a *trait locus* determining the phenotypic traits that are affected by frequency-dependent selection, and a *mutation rate locus* determining the rate of mutation on the trait locus. Replicator dynamics with mutation [114, 128, 221] are used to model evolution on the trait locus, while evolution on the mutation rate locus is studied using the adaptive dynamics approach. We assume that that evolution is faster on the trait locus than on the mutation locus, allowing a separation of timescales. Classic theoretical works on the evolution of mutation rates [7, 147, 174] implicitly make a similar assumption, by focusing their analysis on the fate of a given mutation rate modifier over long periods of time involving multiple selective sweeps of beneficial mutations.

We present results on both the upward and downward evolution of mutation rates. We show that when frequency dependence leads to a stable mutation-selection equilibrium, mutation rates will evolve downward. However, when trait dynamics converge to a heteroclinic cycle, as can occur in "rock-paper-scissors" scenarios and other cases of non-transitive competition, positive mutation rates can arise and persist. This suggests that genetic loci coding for such rock-paper-scissors traits may have abnormally high mutation rates.

We also show, perhaps surprisingly, that mutating and non-mutating strains can coexist indefinitely. This coexistence is possible—and occurs generically—whenever there exists an evolutionarily stable state of the replicator

dynamics in which all traits are present in a selective balance.

We caution that since our model focuses on just a single trait locus, it does not incorporate the broader (often deleterious) effects that mutation can have on the rest of the genome. Our work therefore most clearly addresses cases such as short sequence motifs and DNA repeats that modify only the mutation rate in the immediate genetic neighborhood [37, 46, 82, 87, 95, 127, 164, 178, 179, 183, 205, 278, 309, 330]. Mutator alleles of global effect, as have been discovered in a number of bacteria [45, 107, 224, 260, 294, 324], could be incorporated into our model by adding terms that describe genome-wide mutational load [163]; however, we do not pursue this avenue here.

We first review necessary background on dynamical systems theory and standardize our notation in Sect. 6.2. Section 6.3 presents our two-locus model and describes how evolution proceeds on each locus. Our results on the evolution of mutation rates are discussed in Sect. 6.4.

6.2 DYNAMICAL SYSTEMS THEORY BACKGROUND AND NOTATION

We first introduce several notions from dynamical systems theory that are used in our work. For this section, let D be the closure of a bounded open set in $\mathbb{R}^d$. Consider a differential equation on D of the form $\dot{\mathbf{x}} = \mathbf{g}(\mathbf{x})$, where $\mathbf{g} : D \to \mathbb{R}^d$ is smooth. The solution to this equation can be described by a smooth flow on D, that is, a collection of smooth maps $\psi_t : D \to D$ satisfying $\psi_{s+t} = \psi_s \circ \psi_t$. The orbit associated with a given initial condition $\mathbf{x}(0) = \mathbf{x}_0 \in D$ is described by

$\mathbf{x}(t) = \psi_t(\mathbf{x}_0)$. We assume $\mathbf{g}$ is such that $\psi_t(\mathbf{x}) \in D$ for all $\mathbf{x} \in D, t \in \mathbb{R}$ (that is, solutions persist indefinitely both forwards and backwards in time).

6.2.1 ω-LIMIT SETS

The *ω-limit set* of a point $\mathbf{x} \in D$, which we denote $\omega(\mathbf{x}) \subset D$, is the set of points $\mathbf{y} \in D$ for which there exists an increasing sequence of real numbers $\{t_i\}_{i=1}^{\infty}$ with

- $\lim_{i\to\infty} t_i = \infty$,

- $\lim_{i\to\infty} \psi_{t_i}(\mathbf{x}) = \mathbf{y}$.

In words, $\omega(\mathbf{x})$ is the set of points that are asymptotically approached by the orbit $\psi_t(\mathbf{x})$ as $t \to \infty$. ω-Limit sets are invariant (both forwards and backwards); that is, if $\mathbf{y} \in \omega(\mathbf{x})$ then $\psi_t(\mathbf{y}) \in \omega(\mathbf{x})$ for all $t \in \mathbb{R}$ (e.g., 47.

6.2.2 ATTRACTORS

An attractor is a set to which an open set of points converge under a flow as $t \to \infty$. This notion has been mathematically formalized in various ways. The definition we use is adopted from Gyllenberg et al. [113] and Geritz et al. [104].

We define an *attractor* as a subset $A \subset D$ satisfying the following conditions:

1. There exists an open neighborhood U of A for which

$$\lim_{t\to\infty} \text{dist}(\psi_t(\mathbf{x}), A) = 0,$$

for all $\mathbf{x} \in U$.

2. For each open neighborhood V of A, there exists an open neighborhood W of A for which $\psi_t(\mathbf{x}) \in V$ for all $\mathbf{x} \in W$ and $t \geq 0$.

6.2.3   THE SOJOURN TIME MEASURE

We will be also be interested in describing the asymptotic behavior of an orbit in statistical terms. To this end, we define a probability measure that captures, roughly, the probability that an orbit occupies a certain region of D after a long amount of time has passed. This measure was introduced by Takens [301] and was named the *sojourn time measure* by Bonneuil [34].

Given an initial point $\mathbf{x} \in D$, and a closed subset $U \subset D$, the sojourn time measure of U is defined by

$$\sigma_{\mathbf{x}}(U) = \lim_{\varepsilon \to 0} \lim_{T \to \infty} \frac{1}{T}\lambda\big(\big\{t : 0 \leq t \leq T, \mathrm{dist}(\psi_t(\mathbf{x}), U) < \varepsilon\big\}\big), \qquad (6.1)$$

where λ is the Lebesgue measure on $\mathbb{R}$. The closed subsets of D for which the above limits exist are taken as a basis for a σ-algebra, to which the measure $\sigma_{\mathbf{x}}$ is extended. Informally, $\sigma_{\mathbf{x}}(U)$ quantifies the asymptotic proportion of time that the orbit $\psi_t(\mathbf{x})$ spends in or near U, as $t \to \infty$.

The sojourn time measure is a probability measure, meaning that $\sigma_{\mathbf{x}}(D) = 1$, for each $\mathbf{x} \in D$. Moreover, it satisfies $\sigma_{\mathbf{x}}(\omega(\mathbf{x})) = 1$; that is, it is concentrated entirely on the ω-limit set of $\mathbf{x}$. The sojourn time measure is also invariant: for each measurable $U \subset D$, $\sigma_{\mathbf{x}}(\psi_t(U))$ is constant in t.

As a trivial example, we observe that if the orbit associated with $\mathbf{x}$ converges to a fixed point $\tilde{\mathbf{x}}$, then the sojourn time measure is a point mass (Dirac

δ-distribution) at $\tilde{\mathbf{x}}$.

6.3   Model

Here we outline a mathematical framework for studying the evolution of mutation rates under frequency-dependent selection. Section 6.3.1 introduces the two loci in our model: one controlling the traits subject to frequency-dependent selection, and the other controlling the mutation rate. Section 6.3.2 presents our model for evolution on the trait locus, using the framework of replicator dynamics. The model for evolution of the mutation rate, based on the adaptive dynamics approach, is presented in Sect. 6.3.3.

6.3.1   Two-locus description

Our model considers two evolving loci:

- The *trait locus* controls the traits that are directly affected by frequency-dependent selection. We consider n competing alleles on the trait locus, indexed $i = 1, \ldots, n$. Each allele produces a distinct trait.

- The *mutation rate locus* controls the rate of mutation on the trait locus. Alleles on the mutation rate locus are represented by the mutation rate $u \in [0, 1]$.

We assume that mutation on the mutation rate locus is rare, compared to typical mutation rates on the trait locus. Thus, evolution of the mutation rate u occurs on a longer timescale than evolution of the traits $i = 1, \ldots, n$. (Here and

henceforth, the term "mutation rate" refers to the rate of mutation on the trait locus, as determined by the mutation rate locus.) Appealing to this separation of timescales, we suppose that at most two mutation rates are present at any given time. That is, we assume that competition between any two mutation rates will be resolved before any third mutation rate can appear.

To standardize terminology, we use the term "strain" to refer to a subpopulation with a given mutation rate (e.g. "a strain of mutation rate u"). Multiple traits may be present within one strain. The frequency distribution of traits within each strain evolves over time, due to frequency-dependent selection.

6.3.2 Short-term evolution on the trait locus

We model frequency-dependent selection on the trait locus using replicator dynamics [132, 134, 135, 218, 270, 305], which describe the dynamics of competing traits in a large population by a system of ordinary differential equations.

Replicator dynamics

Replicator dynamics study the dynamics of the frequency vector $\mathbf{x} = (x_1, \ldots, x_n)$, where x_i is the frequency of trait i. The state space is the n-simplex

$$\triangle_n = \left\{ (x_1, \ldots, x_n) : x_i \geq 0, \sum_i x_i = 1 \right\}.$$

Frequency dependence is described by the fitness functions $f_i : \triangle_n \to \mathbb{R}$ for $i = 1, \ldots, n$, where $f_i(\mathbf{x})$ gives the fitness of trait i when the population state is

$\mathbf{x} \in \triangle_n$. We require that each f_i be smooth and positive for all $\mathbf{x} \in \triangle_n$.

Replicator dynamics are most often studied using linear f_i; however, none of the results of this work requires linearity.

When mutation is absent, trait dynamics are governed by the *replicator equations*

$$\dot{x}_i = x_i(f_i(\mathbf{x}) - \varphi), \qquad (6.2)$$

for $i = 1, \ldots, n$, with $\varphi = \sum_{i=1}^{n} x_i f_i$ denoting average population fitness. Equation 6.2 can also be written in vector form as

$$\dot{\mathbf{x}} = (F(\mathbf{x}) - \varphi I)\mathbf{x},$$

where $F(\mathbf{x})$ is the $n \times n$ diagonal matrix with ith diagonal entry $f_i(\mathbf{x})$.

REPLICATOR-MUTATOR DYNAMICS

Incorporating mutation of rate u into the replicator equation yields the *replicator-mutator equations* [114, 128, 221]:

$$\dot{x}_i = (1 - u)x_i f_i(\mathbf{x}) + \frac{u}{n}\varphi - \varphi x_i, \qquad (6.3)$$

or in vector form:

$$\dot{\mathbf{x}} = (Q_u F(\mathbf{x}) - \varphi I)\mathbf{x}. \qquad (6.4)$$

Above, the matrix Q_u describes mutation:

$$Q_u = \begin{pmatrix} 1 - \frac{n-1}{n}u & \frac{u}{n} & \cdots & \frac{u}{n} \\ \frac{u}{n} & 1 - \frac{n-1}{n}u & \cdots & \frac{u}{n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{u}{n} & \frac{u}{n} & \cdots & 1 - \frac{n-1}{n}u \end{pmatrix} = (1 - u)I + \frac{u}{n}\mathbf{1}.$$

The symbol 1 above denotes the square matrix of all ones. In Eq. 6.3, and in the remainder of this work, it is assumed that each mutation on the trait locus is equally likely to result in any one of the n traits. We will refer to the fixed points of Eq. 6.3 as mutation-selection equilibria.

6.3.3 Long-term evolution on the mutation rate locus

We study the long-term evolution of the mutation rate under the assumptions of the adaptive dynamics framework [70, 103, 131, 199]. This framework assumes that long-term evolution proceeds by a sequence of substitution events. In each such event, an invading strain successfully displaces the resident strain and becomes the new resident. Mutation (on the mutation rate locus, in our case) is assumed rare, so that the possibility of concurrent invasions by multiple strains can be disregarded.

Section 6.3.3 extends the replicator-mutator equations to competition between two strains of different mutation rates. Then in Eq. 6.3.3 we formally state our model for evolution on the mutation rate locus, inspired by adaptive dynamics. Items 6.3.3 and 6.3.3 introduce the concept of invasion fitness, and

discuss the extent to which it predicts the outcome of invasion events.

Competition between strains

Here we extend the replicator-mutator equations, Eq. 6.3, to the case of two competing strains: a resident strain of mutation rate u, and an invading strain of mutation rate u'. We denote the resident trait frequencies by $\mathbf{r} = (r_1, \ldots, r_n)$, and the invader trait frequencies by $\mathbf{z} = (z_1, \ldots, z_n)$. These frequencies are relative to the whole population; hence the state space is the simplex

$$\triangle_{2n} = \left\{ (\mathbf{r}, \mathbf{z}) : r_i \geq 0, z_i \geq 0, \sum_i (r_i + z_i) = 1 \right\}.$$

Evolutionary dynamics are described by the system of equations

$$
\begin{aligned}
\dot{r}_i &= (1 - u) r_i f_i(\mathbf{r} + \mathbf{z}) + \frac{u}{n} R \varphi_{\mathbf{r}} - \varphi r_i \\
\dot{z}_i &= (1 - u') z_i f_i(\mathbf{r} + \mathbf{z}) + \frac{u'}{n} Z \varphi_{\mathbf{z}} - \varphi z_i,
\end{aligned}
\tag{6.5}
$$

where R and Z, respectively, are the total frequencies of residents and invaders,

$$R = \sum_{i=1}^{n} r_i, \qquad Z = \sum_{i=1}^{n} z_i,$$

and

$$\varphi = \sum_{i=1}^{n}(r_i + z_i)f_i(\mathbf{r} + \mathbf{z}),$$

$$\varphi_{\mathrm{r}} = \frac{1}{R}\sum_{i=1}^{n} r_i f_i(\mathbf{r} + \mathbf{z}), \tag{6.6}$$

$$\varphi_{\mathrm{z}} = \frac{1}{Z}\sum_{i=1}^{n} z_i f_i(\mathbf{r} + \mathbf{z}),$$

are the average fitness functions of the whole population and the two strains, respectively.

Note that when restricted to either of the subsets

$$\triangle_{2n\,|\,Z=0} = \{(\mathbf{r}, \mathbf{z}) \in \triangle_{2n} | \mathbf{z} = 0\},$$

$$\triangle_{2n\,|\,R=0} = \{(\mathbf{r}, \mathbf{z}) \in \triangle_{2n} | \mathbf{r} = 0\}, \tag{6.7}$$

then Eq. 6.5 reduces to Eq. 6.3 (with $u := u'$ in the latter case).

For a given orbit of Eq. 6.5, we say that the invading (resp., resident) strain *fixates* if $Z(t) \to 1$, (resp., $R(t) \to 1$) as $t \to \infty$. The question of whether invaders or residents fixate depends, in general, on the initial conditions $(\mathbf{r}(0), \mathbf{z}(0)) \in \triangle_{2n}$ that define this orbit. It is also possible for neither strain to fixate, as we discuss in Eq. 6.3.3.

The process of mutation rate evolution

This section presents a model of mutation rate evolution as an iterative process. In each iteration, an invading strain arises, competes with the resident strain, and

either displaces the resident strain (becoming the new resident) or goes extinct. This process is motivated by the adaptive dynamics framework discussed in the beginning of Sect. 6.3.3.

This process depends on the following data:

- the fitness functions $f_i : \triangle_n \to \mathbb{R}_{\geq 0}$ for $i = 1, \dots, n$,

- the initial mutation rate $u_0 \in [0, 1]$,

- the initial trait distribution $\mathbf{x}_0 \in \triangle_n$,

- the initial frequency $a \in (0, 1)$ of invading strains,

- a probability measure M giving the probability that mutation on the mutation rate locus yields a change Δu in the mutation rate.

We require that M be supported on the interval $[-v, v]$ for some constant $v > 0$. We also require that M assign positive probability to the intervals $[-v, 0)$ and $(0, v]$, and zero probability to $\{0\}$. (In other words, mutation on the mutation rate locus can either increase or decrease the mutation rate, but does not leave it unchanged.)

To formally define this iterative process, let u_k and $\mathbf{x}_k$ be the mutation rate and trait distribution, respectively, after $k \geq 0$ iterations of this process. The next iteration then proceeds as follows:

1. A trait distribution $\mathbf{x}_k^* \in \triangle_n$ is sampled from the sojourn time measure $\sigma_{\mathbf{x}_k}$ associated with the flow on $\triangle_n$ defined by the replicator-mutator

equations, Eq. 6.3. (The point $\mathbf{x}_k^*$ represents the trait distribution state of the resident population when an invading strain appears.)

2. A vertex $\mathbf{v}_k^*$ of $\triangle_n$ is randomly chosen, with the probability that vertex i is chosen given by

$$(1 - u_k)(\mathbf{x}_k^*)_i \frac{f_i(\mathbf{x}_k^*)}{\varphi(\mathbf{x}_k^*)} + \frac{u_k}{n}.$$

(The vertex $\mathbf{v}_k^*$ represents the initial trait distribution within the invading strain. This distribution is monomorphic, reflecting the biological fact that the invading strain must be seeded by a single individual. The probability associated to vertex i represents the probability that a randomly chosen offspring from the current resident population will have trait i.)

3. A mutation rate change Δu_k is sampled from M. The mutation rate u_k' of the invading strain is defined as

$$u_k' := \begin{cases} u_k + \Delta u_k & \text{if } u_k + \Delta u_k \in [0, 1], \\ 0 & \text{if } u_k + \Delta u_k < 0, \\ 1 & \text{if } u_k + \Delta u_k > 1. \end{cases}$$

4. A point $(\mathbf{r}_k^*, \mathbf{z}_k^*) \in \triangle_{2n}$ is sampled from the sojourn time measure $\sigma_{((1-a)\mathbf{x}_k^*, a\mathbf{v}_k^*)}$ associated with the flow on $\triangle_{2n}$ defined by Eq. 6.5, with $u := u_k$ and $u' := u_k'$. (The point $(\mathbf{r}_k^*, \mathbf{z}_k^*)$ represents the state of the system after resident-invader competition is resolved.)

5. The $(k+1)$st mutation rate and trait distribution are defined as follows:

211

- if $\mathbf{r}_k^* = 0$, then $\mathbf{x}_{k+1} := \mathbf{z}_k^*$ and $u_{k+1} := u'$ (invaders replace

  residents),

- if $\mathbf{z}_k^* = 0$, then $\mathbf{x}_{k+1} := \mathbf{r}_k^*$ and $u_{k+1} := u$ (invaders go extinct).

Iterating this process yields a sequence $\{(u_k, \mathbf{x}_k)\}_{k=0}^{\infty}$. The sequence $\{u_k\}_{k=0}^{\infty}$ describes the long-term evolution of the mutation rate.

It is possible for neither $\mathbf{r}_k^* = 0$ nor $\mathbf{z}_k^* = 0$ to be satisfied in Step 4, as we discuss in Eq. 6.3.3. In this case, the process ceases to adequately describe mutation rate evolution, and is formally terminated.

Although this process is well-defined for any values of the required data, we will be most interested in the case where v and a are small. This means that we consider mutation on the mutation rate locus to be incremental, and invading strains to initially comprise only a small fraction of the population.

INVASION FITNESS

The invasion fitness of a strain is defined as its exponential growth rate when rare [70, 120, 198, 199, 252]. Invasion fitness is used in adaptive dynamics to predict the outcomes of invasion events.

For so long as the invading strain remains rare, the trait distribution among invaders has negligible effect on the trait dynamics within the resident strain. The resident trait dynamics are in this case closely approximated by the single-strain replicator-mutator equations, Eq. 6.3.

We can therefore quantify invasion fitness by assuming that $\mathbf{x} = \mathbf{x}(t)$ has dynamics given by Eq. 6.3, and that $\mathbf{z}$ evolves according to the time-dependent

linear equation,

$$\dot{\mathbf{z}} = \left(Q_{u'} F(\mathbf{x}(t)) - \varphi_{\mathrm{r}}(t)I \right) \mathbf{z}, \qquad (6.8)$$

where $\varphi_{\mathrm{r}}(t) = \sum_{i=1}^{n} f_i(\mathbf{x}(t))x_i(t)$ is the average fitness of the resident strain at time t. The state space for Eq. 6.8 is the set of all nonnegative vectors in $\mathbb{R}^n$.

Using this simplification, we define the invasion fitness of an invading strain with mutation rate u', which appears with internal trait distribution $\mathbf{x}'$ when the resident strain (of mutation rate u) has trait distribution $\mathbf{x}$, by

$$s_{u,\mathbf{x}}(u', \mathbf{x}') = \lim_{T \to \infty} \frac{1}{T} \int_0^T \frac{\dot{Z}(t)}{Z(t)} \, dt, \qquad (6.9)$$

if this limit exists (otherwise the invasion fitness is undefined). Above, $Z(t) = \sum_{i=1}^{n} z_i(t)$, the dynamics of $\mathbf{z}$ are given by Eq. 6.8 with initial condition $\mathbf{z}(0) = \mathbf{x}'$, and $\mathbf{x}$ gives the initial condition for Eq. 6.3. This extends previous definitions [70, 198, 252] to the case in which both the resident and invading strains contain a number of sub-populations (in our case, bearers of different traits), which are themselves evolving on a fast timescale. We caution, however, that time-averaged quantities of the form 6.9 may be undefined for orbits approaching heteroclinic or strange attractors; see Ref. [282] for a general discussion.

The invasion implies fixation conjecture

Positive invasion fitness does not necessarily imply that the invading strain will ultimately displace the resident strain. The two strains may evolve towards a

stable or dynamic coexistence.

This issue arises more generally in studying the adaptive dynamics of an arbitrary quantitative trait x. In general, coexistence of resident and invader strains is possible. However, it is reasonable to suppose that if the invasion fitness function $s_x(x')$ is nonsingular at $x' = x$, and if x' is sufficiently close to x, then positive invasion fitness implies fixation of the invading strain [67]. This proposition, known as "invasion implies fixation" (hereafter, IIF), was a long-standing conjecture until proofs were discovered by Dercole [66] and Geritz [102].

Neither proof applies directly to our situation of two loci evolving on different timescales, though it appears likely that the techniques used in these proofs might extend to the present case. For our purposes, we state IIF as a formal conjecture:

Conjecture 1 (Invasion Implies Fixation; IIF)**.** *Consider given fitness functions $f_i : \triangle_n \to \mathbb{R}$, mutation rate $u \in [0, 1]$, and trait distributions $\mathbf{x}, \mathbf{x}' \in \triangle_n$. Suppose that*

- *$\mathbf{x}$ lies in the basin of attraction of an attractor of the replicator-mutator equations, Eq. 6.3, with mutation rate u,*

- *there is a $\delta > 0$ such that $s_{u,\mathbf{x}}(u', \mathbf{x}')$ is strictly monotonic (either increasing or decreasing) in u', on the interval $u' \in (u - \delta, u + \delta) \cap [0, 1]$.*

Then there exists $\varepsilon > 0$ such that for all u' with $|u' - u| < \varepsilon$ and all $0 < a < \varepsilon$, invaders (resp., residents) fixate in Eq. 6.5 from the initial conditions

$$(\mathbf{r}(0), \mathbf{z}(0)) = ((1 - a)\mathbf{x}, a\mathbf{x}'),\qquad(6.10)$$

if $s_{u,\mathbf{x}}(u', \mathbf{x}') > 0$ (resp., $s_{u,\mathbf{x}}(u', \mathbf{x}') < 0$).

We will specify, in the remainder of this work, those of our results that depend on this conjecture.

6.4 RESULTS

Replicator-mutator dynamics can yield a wide variety of behaviors, including multiple attractors, limit cycles, and chaos [132, 194, 268]. Invasion fitness is difficult to compute analytically when the replicator-mutator dynamics are complex. We therefore limit our focus to simple cases for which analytical results can be obtained.

We first consider cases in which the mutation rate evolves downwards. Section 6.4.1 shows that such downward evolution can be expected whenever residents are at mutation-selection equilibrium. This downward evolution continues until either this equilibrium is lost or the mutation rate reaches zero.

We then investigate the upward evolution of mutation rates in Sect. 6.4.2. We derive conditions under which a mutator strain can invade and persist in a non-mutating resident population initially at equililbrium. We then show that if the replicator dynamics, Eq. 6.2, admits a globally attracting heteroclinic cycle between monomorphic states, the zero mutation rate is evolutionarily unstable.

Last, Sect. 6.4.3 explores cases in which mutators and non-mutators can stably coexist. This coexistence can occur when the replicator dynamics, Eq. 6.2, admits a stable polymorphic fixed point.

6.4.1 Downward evolution of mutation rate under stable mutation-selection equilibria

This section shows that mutation-selection equilibria lead to the downward evolution of mutation rates. This extends the classical observation that mutation rates evolve to zero in constant fitness landscapes, if no new beneficial mutations are available [163, 174, 175].

We start by showing how invasion fitness can be related to an eigenvalue characterizing the growth rate of the invading strain. Consider a resident strain of mutation rate $u \geq 0$, and suppose the trait distribution within this resident strain is at a fixed point $\tilde{\mathbf{x}}$ of the replicator-mutator equations, Eq. 6.3. We introduce the notation $\tilde{f}_i = f_i(\tilde{\mathbf{x}})$ for the fitness of trait i at this equilibrium, $\tilde{\varphi}_r = \sum_i \tilde{f}_i \tilde{x}_i$ for the average fitness of the resident strain, and $\tilde{F}$ for the diagonal matrix with entries $\tilde{f}_i$.

Lemma 1. *For any invading strain with trait distribution $\mathbf{x}' \in \triangle_n$ and mutation rate $u' > 0$, the invasion fitness of this strain is given by*

$$ s_{u,\tilde{\mathbf{x}}}(u', \mathbf{x}') = \hat{\varphi}_z - \tilde{\varphi}_r, $$

where $\hat{\varphi}_z$ denotes the largest eigenvalue of the matrix $Q_{u'}\tilde{F}$.

Proof. Since the resident strain is at equilibrium, Eq. 6.8 for invader dynamics is time-independent:

$$ \dot{\mathbf{z}} = \left(Q_{u'}\tilde{F} - \tilde{\varphi}_r I \right) \mathbf{z}. \tag{6.11} $$

The solution is given by

$$\mathbf{z}(t) = \exp\left(t(Q_{u'}\tilde{F} - \tilde{\varphi}_r I)\right)\mathbf{x}'Z_0.$$

Above, Z_0 is the (small) initial frequency of invaders.

Since $u' > 0$ and $\tilde{f}_i > 0$ for each i, then $Q_{u'}\tilde{F}$ has strictly positive entries. By the Perron-Frobenius theorem, $Q_{u'}\tilde{F}$ has a positive simple largest eigenvalue $\hat{\varphi}_z$, with associated strictly positive eigenvector $\hat{\mathbf{z}}$ (which we normalize so that $\sum_i \hat{z}_i = 1$). The matrix $\exp\left(Q_{u'}\tilde{F}\right)$ is also strictly positive, has Perron-Frobenius eigenvector $\hat{\mathbf{z}}$, and associated largest eigenvalue $e^{\hat{\varphi}_z}$. Applying the Perron-Frobenius theorem to $\exp\left(Q_{u'}\tilde{F}\right)$ further yields that

$$\lim_{t\to\infty} \frac{\exp\left(tQ_{u'}\tilde{F}\right)}{e^{t\hat{\varphi}_z}}\mathbf{x}' = k\hat{\mathbf{z}},$$

where $k \in \mathbb{R}$ is a constant. (Specifically, $k\hat{\mathbf{z}}$ is the Perron projection of $\mathbf{x}'$; that is, the eigenspace projection of $\mathbf{x}'$ onto $\hat{\mathbf{z}}$.)

We now consider the limit

$$\begin{aligned}
\lim_{t\to\infty} \frac{\mathbf{z}(t)}{\exp\left(t(\hat{\varphi}_z - \tilde{\varphi}_r)\right)} &= \lim_{t\to\infty} \frac{\exp\left(t(Q_{u'}\tilde{F} - \tilde{\varphi}_r I)\right)\mathbf{x}'Z_0}{\exp\left(t(\hat{\varphi}_z - \tilde{\varphi}_r)\right)} \\
&= Z_0 \lim_{t\to\infty} \frac{\exp\left(tQ_{u'}\tilde{F}\right)}{e^{t\hat{\varphi}_z}}\mathbf{x}' \qquad (6.12) \\
&= Z_0 k\hat{\mathbf{z}}.
\end{aligned}$$

Summing Eq. 6.12 over all components yields

$$\lim_{t\to\infty} \frac{Z(t)}{\exp\left(t(\hat{\varphi}_z - \tilde{\varphi}_r)\right)} = Z_\circ k. \tag{6.13}$$

Separately, left-multiplying both sides of Eq. 6.12 by $Q_{u'}\tilde{F} - \tilde{\varphi}_r I$ and comparing to Eq. 6.11 yields

$$\lim_{t\to\infty} \frac{\dot{\mathbf{z}}(t)}{\exp\left(t(\hat{\varphi}_z - \tilde{\varphi}_r)\right)} = Z_\circ k(Q_{u'}\tilde{F} - \tilde{\varphi}_r I)\hat{\mathbf{z}}$$
$$= Z_\circ k(\hat{\varphi}_z - \tilde{\varphi}_r)\hat{\mathbf{z}}. \tag{6.14}$$

Summing Eq. 6.14 over all components, we obtain

$$\lim_{t\to\infty} \frac{\dot{Z}(t)}{\exp\left(t(\hat{\varphi}_z - \tilde{\varphi}_r)\right)} = Z_\circ k(\hat{\varphi}_z - \tilde{\varphi}_r). \tag{6.15}$$

Finally, dividing Eq. 6.15 by Eq. 6.13 yields

$$\lim_{t\to\infty} \frac{\dot{Z}(t)}{Z(t)} = \hat{\varphi}_z - \tilde{\varphi}_r.$$

Comparing to the definition of invasion fitness, Eq. 6.9, we conclude that

$$s_{u,\tilde{\mathbf{x}}}(u',\mathbf{x}') = \hat{\varphi}_z - \tilde{\varphi}_r,$$

as desired. $\qquad\square$

Intuitively, the Perron-Frobenius eigenvector $\hat{\mathbf{z}}$ represents the invader trait distribution that is stable under mutation and selection when the invader

frequency is small. The corresponding eigenvalue $\hat{\varphi}_z$ equals the average fitness of invaders at this trait distribution. Thus Lemma 1 confirms that an invading strain succeeds when its average fitness exceeds that of the resident.

Using the above lemma, we now show that, when the resident trait distribution is at a mutation-selection equilibrium, an invading strain succeeds if and only if its mutation rate is lower than that of the resident population. The only exception is the non-generic case in which all traits are equally fit, in which case all invading strains are selectively neutral (we will discuss this case later in this section and in Sect. 6.4.3).

Theorem 1. *Consider a resident strain of mutation rate $u > 0$, and suppose the trait distribution within this resident strain is at a fixed point $\tilde{\mathbf{x}}$ of the replicator-mutator equations, Eq. 6.3. Then for all $\mathbf{x}' \in \triangle_n$ and $u' > 0$,*

- *If the $\tilde{f}_i = f_i(\tilde{\mathbf{x}})$ are all equal, $s_{u,\tilde{\mathbf{x}}}(u', \mathbf{x}') = 0$.*

- *If the $\tilde{f}_i$ are not all equal, $s_{u,\tilde{\mathbf{x}}}(u', \mathbf{x}') > 0 \Leftrightarrow u' < u$.*

Proof. We claim, and prove below, that $\frac{d\hat{\varphi}_z}{du'} \leq 0$, with equality if and only if all the $\tilde{f}_i$ are equal. Thus by Lemma 1, $s_{u,\tilde{\mathbf{x}}}(u', \mathbf{x}') = \hat{\varphi}_z - \tilde{\varphi}_r$ is decreasing in u' if the $\tilde{f}_i$ are not all equal, and constant in u' otherwise. We also note that by Eq. 6.3, the fixed point $\tilde{\mathbf{x}}$ is an eigenvector of the (strictly positive) matrix $Q_u \tilde{F}$, with eigenvalue $\tilde{\varphi}_r$. By the Perron-Frobenius theorem, $\tilde{\varphi}_r$ is the unique eigenvalue of $Q_u \tilde{F}$ whose associated eigenvector is nonnegative. In the case $u' = u$, this eigenvalue is, by definition, $\hat{\varphi}_z$. Thus for $u' = u$, $\hat{\varphi}_z = \tilde{\varphi}_r$, and therefore $s_{u,\tilde{\mathbf{x}}}(u, \mathbf{x}') = 0$. Combining this with the above-mentioned decreasing behavior of $s_{u,\tilde{\mathbf{x}}}(u', \mathbf{x}')$ in u' proves the

theorem.

To demonstrate the claim regarding $\frac{d\hat{\varphi}_z}{du'}$, we first observe that the positive and symmetric matrix $\tilde{F}^{1/2} Q_{u'} \tilde{F}^{1/2}$ has Perron-Frobenius eigenvalue $\hat{\varphi}_z$ and associated eigenvector $\tilde{F}^{1/2}\hat{\mathbf{z}}$. Applying a standard formula for the derivative of the eigenvalues of a real symmetric matrix [138, 186], we write

$$
\begin{aligned}
\frac{d\hat{\varphi}_z}{du'} &= \frac{\left(\tilde{F}^{1/2}\hat{\mathbf{z}}\right)^T \tilde{F}^{1/2} \left(\frac{d}{du'} Q_{u'}\right) \tilde{F}^{1/2} \left(\tilde{F}^{1/2}\hat{\mathbf{z}}\right)}{\left(\tilde{F}^{1/2}\hat{\mathbf{z}}\right)^T \left(\tilde{F}^{1/2}\hat{\mathbf{z}}\right)} \\
&= \frac{\hat{\mathbf{z}}^T \tilde{F} \left(\frac{d}{du'} Q_{u'}\right) \tilde{F}\hat{\mathbf{z}}}{\hat{\mathbf{z}}^T \tilde{F}\hat{\mathbf{z}}} \\
&= \frac{-\sum_{i=1}^{n} \left(\tilde{f}_i \hat{z}_i\right)^2 + \frac{1}{n}\left(\sum_{i=1}^{n} \tilde{f}_i \hat{z}_i\right)^2}{\hat{\mathbf{z}}^T \tilde{F}\hat{\mathbf{z}}}.
\end{aligned}
\tag{6.16}
$$

But Jensen's inequality implies that

$$
\left(\frac{\sum_i \tilde{f}_i \hat{z}_i}{n}\right)^2 \leq \frac{\sum_i (\tilde{f}_i \hat{z}_i)^2}{n}.
$$

Hence the numerator of Eq. 6.16 is ≤ 0, with equality if and only if the terms $\tilde{f}_i \hat{z}_i$ are equal for all i. In the latter case, we write $\tilde{F}\hat{\mathbf{z}} = a\mathbf{1}$, for some constant a, where $\mathbf{1}$ denotes the vector with each entry equal to one. We then have $\hat{\varphi}_z \hat{\mathbf{z}} = Q_{u'} \tilde{F}\hat{\mathbf{z}} = Q_{u'}(a\mathbf{1}) = a\mathbf{1}$. Thus the trait abundances $\hat{z}_i = a/\hat{\varphi}_z$ are constant over i, and the fitness functions $\tilde{f}_i$ must be constant over i as well. We also note that the denominator of Eq. 6.16 is positive in all cases, since $\tilde{F}$ is a diagonal matrix with strictly positive elements and is therefore positive definite. This proves the claim that $\frac{d\hat{\varphi}_z}{du'} \leq 0$, with equality if and only if the $\tilde{f}_i$ are all equal. $\qquad\square$

A similar result holds when the invading strain is non-mutating ($u' = 0$). However, in this case, the invasion fitness $s_{u,\tilde{\mathbf{x}}}(0, \mathbf{x}')$ depends on the initial distribution $\mathbf{x}'$ of traits within the invading strain, as we state in the following proposition.

Proposition 1. *Consider a resident strain of mutation rate $u > 0$, and suppose the trait distribution within this resident strain is at a fixed point $\tilde{\mathbf{x}}$ of the replicator-mutator equations, Eq. 6.3. For any $\mathbf{x}' \in \triangle_n$, $s_{u,\tilde{\mathbf{x}}}(0, \mathbf{x}') > 0$ if and only if $\max\{\tilde{f}_i : x'_i > 0\}$ exceeds $\tilde{\varphi}_r = \sum_i \tilde{f}_i \tilde{x}_i$.*

Proof. This follows immediately from observing that for $u' = 0$, Eq. 6.11 simplifies to

$$\dot{z}_i = z_i(\tilde{f}_i - \tilde{\varphi}_r),$$

for each $i = 1, \ldots, n$. $\qquad\square$

In particular, if $\mathbf{x}' \in \text{int } \triangle_n$, then the condition $\max\{\tilde{f}_i : x'_i > 0\} > \tilde{\varphi}_r$ of Proposition 1 becomes equivalent to the condition of Theorem 1 that the $\tilde{f}_i$ are not all equal. This is because for $\mathbf{x}' \in \text{int } \triangle_n$, $\max\{\tilde{f}_i : x'_i > 0\} > \tilde{\varphi}_r$ holds if and only if there is any i for which $\tilde{f}_i > \tilde{\varphi}_r$. Since $\tilde{\mathbf{x}}$ must also be in int $\triangle_n$ by the fact that $\tilde{\mathbf{x}}$ is a fixed point of Eq. 6.3 with $u > 0$, then $\tilde{\varphi}_r$ is a weighted average of the $\tilde{f}_i$ with positive weighting for each component. Thus $\tilde{f}_i > \tilde{\varphi}_r$ for some i if and only if the $\tilde{f}_i$ are not all equal.

At the other extreme, if $\mathbf{x}'$ is a vertex of $\triangle_n$ (as is always the case for invading strains that appear in the mutation rate evolution process described in Eq. 6.3.3), then the condition of Proposition 1 becomes simply $f_i > \tilde{\varphi}_r$, where i is the index

of the trait represented by $\mathbf{x}'$.

We remark that, according to Theorem 1, if the trait fitness functions $f_i(\tilde{\mathbf{x}})$ are all equal at a mutation-selection equilibrium $\tilde{\mathbf{x}}$, then all invading strains of different mutation rate are selectively neutral. This can be explained by observing that when all traits are equally fit, all mutations are themselves selectively neutral. This situation can only occur in very particular circumstances. Indeed, setting $f_i(\tilde{\mathbf{x}}) = \varphi$ for each i at a fixed point $\tilde{\mathbf{x}}$ of the replicator-mutator equations, Eq. 6.3, yields

$$0 = u\varphi \left(\frac{1}{n} - \tilde{x}_i \right),$$

for each i. The average fitness φ is positive since $f_i(\tilde{\mathbf{x}}) \geq 0$ for each i and $\sum_{i=1}^{n} f_i(\tilde{\mathbf{x}}) > 0$. Hence if all traits are equally fit at $\tilde{\mathbf{x}}$, then either $u = 0$ (mutation is absent) or $\tilde{\mathbf{x}}$ is located at the centroid $1/n = (1/n, \ldots, 1/n) \in \triangle_n$ (all traits are equally abundant). We explore the $u = 0$ case further in Sect. 6.4.3.

Our next result applies Theorem 1 to the mutation rate evolution process defined in Eq. 6.3.3. It shows that if, at some step k in the process, the resident trait distribution is at a *stable* mutation-selection equilibrium, and if v and a are sufficiently small, then not only will the mutation rate at step $k + 1$ be less than or equal to the mutation rate at step k, but the new resident trait distribution will again be at a stable mutation-selection equilibrium.

Theorem 2. *In the process of mutation rate evolution outlined in Eq. 6.3.3, suppose that for some $k \geq 0$, $u_k > 0$ and that $\mathbf{x}_k \neq 1/n$ is a stable hyperbolic fixed point of the replicator-mutator equations, Eq. 6.3, with mutation rate u_k. Then if v and a are sufficiently small and IIF is assumed, then $u_{k+1} \leq u_k$ and $\mathbf{x}_{k+1}$ is a stable hyperbolic*

fixed point of Eq. 6.3 with mutation rate u_{k+1}.

Proof. We begin by noting that, as a consequence of the implicit function theorem, there is some relatively open neighborhood $U \subset [0, 1]$ of u_k, and a differentiable curve $\gamma : U \to \triangle_n$ such that for each $u \in U$, $\gamma(u)$ is a stable hyperbolic fixed point of Eq. 6.3 with mutation rate u.

We will prove that, for the mutational step Δu_k sampled in Step 3,

(a) If $\Delta u_k < 0$ and $u'_k > 0$, then $u_{k+1} = u'_k$ with probability one.

(b) If $\Delta u_k < 0$ and $u'_k = 0$ then $u_{k+1} = 0$ with positive probability (otherwise $u_{k+1} = u_k$).

(c) If $\Delta u_k > 0$ then $u_{k+1} = u_k$ with probability one.

(d) In all cases, $\mathbf{x}_{k+1} = \gamma(u_{k+1})$.

The claims of the theorem follow from the above statements.

We note that since $\mathbf{x}_k$ is a fixed point of Eq. 6.3, then $\mathbf{x}^*_k = \mathbf{x}_k$. By the remarks following Proposition 1, since $\mathbf{x}^*_k \neq 1/n$ and $u_k > 0$, then the fitness functions $f_i(\mathbf{x}^*_k)$ are not all equal. Therefore, in the case $u'_k > 0$, Theorem 1 implies that $s_{u_k, \mathbf{x}^*_k}(u'_k, \mathbf{v}^*_k) > 0$ if and only if $\Delta u_k < 0$, regardless of which vertex $\mathbf{v}^*_k$ is sampled in Step 2. In the case $u'_k = 0$, Proposition 1 implies that there is at least one vertex $\mathbf{v}$ of $\triangle_n$ such that $s_{u_k, \mathbf{x}^*_k}(u'_k, \mathbf{v}^*_k) > 0$, and this vertex has nonzero probability of being sampled in Step 2 since $\mathbf{x}^*_k \in \text{int } \triangle_n$. In summary,

- If $\Delta u_k < 0$ and $u'_k > 0$, then $s_{u_k, \mathbf{x}^*_k}(u'_k, \mathbf{v}^*_k) > 0$ with probability one.

- If $\Delta u_k < 0$ and $u'_k = 0$, then $s_{u_k, \mathbf{x}^*_k}(u'_k, \mathbf{v}^*_k) > 0$ with positive probability.

- If $\Delta u_k > 0$ then $s_{u_k, \mathbf{x}_k^*}(u_k', \mathbf{v}_k^*) < 0$ with probability one.

IIF now guarantees that (a), (b), and (c) are satisfied as long as v and a are sufficiently small.

To verify claim (d), we invoke the Tube Theorem [104]. Applied to Eq. 6.3 and Eq. 6.5, the Tube Theorem guarantees the following: Let $\tilde{\mathbf{x}}$ be a stable hyperbolic fixed point of Eq. 6.3. Then for each $\varepsilon > 0$ there exists a $\delta > 0$ such that if $|u' - u| < \delta$ and $|\mathbf{r}(0) + \mathbf{z}(0) - \tilde{\mathbf{x}}| < \varepsilon$, then $|\mathbf{r}(t) + \mathbf{z}(t) - \tilde{\mathbf{x}}| < \varepsilon$ for all $t \geq 0$ under the dynamics of Eq. 6.5. In words, the trait distribution in the whole population (residents and invaders combined) stays close to the fixed point $\tilde{\mathbf{x}}$, as long as the difference in mutation rates $|u' - u|$ is sufficiently small.

Consider the ω-limit set $\omega_k := \omega((1 - a)\mathbf{x}_k^*, a\mathbf{v}_k^*) \subset \triangle_{2n}$ associated with the flow defined by Eq. 6.5 with $u := u_k$ and $u' := u_k'$ (that is, the flow considered in Step 4). IIF and the Tube Theorem jointly imply that for each $\varepsilon > 0$ there exists $\delta > 0$ such that if $v < \delta$ and $a < \delta$, then

$$
\begin{cases}
\omega_k \subset \left\{ (\mathbf{r}, \mathbf{z}) \in \triangle_{2n} : \mathbf{r} = 0 \text{ and } |\mathbf{z} - \gamma(u_k)| < \varepsilon \right\} & \text{if } s_{u_k, \mathbf{x}_k^*}(u_k', \mathbf{v}_k^*) > 0, \\[2mm]
\omega_k \subset \left\{ (\mathbf{r}, \mathbf{z}) \in \triangle_{2n} : \mathbf{z} = 0 \text{ and } |\mathbf{r} - \gamma(u_k)| < \varepsilon \right\} & \text{if } s_{u_k, \mathbf{x}_k^*}(u_k', \mathbf{v}_k^*) < 0.
\end{cases}
$$
$$(6.17)$$

Let $\bar{\omega}_k \subset \triangle_n$ denote the image of ω_k under the identifications $\triangle_n \leftrightarrow \triangle_{2n|R=0}$ or $\triangle_n \leftrightarrow \triangle_{2n|Z=0}$, in the cases $s_{u_k, \mathbf{x}_k^*}(u_k', \mathbf{v}_k^*) > 0$ and $s_{u_k, \mathbf{x}_k^*}(u_k', \mathbf{v}_k^*) < 0$, respectively. By Eq. 6.17, $\bar{\omega}_k$ is a subset of the open ball $B(\gamma(u_k), \varepsilon)$ of radius ε around $\gamma(u_k)$. Additionally, since ω_k is an invariant set of Eq. 6.5, then $\bar{\omega}_k$ is an invariant set of Eq. 6.3. Since $\gamma(u_{k+1})$ is a nondegenerate fixed point of Eq. 6.3,

there is some neighborhood of $U \subset \triangle_n$ of $\gamma(u_{k+1})$ such that $\{\gamma(u_{k+1})\}$ is the only invariant subset of U. The continuity of γ implies that $B(\gamma(u_k), \varepsilon) \subset U$ for sufficiently small ε. Consequently there exists a $\delta > 0$ such that $|\Delta u_k| < \delta$ implies $\bar{\omega}_k \subset U$. Since $\bar{\omega}_k$ is an invariant set of Eq. 6.3 and a subset of U, then $\bar{\omega}_k$ must consist only of the single point $\gamma(u_{k+1})$. Finally, since sojourn time distributions are concentrated on the corresponding ω-limit sets, the point $\mathbf{x}_{k+1}$ assigned in Step 5 can only be $\gamma(u_{k+1})$, as long as ν is sufficiently small. This proves claim (d), completing the proof. $\qquad\square$

One can guarantee that the downward evolution of mutation rate will continue for any desired finite number of steps by setting ν and a sufficiently small and applying Theorem 2 inductively. However, it is not possible in general to guarantee that this downward evolution will continue indefinitely or until the zero mutation rate is reached. This is because the bounds on ν and a needed to guarantee that $u_{k+1} \leq u_k$ and $\mathbf{x}_{k+1} = \gamma(u_{k+1})$ are not necessarily uniform in u_k. In less technical language, for any fixed values of ν and a, the mutation rate may evolve downwards to the point that the mutation-selection equilibrium disappears or loses stability, at which point upward mutation rate evolution is again possible.

6.4.2   EVOLUTION OF POSITIVE MUTATION RATES

Having identified conditions under which mutation rates evolve downwards, we now turn our attention to the upward evolution of mutation rates. In Sect. 6.4.2, Theorem 3 and Corollary 1 give conditions under which a strain of positive

mutation rate can invade a non-mutating resident population which is at a (possibly unstable) selection-induced equilibrium. Theorem 4 shows that, in this case, successful invasion implies persistence of the mutator strain.

Equations 6.4.2 focuses on the case where the replicator equations, Eq. 6.2, admit a globally attracting heteroclinic cycle. In this case, the zero mutation rate is evolutionarily unstable.

Conditions for mutator invasion of non-mutating residents at equilibrium

In this section we ask whether a strain of positive mutation rate u' can invade a non-mutating resident strain at equilibrium. As discussed in the proof of Theorem 1 (and using the notation defined there), the invasion fitness is equal to the largest eigenvalue of $Q_{u'}\tilde{F} - \tilde{\varphi}_r I$. The following theorem gives necessary and sufficient conditions for this invasion fitness to be positive, avoiding the need for an eigenvalue calculation:

Theorem 3. *Let $\tilde{\mathbf{x}}$ be a fixed point of the replicator equations, Eq. 6.2, and let $\tilde{f}_i = f_i(\tilde{\mathbf{x}})$ and $\tilde{\varphi}_r = \sum_{i=1}^n \tilde{f}_i \tilde{x}_i$. Then for any $\mathbf{x}' \in \triangle_n$ and $u' > 0$, $s_{0,\tilde{\mathbf{x}}}(u', \mathbf{x}') > 0$ if and only if either*

(a) $(1 - u')\tilde{f}_i \geq \tilde{\varphi}_r$ for some i, or

(b) $\displaystyle\sum_{i=1}^n \frac{\tilde{\varphi}_r}{\tilde{\varphi}_r - (1 - u')\tilde{f}_i} > \frac{n}{u'}.$

Proof. Let the matrix $\tilde{F}$, the Perron-Frobenius eigenvalue $\hat{\varphi}_z$ of $Q_{u'}\tilde{F}$, and the associated eigenvector $\hat{\mathbf{z}}$ be defined as in Sect. 6.4.1. The eigenvector equation

$Q_{u'}\tilde{F}\hat{\mathbf{z}} = \hat{\varphi}_z\hat{\mathbf{z}}$ can be rewritten as

$$\frac{n}{u'}\hat{z}_i = \frac{\hat{\varphi}_z}{\hat{\varphi}_z - (1 - u')\tilde{f}_i}, \tag{6.18}$$

for each i. Summing over all i yields

$$\frac{n}{u'} = \sum_{i=1}^{n} \frac{\hat{\varphi}_z}{\hat{\varphi}_z - (1 - u')\tilde{f}_i}. \tag{6.19}$$

Equation 6.18 and the positivity of $\hat{\mathbf{z}}$ and $\hat{\varphi}_z$ imply that $(1 - u')\tilde{f}_i < \hat{\varphi}_z$ for all i. Thus, if $(1 - u')\tilde{f}_i \geq \tilde{\varphi}_r$ for any i, then $\hat{\varphi}_z > \tilde{\varphi}_r$ and hence $s_{o,\tilde{\mathbf{x}}}(u', \mathbf{x}') > 0$ by Lemma 1. This proves that condition (a) is sufficient.

We now restrict to the case that condition (a) is false (that is, $(1 - u')\tilde{f}_i < \tilde{\varphi}_r$ for all i), and show that in this case condition (b) is necessary and sufficient. Consider the real-valued function

$$g(y) = \sum_{i=1}^{n} \frac{y}{y - (1 - u')\tilde{f}_i}.$$

As long as $(1 - u')\tilde{f}_i < y$ for all i, $g(y)$ is monotone decreasing in y. Equation 6.19 implies that $g(\hat{\varphi}_z) = n/u$; hence, $\hat{\varphi}_z > \tilde{\varphi}_r$ if and only if $g(\tilde{\varphi}_r) > n/u$, which is equivalent to condition (b). Since $\hat{\varphi}_z > \tilde{\varphi}_r$ is equivalent to $s_{o,\tilde{\mathbf{x}}}(u', \mathbf{x}') > 0$, this proves the theorem. $\qquad\square$

We can also ask whether there exists *any* strain of positive mutation rate that can invade non-mutating residents at the equilibrium $\tilde{\mathbf{x}}$. This question has a much simpler answer: if there is any trait whose fitness is larger than the average

resident fitness, then there exists a mutator strain that can invade. We state this formally as:

Corollary 1. *Let $\tilde{\mathbf{x}}$ be a fixed point of the replicator equations, Eq. 6.2. Then $\tilde{\hat{f}}_i > \tilde{\varphi}_r$ for some $i \in \{1, \ldots, n\}$ if and only if there exists $u' > 0$ such that $s_{0,\tilde{\mathbf{x}}}(u', \mathbf{x}') > 0$ for all $\mathbf{x}' \in \triangle_n$.*

Proof. The "only if" direction follows immediately from Theorem 3. For the "if" direction, suppose $\tilde{\hat{f}}_i \leq \tilde{\varphi}_r$ for all i. Then $\hat{\varphi}_z = \sum_i \tilde{\hat{f}}_i \hat{z}_i$ and $\hat{\mathbf{z}} \in \triangle_n$ imply that $\hat{\varphi}_z \leq \tilde{\varphi}_r$, and hence $s_{0,\tilde{\mathbf{x}}}(u', \mathbf{x}') \leq 0$ for all $\mathbf{x}' \in \triangle_n$. $\square$

The IIF conjecture does not necessarily apply to the scenarios considered in Theorem 3 and Corollary 1. In particular, if the non-mutating resident strain is monomorphic (that is, if $\tilde{\mathbf{x}} = \mathbf{v}$ is a vertex of $\triangle_n$) and $s_{0,\mathbf{v}}(u', \mathbf{v}) > 0$, then the monotonicity condition of the IIF conjecture is not satisfied: there is no $\delta > 0$ for which $s_{0,\mathbf{v}}(u', \mathbf{v})$ is monotonic in u' for $u' \in [0, \delta)$. This follows from observing that (a) $s_{0,\mathbf{v}}(0, \mathbf{v}) = 0$, (b) $s_{0,\mathbf{v}}(u', \mathbf{v}) > 0$ for all sufficiently small $u' > 0$ by Theorem 3, and (c) $s_{0,\mathbf{v}}(u', \mathbf{v})$ is decreasing in u' for $u' \in (0, 1]$ by the proof of Theorem 1.

We therefore cannot use IIF to argue that a mutator strain that successfully invades a monomorphic non-mutating resident strain (in the sense of having positive invasion fitness) will eventually rise to fixation. We can, however, prove a weaker result: if $s_{0,\mathbf{v}}(u', \mathbf{v}) > 0$, then an invading strain of mutation rate u' will persist indefinitely. We state this formally in the following theorem:

Theorem 4. *Let* $\mathbf{v}$ *be a vertex of* $\triangle_n$ *and suppose* $s_{0,\mathbf{v}}(u', \mathbf{v}) > 0$ *for some* $u' > 0$.
Then, for the orbit of the two-strain replicator-mutator equation, Eq. 6.5, with $u = 0$
and initial conditions $(\mathbf{r}(0), \mathbf{z}(0)) = ((1 - a)\mathbf{v}, a\mathbf{v})$, *the resident strain does not
fixate.*

Proof. Without loss of generality, suppose $\mathbf{v} = (1, 0, \ldots, 0)$. Note that for $i \neq 1$,
$r_i(t) = 0$ for all times $t \geq 0$. We can therefore disregard the coordinates
$r_2, \ldots, r_n$ and rewrite Eq. 6.5 as

$$\dot{r}_1 = r_1\big(f_1(\mathbf{r} + \mathbf{z}) - \varphi\big)$$
$$\dot{z}_i = (1 - u')z_i f_i(\mathbf{r} + \mathbf{z}) + \frac{u'}{n}Z\varphi_{\mathbf{z}} - \varphi z_i, \qquad (6.20)$$

for $i = 1, \ldots, n$.

Assume for contradiction that the resident strain fixates; then the orbit in
question must converge to the fixed point $(r_1, \mathbf{z}) = (1, 0)$. By the
Hartman-Grobman theorem, the dynamics of Eq. 6.20 are, in a neighborhood of
this fixed point, conjugate to the dynamics of the linearized system

$$\dot{r}_1 = \tilde{\varphi}(1 - r_1) - \sum_{i=1}^{n} \tilde{f}_i z_i$$
$$\dot{\mathbf{z}} = (Q_{u'}\tilde{F} - \tilde{\varphi}I)\mathbf{z}. \qquad (6.21)$$

Above, $\tilde{f}_i$, $\tilde{\varphi}$ and $\tilde{F}$ are the values of f_i, φ and F, respectively, at the fixed point
$(r_1, \mathbf{z}) = (1, 0)$.

From any initial conditions $(r_1(0), \mathbf{z}(0))$ with $\mathbf{z}(0) \neq 0$, the dynamics of

Eq. 6.21 will converge to $(r_1, \mathbf{z}) = (1, 0)$ if and only if the dynamics of

$$\dot{\mathbf{z}} = (Q_{u'}\tilde{F} - \tilde{\varphi}I)\mathbf{z},$$

with initial condition $\mathbf{z}(0)$, converge to $\mathbf{z} = 0$. By the reasoning used in the first part of the proof of Theorem 1, this can only occur if the Perron-Frobenius eigenvalue $\hat{\varphi}_{\mathbf{z}}$ of $Q_{u'}\tilde{F}$ is less than or equal to than $\tilde{\varphi}$, which contradicts $s_{0,\mathbf{v}}(u', \mathbf{v}) > 0$. $\qquad\square$

REMARKS ON THE SPECIAL CASE OF A HETEROCLINIC CYCLE

Another interesting case occurs when the replicator dynamics admit a globally attracting heteroclinic cycle, the fixed points of which are vertices of $\triangle_n$. Such cycles may occur in "rock-paper-scissors" systems, in which each trait is bested by another [130, 132].

Let the fitness functions f_i be such that the replicator equations, Eq. 6.2, admit a globally attracting heteroclinic cycle whose fixed points are vertices of $\triangle_n$. In this case, the invariance of the sojourn time measure implies that for each $\mathbf{x} \in \triangle_n$, $\sigma_{\mathbf{x}}$ is concentrated entirely on the fixed points of this cycle, i.e. the vertices of $\triangle_n$.[1]

Consider the mutation rate evolution process defined in Eq. 6.3.3, and suppose

---

[1]Perhaps contrary to intuition, the value of $\sigma_{\mathbf{x}}$ is in general undefined on the *individual* fixed points of such a heteroclinic cycle. This is because, if $\mathbf{v}$ is a fixed point on this cycle, and if all fixed points are hyperbolic—the generic case—then $\sigma_{\mathbf{x}}(\{\mathbf{v}\})$ will not converge according to the limit definition 6.1 [101, 282]. Thus, according to our definition, the singleton $\{\mathbf{v}\}$ is a non-measurable subset, as is any proper subset of the vertex set. However, this does not affect our current argument, which requires only that the entire vertex set is assigned probability one.

that $u_k = 0$ for some $k \geq 0$. By the above remarks, the state $\mathbf{x}_k^*$ sampled in Step 1 will be a vertex of $\triangle_n$. By Theorem 4, the invading strain of mutation rate $u_k' > 0$ that arises in Step 2 will persist indefinitely, so long as u_k' is small enough that the conditions of Theorem 3 are satisfied. In this case, either the invading strain will fixate, or neither strain will fixate and the process will be formally terminated.

In short, the zero mutation rate is evolutionarily unstable for such fitness functions. Any non-mutating strain is vulnerable to invasion by strains of positive but sufficiently small mutation rate. This result is intuitive: in rock-paper-scissors type scenarios, any pure strategy is bested by another. Thus if a non-mutating resident population converges to a monomorphic state, there must be at least one beneficial mutation available, upon which mutator alleles may hitchhike.

This result raises the question of how mutation rates ultimately evolve under the process outlined in Eq. 6.3.3, in the case that the fitness functions f_i induce an attracting heteroclinic cycle in the replicator system, Eq. 6.2. In our preliminary simulations, using n-strategy games of rock-paper-scissors type, we have found that mutation rates tend to converge under this process to a stable intermediate rate. We will report these results in full in forthcoming work.

6.4.3   COEXISTENCE BETWEEN MUTATING AND NON-MUTATING STRAINS

An interesting situation occurs when the replicator equations, Eq. 6.2, admit an equilibrium $\tilde{\mathbf{x}}$ at which all types are present (that is, $\tilde{\mathbf{x}} \in \text{int } \triangle_n$). In this case, Eq. 6.2 implies that all trait fitness functions $f_i(\tilde{\mathbf{x}})$ are equal.

Stable coexistence between mutating and non-mutating strains is possible in

this case. To see this, consider the two-strain replicator-mutator equation, Eq. 6.5, and suppose that only the invading strain has positive mutation rate (i.e. $u = 0, u' > 0$). Then the point $(\mathbf{r}, \mathbf{z}) \in \triangle_{2n}$ is a fixed point of Eq. 6.5 whenever $\mathbf{z}/Z = 1/n$ (that is, all traits are equally present in the invading strain), and $\mathbf{r} + \mathbf{z} = \tilde{\mathbf{x}}$. Note that the set of all such fixed points of Eq. 6.5 comprise a one-dimensional submanifold of $\triangle_{2n}$; therefore these fixed points are all non-hyperbolic.

The linear snowdrift game (also known as the hawk-dove game) is a simple example admitting such stable coexistence. In this game there are two traits, cooperator (C) and defector (D), and the fitness functions of the types are given by

$$\begin{pmatrix} f_C(\mathbf{x}) \\ f_D(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} R & S \\ T & P \end{pmatrix} \begin{pmatrix} x_C \\ x_D \end{pmatrix}, \tag{6.22}$$

with $T > R > S > P > 0$. This game admits a mixed evolutionarily stable state—that is, a stable hyperbolic fixed point of the replicator equations, Eq. 6.2—with trait frequencies

$$\tilde{x}_C = \frac{S - P}{T - R + S - P}, \qquad \tilde{x}_D = \frac{T - R}{T - R + S - P}.$$

Then the fixed points of Eq. 6.5 that describe coexistence between a resident

non-mutating strain and an invading mutator strain are given by

$$\tilde{r}_C = \frac{S - P}{T - R + S - P} - \frac{Z}{2} \qquad \tilde{z}_C = \frac{Z}{2}$$

$$\tilde{r}_D = \frac{T - R}{T - R + S - P} - \frac{Z}{2} \qquad \tilde{z}_D = \frac{Z}{2}, \tag{6.23}$$

valid for any value of Z for which all of the above quantities are positive.

This raises the question of whether an invading strain of positive mutation rate will generally coexist indefinitely with a non-mutating resident strain in such situations, or whether one of the two strains will generally fixate. The IIF conjecture is uninformative on this question, since if $\tilde{\mathbf{x}}$ is a fixed point of Eq. 6.2 for which all traits are present, then all traits must be equally fit according to Eq. 6.2, and thus Theorem 1 implies that $s_{0,\tilde{\mathbf{x}}}(u', \mathbf{x}') = 0$, for any $u' \in [0, 1]$, $\mathbf{x}' \in \triangle_n$. So such an invading strain will have zero invasion fitness.

To investigate this question, we simulated Eq. 6.5 numerically in the case of the snowdrift game discussed above (see Fig. 6.4.1). We found that whenever both traits are present in the initial non-mutating resident population ($r_C(0) > 0$ and $r_D(0) > 0$), the dynamics of invading mutators and resident non-mutators converge to a fixed point of the form 6.23. In this system, the asymptotic frequency of invaders depends on the initial conditions. This result is robust to variations in the payoff matrix, initial strain abundances, initial trait distributions, and the mutation rate of the invading strain.

Figure 6.4.1: Numerical simulation of the two-strain replicator-mutator equations, Eq. 6.5, for the snowdrift game, varying the initial proportion of trait C (cooperation) within the non-mutating resident strain $(r_C(0)/(1-a))$. Panels show features of the stable equilibrium $(\tilde{\mathbf{r}}, \tilde{\mathbf{z}})$. (a) Equilibrium proportion of strategy C within the resident strain $(\tilde{r}_C/\tilde{R}$, dashed line and circle), within the invading mutator strain $(\tilde{z}_C/\tilde{Z}$, dotted line and triangle), and within the population overall $(\tilde{r}_C + \tilde{z}_C$, solid line). The value $\tilde{x}_C = \tilde{r}_C + \tilde{z}_C = 1/3$ is the stable fixed point of the replicator equations, Eq. 6.2. (b) Equilibrium frequency of the invading mutator strain $(\tilde{Z}$, solid line and square). The open/closed symbols in both panels indicate the discontinuity at $r_C(0)/(1-a) = 1$; at this point, the non-mutating resident strain contains only trait C and goes extinct, while the invading strain converges to the mutation-selection equilibrium of the single-strain replicator-mutator equations, Eq. 6.3. Parameters: game payoffs $R = 3$, $S = 2$, $T = 5$, $P = 1$; mutation rates $u = 0$, $u' = 0.01$; initial invader frequency $a = 0.1$; initial invader trait distribution $z_C(0) = 0$, $z_D(0) = a$ (i.e., the invading strain starts with only trait D, defection).

6.5   DISCUSSION

Our work investigates how frequency-dependent selection within a population affects the evolution of mutation rates. The framework we present in Sect. 6.3 for studying this question merges the fields of replicator dynamics (including mutation) and adaptive dynamics. This framework is simple to state, but complex in the variety of behaviors it can yield.

Our work shows that frequency-dependent selection can induce either upward or downward selective pressure on the mutation rate. Specifically, when frequency dependence leads to a stable equilibrium trait distribution, mutation rates evolve downwards (possibly converging to zero); but when these dynamics converge to a heteroclinic cycle, the zero mutation rate is evolutionarily unstable. Our findings extend the general understanding that stable (resp., unstable) fitness landscapes lead to downward (resp., upward) mutation rate evolution [85, 107, 141, 163, 174, 175, 294, 295, 299, 307, 313, 326] to the case where the fitness landscape depends on the population itself.

Our work also highlights an interesting scenario of coexistence between mutating and non-mutating strains. While previous work [69, 148] has considered coexistence of mutation rates due to mutation-selection balance between them, we show that such coexistence can occur in the absence of mutation between the mutation rates. This coexistence occurs only at equilibria of frequency-dependent selection, in which all traits have equal fitness. Since it relies on one strain having an exactly zero mutation rate—a biologically unlikely

scenario—the relevance of this result remains to be shown.

Of course, the replicator and replicator-mutator equations can yield a wide variety of dynamics beyond the cases examined here [132, 194, 268]. It will be especially interesting to see whether cyclical or chaotic dynamics lead to the upward evolution of the mutation rate. Such questions may be beyond the reach of analytical treatment, and are perhaps best studied using simulations.

Of particular theoretical interest is the identification of evolutionarily stable mutation rates, which cannot be invaded by strains of either lower or higher mutation rate. This extends the concept of an evolutionarily stable strategy (ESS) [196] to mutation rate evolution. More generally, adaptive dynamics can lead to a wide variety of behaviors, including limit cycles, chaos, and evolutionary branching [67, 103]. Identifying which of these behaviors can occur for mutation rate evolution, and then relating these behaviors back to the underlying trait dynamics, is a promising avenue for future research.

Our framework makes use of the mathematical notion of invariant measures [86, 262]. Invariant measures are a formal mathematical representation of the ergodic principle: that the long-time behavior of a dynamical system can be understood from a statistical point of view by studying an invariant probability distribution over the set of possible states. Ergodic theory has long been applied to many areas of physics, but has seen only sporadic application in ecology and evolution [70, 75, 76, 252]. Invariant measures arise naturally when considering two dynamical processes on different timescales. An invariant measure can serve as probability distribution for the state of a fast process (e.g., demographic

dynamics) when a rare event (e.g., mutation, environmental change) occurs.

Finally, our work suggests new empirical hypotheses. A number of species, including *Escherichia coli* [157, 165], the common side-blotched lizard *Uta stansburiana* [289, 290], and the yeast *Saccharomyces cerevisiae* [231], show nontransitive competition between bearers of different traits. We have shown in Eq. 6.4.2 that in some cases, such nontransitive competition can induce upward evolutionary pressure on the mutation rate. We therefore conjecture that in such species, the mutation rate on loci coding for these traits may be abnormally high.

6.6 Acknowledgements

6.7 Manuscript information

6.7.1 Previously published as

This manuscript appeared in [6]:

Allen B, Rosenbloom DIS (2012). Mutation rate evolution in replicator dynamics. *Bull. Math. Biol.*, 74(11): 2650–2675.

7
Appendix

7.1 Supplementary tables referenced in Chapter 2

Table 7.1.1: Pharmacokinetic and pharmacodynamic parameters for anti-HIV drugs used in the study [265]

Class	Drug	IC_{50} (μMol)	slope	C_{max} (μMol)	half-life (hrs)	dosing (d^{-1})
NRTI	3TC	0.0298	1.15	15.3	10.0	2
	ABC	0.0381	0.93	10.5	21.0	2
	AZT	0.1823	0.85	4.5	8.5	2
	d4T	0.5524	1.13	2.3	3.5	2
	ddI	0.1795	1.07	39.4	18.0	1
	FTC	0.0079	1.20	7.3	39.0	1
	TDF	0.0561	0.97	1.1	60.0	1
NNRTI	EFV	0.0035	1.69	12.9	35.8	1
	ETV	0.0050	1.75	1.6	41	2
	NVP	0.0490	1.49	25.2	21.5	1
PI	ATV	0.0150	2.90	3.3	6.5	1
	ATV/r	0.0150	2.90	6.3	8.6	1
	DRV/r	0.0265	3.55	14.8	15.0	2
	IDV	0.0550	4.5	10.9	1.8	3
	IDV/r	0.0550	4.5	12.5	3.5	2
	LPV/r	0.0380	2.1	15.6	9.9	2
	NFV	0.2360	1.88	5.1	4.0	3
	SQV	0.0550	3.74	3.1	4.3	3
	SQV/r	0.0550	3.74	7.9	4.3	2
	TPV/r	0.2500	2.55	77.6	6.0	2
II	EVG	0.0280	0.94	1.7	8.6	1
	RAL	0.0150	1.03	4.0	10.0	2
FI	ENF	0.0349	1.60	1.1	3.8	2

Table 7.1.2: Parameters for all single-point mutations considered in the study [265]

Class	Mutation	Cost (s)	u	Drug	ρ	σ
NRTI	K65R	0.41	1.1×10^{-5}	3TC	61	-0.19
				ABC	47	0.01
				ddI	20	-0.09
				FTC	29	-0.04
				TDF	43	0.00
	M184V	0.46	1.1×10^{-5}	3TC	963	-0.58
				ABC	9.5	-0.44
				AZT	0.28	-0.03
				ddI	9.5	-0.21
				FTC	1186	-0.49
				TDF	3.0	-0.27
	M41L	0.17	1.3×10^{-6}	AZT	2.2	0.07
				d4T	1.0	0.07
	T215Y	0.05	*	AZT	3.1	-0.34
				d4T	1.08	-0.12
NNRTI	G190S	0.79	2.2×10^{-5}	EFV	70	-0.40
				NVP	237	-0.34
	K101P	0.7	*	ETV	5.00	-0.27
	K103N	0.3	1.5×10^{-6}	EFV	85	-0.17
				NVP	94	-0.15
	Y181C	0.26	1.1×10^{-5}	EFV	2.6	-0.11
				ETV	11	-0.26
				NVP	234	-0.40
	Y181I	0.44	*	ETV	100	-0.37
				NVP	1309	-0.50

* Indicates mutation that requires two nucleotide changes;
mutation rate depends on prevalence of intermediate states.

Table 7.1.3: Parameters for all single-point mutations considered in the study (Cont'd) [265]

Class	Mutation	Cost (s)	u	Drug	ρ	σ
PI	D30N	0.27	5.5×10^{-5}	NFV	2.3	-0.29
	G48V	0.45	8.5×10^{-7}	SQV	2.0	-0.23
	I47A	0.9	*	LPV	5.8	-0.40
	I47V	0.05	1.1×10^{-5}	LPV	1.8	-0.29
	I50L	0.75	9.0×10^{-7}	ATV	1.2	-0.34
	I50V	0.93	1.1×10^{-5}	DRV	0.68	-0.07
	I54L	0.05	9.0×10^{-7}	DRV	0.98	-0.01
	I84V	0.82	1.1×10^{-5}	ATV	0.60	-0.34
				DRV	0.94	-0.01
				IDV	0.73	-0.39
				TPV	0.26	-0.39
	L33F	0.49	6.3×10^{-6}	TPV	1.4	0.02
	L90M	0.30	3.2×10^{-6}	NFV	1.5	0.01
				SQV	1.1	-0.28
	M46I	0.05	5.6×10^{-5}	IDV	1.0	-0.29
	M46L	0.05	1.3×10^{-6}	IDV	0.76	-0.24
	N88S	0.55	1.1×10^{-5}	ATV	3.1	-0.31
	V32I	0.09	4.1×10^{-5}	LPV	0.53	-0.16
	V82A	0.59	1.1×10^{-5}	LPV	1.03	-0.33
	V82F	0.79	3.4×10^{-7}	IDV	0.89	-0.58
				LPV	1.45	-0.44
	V82T	0.22	*	IDV	0.98	-0.34
				LPV	0.87	-0.17
				TPV	0.68	-0.20
II	G140S	0.71	2.2×10^{-5}	RAL	2.1	0.03
	N155H	0.55	5.3×10^{-7}	EVG	20	0.00
				RAL	27	0.02
	Q148H	0.73	2.0×10^{-6}	EVG	6.8	-0.04
				RAL	86	0.06
	Q148K	0.76	6.5×10^{-6}	EVG	19	0.03
				RAL	128	-0.06
	Q148R	0.61	1.1×10^{-5}	EVG	68	0.06
				RAL	90	0.04
	Y143C	0.74	1.1×10^{-5}	RAL	3.6	0.06
	Y143H	0.55	1.1×10^{-5}	RAL	2.7	-0.04
	Y143R	0.32	*	RAL	75	-0.01
FI	G36D	0.12	2.2×10^{-5}	ENF	1.7	-0.45
	N42T	0.54	5.3×10^{-7}	ENF	2.9	-0.13
	N43D	0.88	1.1×10^{-5}	ENF	13	-0.06
	Q40H	0.26	1.5×10^{-6}	ENF	12	-0.31
	V38A	0.17	1.1×10^{-5}	ENF	11	-0.32

* Indicates mutation that requires two nucleotide changes; mutation rate depends on prevalence of intermediate states.

Table 7.1.4: Parameters for pre-existing frequency of mutations and exit rate from the latent reservoir. See **Methods** and **Supplementary Methods** for explanations.

class	mutation	equilibrium frequency	reservoir exit (days)
NRTI	K65R	2.7×10^{-5}	12
	M184V	2.4×10^{-5}	14
	M41L	7.8×10^{-6}	43
	T215Y	*	*
NNRTI	G190S	2.8×10^{-5}	12
	K101P	*	*
	K103N	4.9×10^{-6}	68
	Y181C	4.3×10^{-5}	8
	Y181I	*	*
PI	D30N	2.0×10^{-4}	2
	G48V	1.9×10^{-6}	177
	I47A	*	*
	I47V	2.2×10^{-4}	2
	I50L	1.2×10^{-6}	279
	I50V	1.2×10^{-5}	28
	I54L	1.8×10^{-5}	19
	I84V	1.4×10^{-5}	25
	L33F	1.3×10^{-5}	26
	L90M	1.1×10^{-5}	31
	M46I	1.1×10^{-3}	< 0.5
	M46L	2.6×10^{-5}	13
	N88S	2.0×10^{-5}	17
	V32I	4.6×10^{-4}	1
	V82A	1.9×10^{-5}	18
	V82F	4.3×10^{-7}	769
	V82T	*	*
II	G140S	3.1×10^{-5}	11
	N155H	9.6×10^{-7}	349
	Q148H	2.0×10^{-6}	166
	Q148K	8.5×10^{-6}	39
	Q148R	1.8×10^{-5}	18
	Y143C	1.5×10^{-5}	22
	Y143H	2.0×10^{-5}	17
	Y143R	*	*
FI	G36D	1.8×10^{-4}	2
	N42T	9.8×10^{-7}	342
	N43D	1.3×10^{-5}	27
	Q40H	5.6×10^{-6}	59
	V38A	6.5×10^{-5}	5

* Indicates mutation that requires two nucleotide changes;
equilibrium frequency depends on prevalence of intermediate states.

Table 7.1.5: Parameters for pre-existing frequency and exit rate from the latent reservoir for best "synthetic" mutation for each drug. See Methods and Supplementary Methods for explanations.

class	drug	equilibrium frequency	reservoir exit (days)
NRTI	AZT	2.4×10^{-5}	14
	d4T	7.8×10^{-6}	43
	3TC	2.7×10^{-5}	12
	FTC	2.7×10^{-5}	12
	ABC	2.7×10^{-5}	12
	ddI	2.7×10^{-5}	12
	TDF	2.7×10^{-5}	12
NNRTI	EFV	4.3×10^{-5}	8
	NVP	4.3×10^{-5}	8
	ETV	4.3×10^{-5}	8
PI	DRV	1.8×10^{-5}	19
	NFV	2.0×10^{-4}	2
	SQV	1.1×10^{-5}	31
	LPV	4.6×10^{-4}	1
	ATV	2.0×10^{-5}	17
	IDV	1.1×10^{-3}	< 0.5
	TPV	1.4×10^{-5}	25
II	RAL	3.1×10^{-5}	11
	EVG	1.8×10^{-5}	18
FI	ENF	1.8×10^{-4}	2

Table 7.1.6: Viral dynamics parameters in the absence of drug therapy

	Parameter	Value	Units	Reference
R_{oo}	Baseline basic reproduction ratio	10	(unitless)	See text
d_y	Death rate of actively infected cells	1	d^{-1}	[188]
v_o	Residual plasma viral load maintained by activation from latent reservoir, absent viral replication	2	RNA copies per ml plasma	[74]
A	Latent reservoir exit rate	3000	cells.d^{-1}	Based on v_o, see text

Table 7.1.7: Nucleotide substitution rate parameters for HIV. Each entry gives the per-site transition probability from row base to column base in one round of viral replication. For derivation and source see Section 2.5.2. The extraordinary skew of this matrix (the largest entry, G-to-A mutation, is more than 300 times the smallest, C-to-G mutation) reflects the base composition of the genome, particularly the bias towards A. Values less than 10^{-6} are particularly uncertain, as they were computed from fewer than 5 substitution observations each.

	U	C	A	G
U		1.1×10^{-5}	1.3×10^{-5}	3.6×10^{-6}
C	2.4×10^{-5}		6.5×10^{-6}	1.7×10^{-7}
A	7.9×10^{-7}	5.3×10^{-7}		1.1×10^{-5}
G	8.5×10^{-7}	8.5×10^{-7}	5.5×10^{-5}	

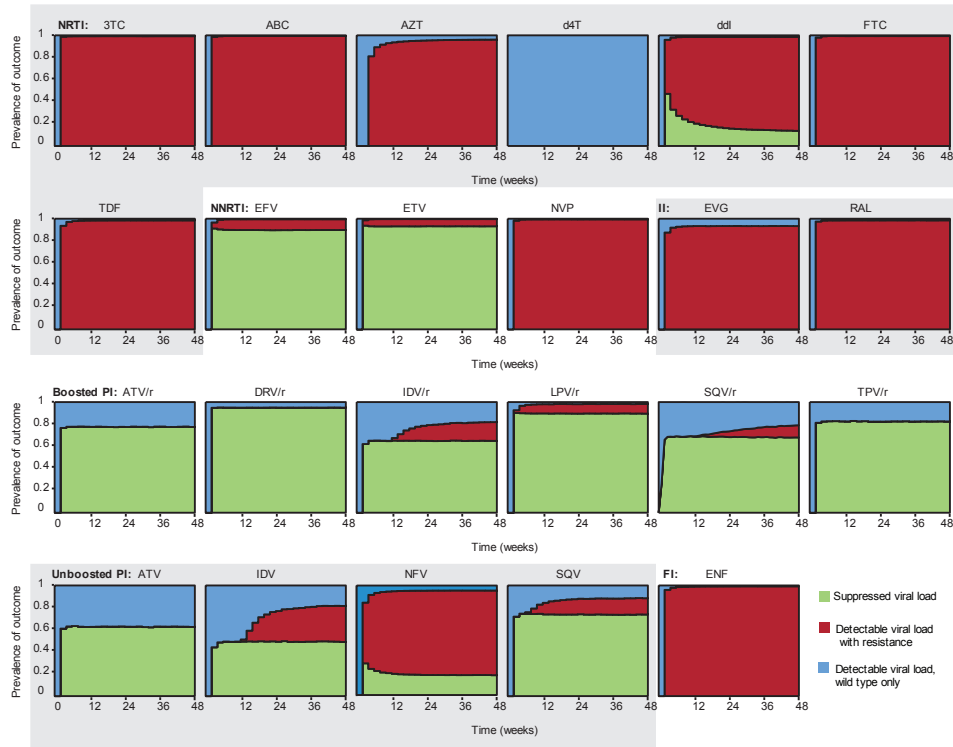## 7.2 Supplementary figures referenced in Chapter 2

Figure 7.2.1: Simulated clinical outcomes versus adherence for all drugs. In "Suppression" trials, patients begin with a realistic distribution of treatment-naive viral loads (between 3000 and 10^6 c.ml^{-1}) and undergo monotherapy for a full 48 weeks. Virologic failure (VF) is defined as a viral load above 50 c.ml^{-1}at Week 48. VF is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. Adherence (x-axis) is measured as the fraction of scheduled doses taken. The height of the area shaded indicates probability of the corresponding outcome at that adherence level. 3TC, lamivudine; ABC, abacavir; AZT, zidovudine; d4T, stavudine; ddI, didanosine, FTC, emtricitabine; TDF, tenofovir disoproxil fumarate; EFV, efavirenz; ETV, etravirine; NVP, nevirapine; ATV, atazanavir; DRV, darunavir; IDV, indinavir; LPV, lopinavir; NFV, nelfinavir; SQV, saquinavir; TPV, tipranavir; EVG, elvitegravir; RAL, raltegravir; ENF, enfuvirtide.
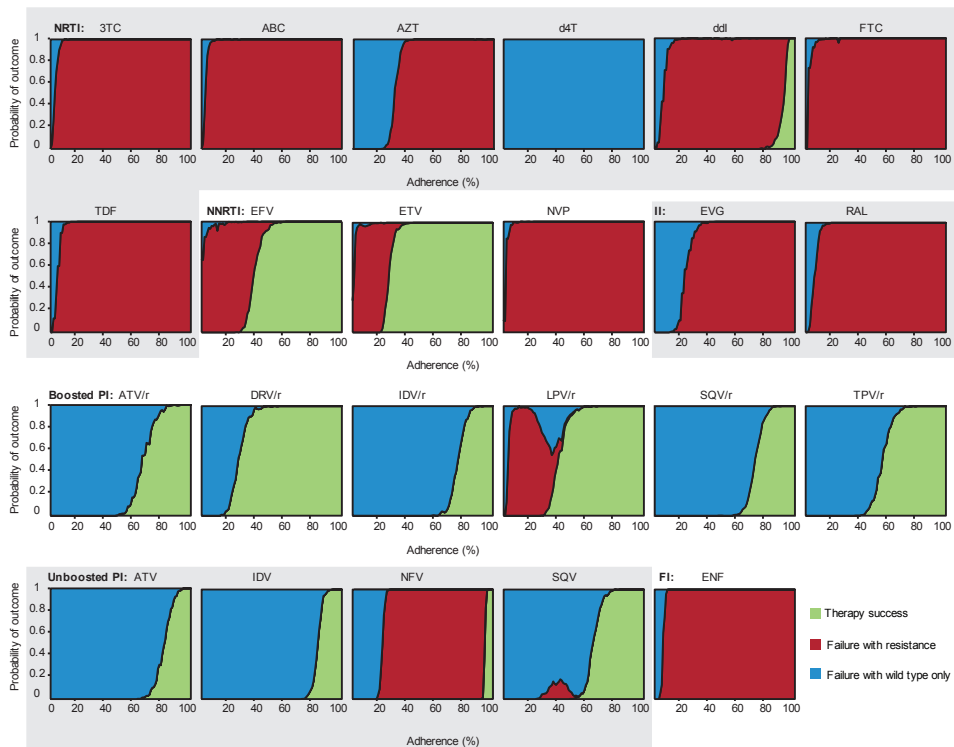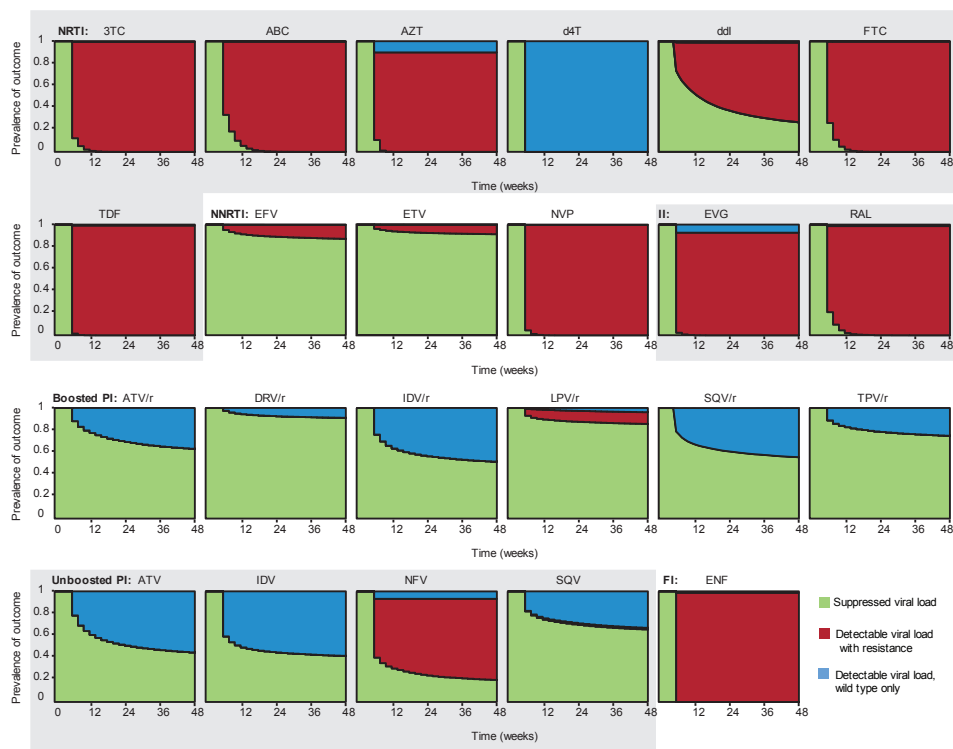
Figure 7.2.2: Simulated clinical outcomes versus adherence for all drugs. In "Maintenance" trials, patients begin with full viral suppression and undergo monotherapy for 48 weeks or until virologic failure (VF), whichever occurs first. VF is defined as "confirmed rebound": two consecutive weekly measurements (starting at week 5) with viral load above 200 c.ml^{-1}. VF is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. Adherence (x-axis) is measured as the fraction of scheduled doses taken. The height of the area shaded indicates probability of the corresponding outcome at that adherence level.

247

Figure 7.2.3: Simulated clinical outcomes versus time for all drugs. In "Suppression" trials, patients begin with a realistic distribution of treatment-naive viral loads (between 3000 and 10^6 c.ml^{-1}) and undergo monotherapy for a variable time (x-axis). "Detectable viral load" is defined as above 50 c.ml^{-1}and is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. The height of the area shaded indicates prevalence of the corresponding outcome at that time. Patients have a realistic distribution of adherence levels with an average of 70%.

**Figure 7.2.4:** Simulated clinical outcomes versus time for all drugs. In "Maintenance" trials, patients begin the trial with full viral suppression and undergo monotherapy for a variable amount of time (x-axis) or until "detectable viral load" is observed, whichever occurs first. "Detectable viral load" is defined as "confirmed rebound": two consecutive weekly measurements (starting at Week 5) above 200 c.ml^{-1}. It is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. The height of the area shaded indicates prevalence of the corresponding outcome at that time. Patients have a realistic distribution of adherence levels with an average of 70%.

Figure 7.2.5: Simulated clinical outcomes versus time for all drugs. In "Maintenance with recovery" trials, patients begin the trial with full viral suppression and undergo monotherapy for a variable amount of time (x-axis). "Detectable viral load" is defined as "confirmed rebound": two consecutive weekly measurements (starting at Week 5) with viral load above 200 c.ml^{-1}. It is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. We allow recovery, meaning that patients stay in the trial to see if they will re-suppress, instead of being removed immediately like in regular "Maintenance" trials. The height of the area shaded indicates prevalence of the corresponding outcome at that time-point. Patients have a realistic distribution of adherence levels with an average of 70%.
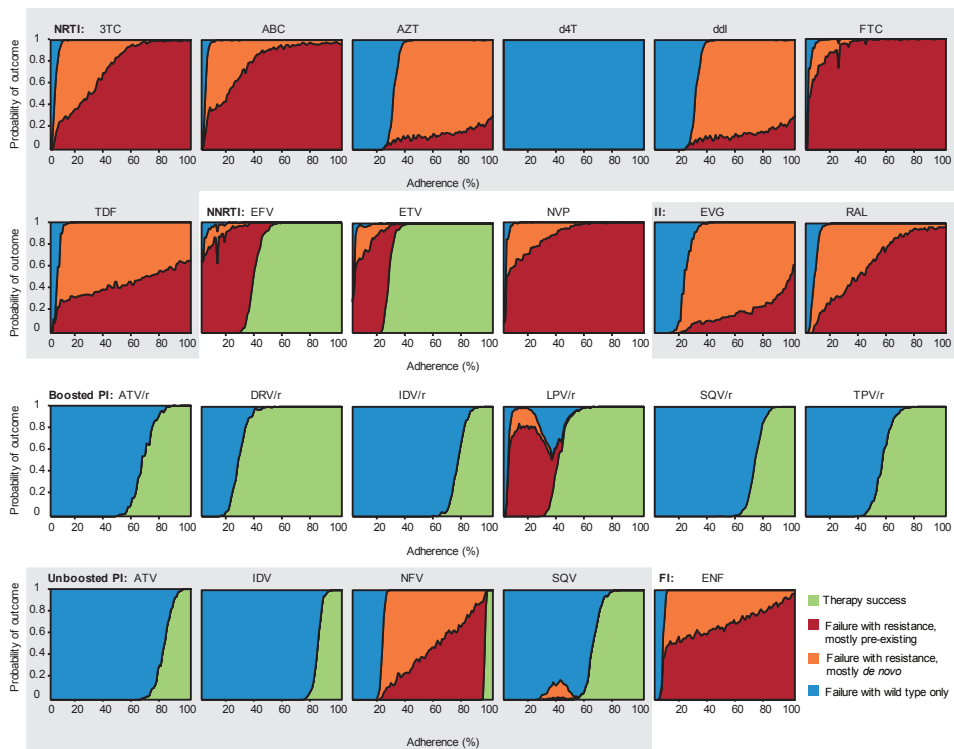
Figure 7.2.6: Simulated clinical outcomes versus adherence for all drugs, distinguishing pre-existing from *de novo* mutations. In the "Suppression" trials shown, patients begin with a realistic distribution of treatment-naive viral loads (between 3000 and 10^6 c.ml^{-1}) and undergo monotherapy for a full 48 weeks. Virologic failure (VF) is defined as a viral load above 50 c.ml^{-1} at Week 48. VF is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. Resistance is classified as *de novo* if the majority of mutants at the time of failure descended from a mutation event that occurred during replication since the start of the trial. Otherwise, resistance is classified as "pre-existing," which includes mutants arising from both the pre-treatment plasma population and the latent reservoir. Adherence (x-axis) is measured as the fraction of scheduled doses taken. The height of the area shaded indicates probability of the corresponding outcome at that adherence level.
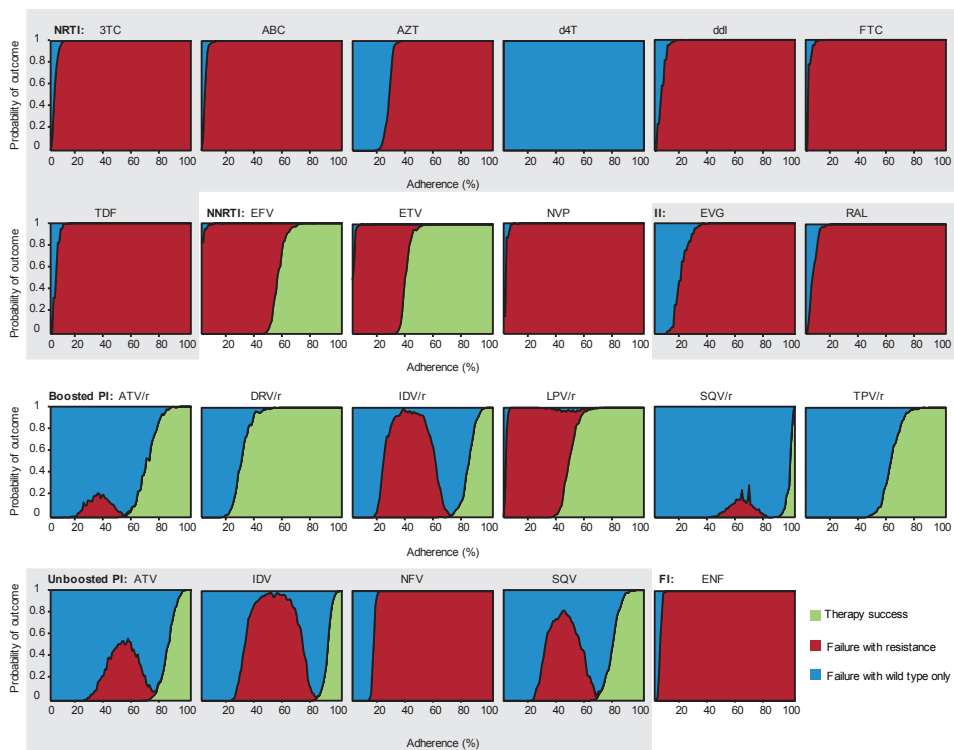
Figure 7.2.7: Simulated clinical outcomes versus adherence for all drugs, distinguishing pre-existing from *de novo* mutations. In the "Maintenance" trials shown, patients begin with full viral suppression and undergo monotherapy for 48 weeks or until virologic failure (VF), whichever occurs first. VF is defined as "confirmed rebound": two consecutive weekly measurements (starting at Week 5) with viral load above 200 c.ml^{-1}. VF is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. Resistance is classified as *de novo* if the majority of mutants at the time of failure descended from a mutation event that occurred during replication since the start of the trial. Otherwise, resistance is classified as "pre-existing," which includes mutants arising from both the pre-treatment plasma population and the latent reservoir. Adherence (x-axis) is measured as the fraction of scheduled doses taken. The height of the area shaded indicates probability of the corresponding outcome at that adherence level.
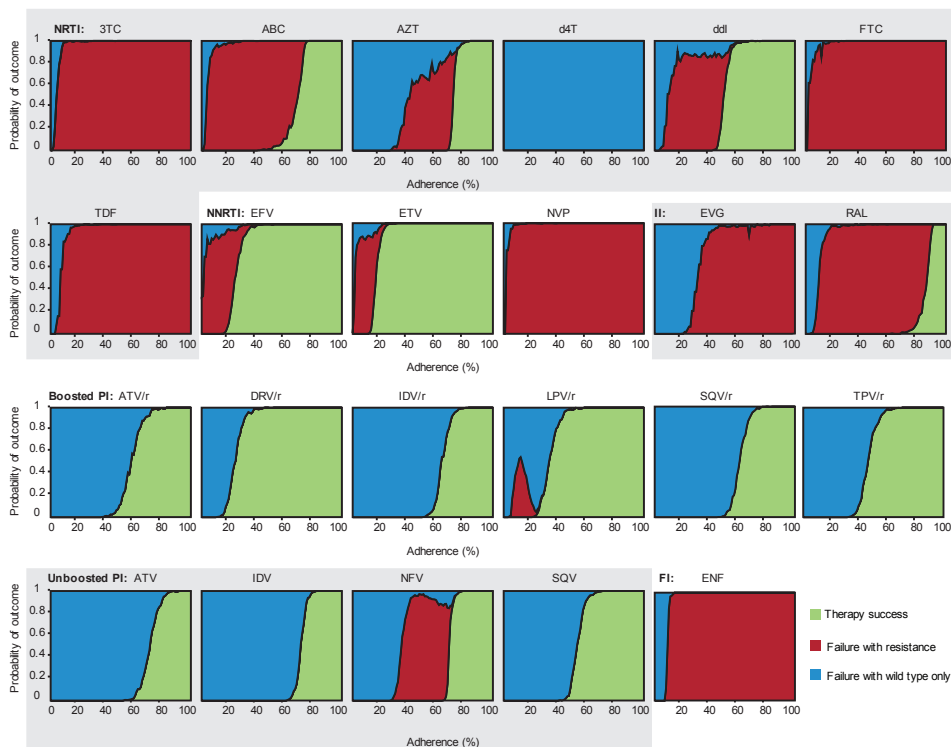
Figure 7.2.8: Simulated clinical outcomes versus adherence for all drugs, $R_{oo}=20$. Results are shown for "Maintenance" trials only. In the "Maintenance" trials shown, patients begin with full viral suppression and undergo monotherapy for 48 weeks or until virologic failure (VF), whichever occurs first. VF is defined as "confirmed rebound": two consecutive weekly measurements (starting at Week 5) with viral load above 200 c.ml^{-1}. VF is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. Adherence (x-axis) is measured as the fraction of scheduled doses taken. The height of the area shaded indicates probability of the corresponding outcome at that adherence level. As compared to $R_{oo}=10$, increasing R_{oo} to 20 leads to higher adherence levels being required for treatment success, and it extends the range of adherence levels (in both directions) for which resistant strains can cause failure. Mutant VF becomes a possible outcome for the PIs ATV, ATV/r, IDV, IDV/r, and SQV/r, and treatment success cannot occur at any adherence level for ddI and NFV.

Figure 7.2.9: Simulated clinical outcomes versus adherence for all drugs, $R_{oo}{=}5$. Results are shown for "Maintenance" trials only. In the "Mainte-nance" trials shown, patients begin with full viral suppression and undergo monotherapy for 48 weeks or until virologic failure (VF), whichever occurs first. VF is defined as "confirmed rebound": two consecutive weekly measure-ments (starting at Week 5) with viral load above 200 c.ml^{-1}. VF is classified as "via resistance" if at least 20% of the viral population at the time of de-tection is mutant. Adherence (x-axis) is measured as the fraction of sched-uled doses taken. The height of the area shaded indicates probability of the corresponding outcome at that adherence level. As compared to $R_{oo}{=}10$, de-creasing R_{oo} to 5 leads to lower adherence levels being required for treatment success, and it reduces the range of adherence levels for which resistant strains can cause failure. A range of high adherence levels appears where there is treatment success for ABC and AZT, and near-perfect adherence is no longer required for ddI and NFV success. Mutant VF no longer occurs for SQV, and for AZT and ddI, wild-type failure may be the first outcome to occur as ad-herence levels decrease from the successful range.

Figure 7.2.10: Simulated clinical outcomes versus adherence for NRTIs with large inter-experimental variation in half-life. The ranges included were {10, 16, 22} for 3TC, {4, 8.5, 11} for AZT, {3.5, 7, 10} for d4T, {15, 18, 20} for ddI, and {60, 120, 180} for TDF. Results are shown for "Maintenance" trials only. In the "Maintenance" trials shown, patients begin with full viral suppression and undergo monotherapy for 48 weeks or until virologic failure (VF), whichever occurs first. VF is defined as "confirmed rebound": two consecutive weekly measurements (starting at Week 5) with viral load above 200 c.ml^{-1}. VF is classified as "via resistance" if at least 20% of the viral population at the time of detection is mutant. Adherence (x-axis) is measured as the fraction of scheduled doses taken. The height of the area shaded indicates probability of the corresponding outcome at that adherence level. Compared to the half-lives used throughout the rest of the paper (see Supplementary Table 7.1.1), the results barely change for 3TC or d4T. For AZT, varying the half-life changes the adherence level where wild-type failure becomes more likely than mutant failure. For ddI, the adherence level where treatment success occurs shifts. For higher TDF half-lives, mutant VF becomes the only outcome, with the exception of rare ($<$ 3%) wild-type failure at the lowest adherence levels for $t_{1/2} = 120$ hours.

Figure 7.2.11: Selection regimes for DRV/r-RAL two-drug therapy. Depending on the length of a treatment interruption to one or both drugs, treatment may be fully suppressive or select for the wild-type strain, a mutant resistant to DRV, a mutant resistant to RAL, or combinations of these strains. The yellow region, where the MSW for both drugs overlap, is barely visible, and it is located where the other MSW regions meet, near the center of the graph.

256

Figure 7.2.12: Relative risk of wild type- vs. mutant-caused virologic failure for anti-HIV drugs, considering the best "synthetic" mutation defined in Section 2.5.3. Two metrics can be used to compare the risk of resistance to the likelihood of wild-type growth, shown on both axes. The x-axis measures the time until a patient interrupting treatment reaches the MSW, divided by the time until that patient reaches the WGW. The y-axis measures the number of days that a patient spends in the MSW during a treatment interruption. Drugs tend to cluster near the endpoints of the x-axis: most NRTIs, the IIs, and the FI are on the left, meaning that the patient enters the MSW immediately or soon after interruption, and most PIs are on the right, meaning that the patient waits relatively long to enter the MSW. Section 2.5.7 further describes the two metrics and explains how they were used in Fig 2.2.2f in the main text to rank the drugs by relative risk of mutant-based versus wild type-based VF. Note that the symbol for DRV/r is obscured behind the symbol for d4T at $(1, 0)$.
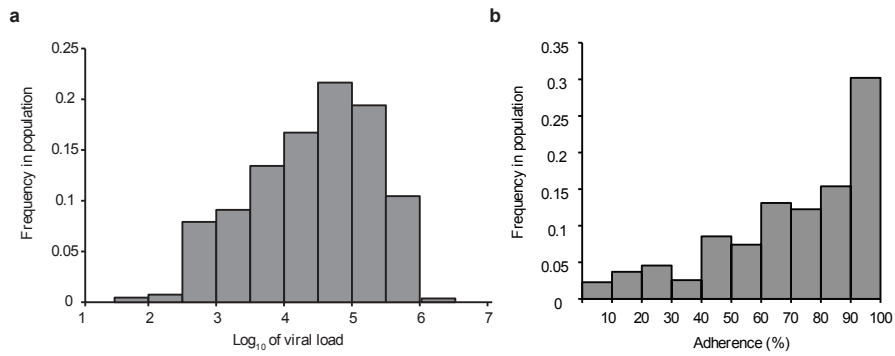
Figure 7.2.13: Distribution of a) viral load setpoints [160] (data available at www.hiv.lanl.gov/content/immunology) and b) adherence levels [16] used in simulations.

## 7.3   Supplementary figures referenced in Chapter 3

Figure 7.3.1: Clearance probabilities and rebound times following LRA therapy predicted from model, versus infectivity variance λ. A) Three cases for the population-level distribution of LR size (see Section 3.4.1). Case i) All patients have the same latent reservoir size, $M_{LR} = 10^6$, estimated from the geometric mean number of cells that are capable of producing infection in laboratory co-culture assays. Case ii) Latent reservoir size is distributed according to variation observed in co-culture assays, with geometric mean 10^6. Case iii) The latent reservoir includes many cells that fail to be detected in co-culture but have intact viral genomes. B) Probability that the reservoir is cleared by LRA. Clearance occurs if all cells in the reservoir die before a reactivating lineage leads to viral rebound. C) Median viral rebound times, among patients who do not have clearance. Each point shows the average of $10^4 - 10^5$ simulated patients. D) Survival curves for patients following LRA therapy. The percentage of patients who have not yet experienced viral rebound is plotted as a function of the time after LRA therapy and interruption of HAART. Curves are colored based on the efficacy of LRA in reducing the size of the LR ($Q = 0$ to 6, see legend). Results are shown for $10^4 - 10^5$ patients, a half-life of 44 months, $R_0 = 2$, and $a = 7 \times 10^{-5}$ day^{-1}.

Figure 7.3.2: LRA therapy efficacies required for different goals, plotted versus reservoir size M_{LR} for different infectivity variances λ. A) The threshold Q that takes the patient into the activation-limited regime, where stochastic waiting time contributes substantially to rebound time (defined in Sec. 3.4.5). B) The target efficacy Q at which at least 50% of patients still have suppressed viral load one month after treatment interruption. C) The target efficacy Q at which at least 50% of patients still have suppressed viral load one year after treatment interruption. D) The target efficacy Q at which at least 50% of patients have eradicated the reservoir without experiencing viral rebound. Because some patients may go for a year without rebound but then rebound later, the target Q for one year off therapy is always less than that for a lifetime off therapy. Results are shown for $10^4 - 10^5$ patients, a half-life of 44 months, $R_0 = 2$, and $a = 7 \times 10^{-5}$ day^{-1}.

Figure 7.3.3: Survival curves for patients following LRA therapy depend on the size of the latent reservoir and the reservoir half-life (varying d_z). The percentage of patients who have not yet experienced viral rebound (viral load $>$ 200 c ml^{-1}) is plotted as a function of the time after LRA therapy and interruption of HAART. Curves are colored based on the efficacy of LRA in reducing the size of the LR ($Q = 0$ to 6, see legend). A) $d_z = 0$ day^{-1}, half-life is 330 months (27 years). B) $d_z = 4.6 \times 10^{-4}$ day^{-1}, half-life is 44 months. C) $d_z = 4 \times 10^{-3}$ day^{-1}, half-life is 6 months. Decreasing the LR half-life (increasing d_z) makes survival times longer and clearance more likely. Including interpatient variation (ii) makes the survival curves fall off more gradually, while allowing for higher reservoir sizes (iii) increases the required drug efficacy. Solid lines represent simulations, and open circles represent approximations from a branching process calculation (Section 3.4.3). Results are shown for $10^4 - 10^5$ patients, $\lambda = 20$, $R_0 = 2$, and $a = 7 \times 10^{-5}$ day^{-1}

Figure 7.3.4: Survival curves for patients following LRA therapy depend on the size of the latent reservoir and the infectivity variance λ. The percentage of patients who have not yet experienced viral rebound (viral load $>$ 200 c ml^{-1}) is plotted as a function of the time after LRA therapy and interruption of HAART. Curves are colored based on the efficacy of LRA in reducing the size of the LR ($Q = 0$ to 6, see legend). Increasing the infectivity variance λ makes survival times longer and clearance more likely. Including interpatient variation (ii) makes the survival curves fall off more gradually, while allowing for higher reservoir sizes (iii) increases the required drug efficacy. Solid lines represent simulations, and open circles represent approximations from a branching process calculation (Section 3.4.3). Results are shown for a half-life of 44 months, $R_0 = 2$, and $a = 7 \times 10^{-5}$ day^{-1}.

Figure 7.3.5 *(following page)*: Clearance probabilities and rebound times after LRA predicted from model for alternate parameter choices. i) Best-estimate parameter values shown in the main text: reservoir size is predicted by the distribution of co-culture results (case ii), $a = 7 \times 10^{-5}$ day^{-1}, $d_z = 4.3 \times 10^{-4}$ day^{-1} (half-life of 44 months), $R_o = 2$. ii) Best-case-scenario, where reservoir size is predicted by mean co-culture results (case i), half-life is short and strong stochastic effects decrease clearance probability. Parameter values are $a = 10^{-5}$ day^{-1}, $d_z = 4 \times 10^{-3}$ day^{-1} (half-life 6 months), $R_o = 2.2$, $\lambda = 50$, and $M_{LR} = 10^6$. This low estimate for a follows from the assumption that all cells with HIV-DNA contribue to the residual viral load observed during HAART. iii) Worst-case-scenario, when LR is large (case iii), cells reactivate frequently, are extremely long lived, and smaller stochastic effects mean most reactivating cells lead to rebound. Parameter values are $a = 7 \times 10^{-5}$ day^{-1}, $d_z = 0$ day^{-1}, $R_o = 2$, $\lambda = 2$, median $M_{LR} \approx 7 \times 10^6$. This high estimate for a follows from the assumption that all virions in the residual viral load come from cells with virus with intact provirus. The lower limit on d_z is realized if homeostatic proliferation or other mechanisms balance reservoir decay caused by mechanisms other than reactivation. R_o is always adjusted to ensure that the baseline rebound time following HAART interruption was constant. All results are for $10^4 - 10^6$ simulated patients.
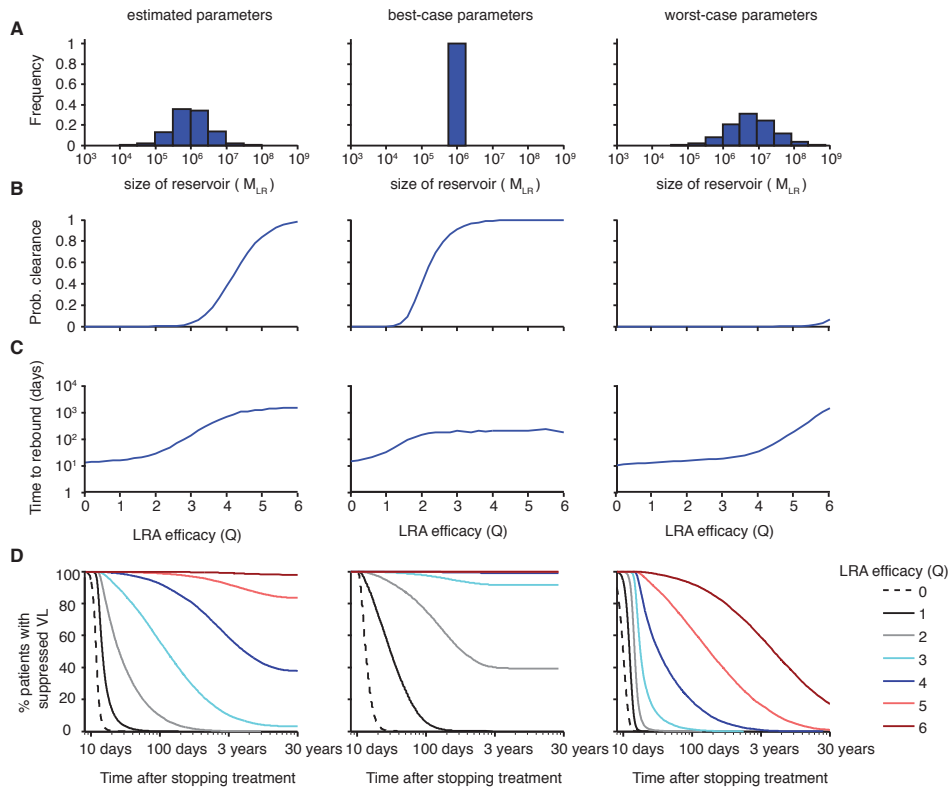
Figure 7.3.5 *(continued)*

Figure 7.4.1: Invasion fitness of a mutating invader against a nonmutating resident, for varying levels of recombination c, computed as the largest eigenvalue of the matrix $CQ_{u'}F$ minus 1, from Eq. (5.13). There is a discontinuity at $u' = 0$ not shown in the graph: the invasion fitness $s_0(0)$ is zero for all parameter values. Parameters are as in Fig. 5.3.4: $n = 4$, $a = 0.4$.

Figure 7.4.2: ESMR (plotted on a log scale) versus fitness benefit a and number of traits n, comparing the estimate presented in Fig. 5.3.5 (thick solid lines) to an alternate estimate (thick lines with white stripe in middle). Points show simulated ESMR values (see Methods) and dashed lines plot u_Ω. The alternate estimate uses Eq. (5.18) to approximate fitness of the optimal trait within the resident strain; it only requires simulation to measure the resident strain's period. The fitness approximation is accurate if $n \geq 6$ and $a \ll 1$. No recombination is shown in this figure.

References

[1] J. Abate, G. L. Choudhury, and W. Whitt. An introduction to numerical transform inversion and its application to probability models. In W. K. Grassmann, editor, *Computational Probability*, volume 24 of *International Series in Operations Research & Management Science*, pages 257–323. Springer US, 2000. doi: 10.1007/978-1-4757-4828-4_8.

[2] M. E. Abram, A. L. Ferris, W. Shao, W. G. Alvord, and S. H. Hughes. Nature, position, and frequency of mutations made in a single cycle of HIV-1 replication. *J. Virol.*, 84:9864–9878, 2010. doi: 10.1128/JVI.00915-10.

[3] G. Achaz, S. Palmer, M. Kearney, F. Maldarelli, J. W. Mellors, J. M. Coffin, and J. Wakeley. A robust measure of HIV-1 population turnover within chronically infected individuals. *Mol. Biol. Evol.*, 21(10):1902–1912, 2004.

[4] A. Aharoni, L. Gaidukov, O. Khersonsky, S. M. Q. Gould, C. Roodveldt, and D. S. Tawfik. The 'evolvability' of promiscuous protein functions. *Nat. Genet.*, 37(1):73–76, 2005.

[5] E. D. Akhunov, A. W. Goodyear, S. Geng, L. L. Qi, B. Echalier, B. S. Gill, J. P. Gustafson Miftahudin, G. Lazo, and et al. The organization and rate of evolution of wheat genomes are correlated with recombination rates along chromosome arms. *Genome Res.*, 13:753–763, 2003.

[6] B. Allen and D. I. S. Rosenbloom. Mutation rate evolution in replicator dynamics. *Bull. Math. Biol.*, 74(11):2650–2675, 2012.

[7] J. B. Andre and B. Godelle. The Evolution of Mutation Rate in Finite Asexual Populations. *Genetics*, 172(1):611–626, 2005.

[8] N. M. Archin, A. Espeseth, D. Parker, M. Cheema, D. Hazuda, and D. M. Margolis. Expression of latent HIV induced by the potent HDAC

inhibitor suberoylanilide hydroxamic acid. *AIDS Res. Hum. Retrovir.*, 25 (2):207–212, 2009. doi: 10.1089/aid.2008.0191.

[9] N. M. Archin, M. Cheema, D. Parker, A. Wiegand, R. J. Bosch, J. M. Coffin, J. Eron, M. Cohen, and D. M. Margolis. Antiretroviral intensification and valproic acid lack sustained effect on residual HIV-1 viremia or resting CD4+ cell infection. *PLoS ONE*, 5(2):e9390, 2010. doi: 10.1371/journal.pone.0009390.

[10] N. M. Archin, A. L. Liberty, A. D. Kashuba, S. K. Choudhary, J. D. Kuruc, A. M. Crooks, D. C. Parker, E. M. Anderson, M. F. Kearney, M. C. Strain, D. D. Richman, M. G. Hudgens, R. J. Bosch, J. M. Coffin, J. J. Eron, D. J. Hazuda, and D. M. Margolis. Administration of vorinostat disrupts HIV-1 latency in patients on antiretroviral therapy. *Nature*, 487(7408):482–485, 2012. doi: 10.1038/nature11286.

[11] N. M. Archin, N. K. Vaidya, J. D. Kuruc, A. L. Liberty, A. Wiegand, M. F. Kearney, M. S. Cohen, J. M. Coffin, R. J. Bosch, C. L. Gay, J. J. Eron, D. M. Margolis, and A. S. Perelson. Immediate antiviral therapy appears to restrict resting CD4+ cell HIV-1 infection without accelerating the decay of latent infection. *Proc. Natl. Acad. Sci. USA*, 109(24):9523–9528, 2012. doi: 10.1073/pnas.1120248109.

[12] J. R. Arribas, A. Horban, J. Gerstoft, G. Fätkenheuer, M. Nelson, N. Clumeck, F. Pulido, A. Hill, Y. van Delft, T. Stark, and C. Moecklinghoff. The MONET trial: darunavir/ritonavir with or without nucleoside analogues, for patients with HIV RNA below 50 copies/ml. *AIDS*, 24(2):223–230, 2010. doi: 10.1097/QAD.0b013e3283348944.

[13] R. Axelrod. *The Evolution Of Cooperation*. Basic Books, New York, 1984.

[14] C. F. Baer, M. M. Miyamoto, and D. R. Denver. Mutation rate variation in multicellular eukaryotes: causes and consequences. *Nat. Rev. Genet.*, 8: 619–631, 2007.

[15] D. R. Bangsberg, A. R. Moss, and S. G. Deeks. Paradoxes of adherence and drug resistance to HIV antiretroviral therapy. *J. Antimicrob. Chemother.*, 53(5):696 –699, 2004. doi: 10.1093/jac/dkh162.

[16] D. R. Bangsberg, E. P. Acosta, R. Gupta, D. Guzman, E. D. Riley, P. R. Harrigan, N. Parkin, and S. G. Deeks. Adherence-resistance relationships for protease and non-nucleoside reverse transcriptase inhibitors explained by virological fitness. *AIDS*, 20(2):223–231, 2006. doi: 10.1097/01.aids.0000199825.34241.49.

[17] D. R. Bangsberg, D. L. Kroetz, and S. G. Deeks. Adherence-resistance relationships to combination HIV antiretroviral therapy. *Current HIV/AIDS Reports*, 4(2):65–72, 2007.

[18] A. V. Bannerjee. A simple model of herd behavior. *Q. J. Econ*, 100: 992–1026, 1992.

[19] S. Bansal, B. Pourbohloul, and L. A. Meyers. A comparative analysis of influenza vaccination programs. *PLoS Med.*, 3(10):e387, 2006.

[20] A. L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[21] A. M. Barton. Application of cascade theory to online systems: a study of email and Google cascades. *Minn. J.L. Sci. & Tech.*, 10:473–502, 2009.

[22] S. Basu, G. B. Chapman, and A. P. Galvani. Integrating epidemiology, psychology, and economics to achieve HPV vaccination targets. *Proc. Natl. Acad. Sci. USA*, 105:19018–19023, 2008.

[23] C. T. Bauch. Imitation dynamics predict vaccinating behaviour. *Proc. R. Soc. Lond. Ser. B*, 272:1669–1675, 2005.

[24] C. T. Bauch and D. J. D. Earn. Vaccination and the theory of games. *Proc. Natl. Acad. Sci. USA*, 101:13391–13394, 2004.

[25] C. T. Bauch, A. P. Galvani, and D. J. D. Earn. Group interest versus self interest in smallpox vaccination policy. *Proc. Natl. Acad. Sci. USA*, 100: 10564–10567, 2003.

[26] B.M. Best, P.P. Koopmans, S.L. Letendre, E.V. Capparelli, S.S. Rossi, D.B. Clifford, A.C. Collier, B.B. Gelman, G. Mbeo, J.A. McCutchan, et al. Efavirenz concentrations in CSF exceed IC$_{50}$ for wild-type HIV. *J. Antimicrob. Chemother.*, 66(2):354, 2011.

[27] A. Biancotto, J.-C. Grivel, F. Gondois-Rey, L. Bettendroffer, R. Vigne, S. Brown, L. B. Margolis, and I. Hirsch. Dual role of prostratin in inhibition of infection and reactivation of human immunodeficiency virus from latency in primary blood lymphocytes and lymphoid tissue. *J. Virol.*, 78(19):10507–10515, 2004. doi: 10.1128/JVI.78.19.10507-10515.2004.

[28] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Pol. Econ.*, 107: 797–817, 1992.

[29] J. N. Blankson, J. D. Siliciano, and R. F Siliciano. The Effect of early treatment on the latent reservoir of HIV-1. *J. Inf. Dis.*, 191(9):1394–1396, 2005.

[30] C. I. Bliss. The toxicity of poisons applied jointly. *Ann. Appl. Biol.*, 26(3): 585–615, 1939. doi: 10.1111/j.1744-7348.1939.tb06990.x.

[31] L. E. Blume. The statistical mechanics of strategic interaction. *Games Econ. Behav.*, 5:387–424, 1993.

[32] V. F. Boltz, Y. Zheng, S. Lockman, F. Hong, E. K. Halvas, J. McIntyre, J. S. Currier, M. C. Chibowa, C. Kanyama, A. Nair, W. Owino-Ong'or, M. Hughes, J. M. Coffin, and J. W. Mellors. Role of low-frequency HIV-1 variants in failure of nevirapine-containing antiviral therapy in women previously exposed to single-dose nevirapine. *Proc. Natl. Acad. Sci. USA*, 108(2):9202–9207, 2011.

[33] S. Bonhoeffer and M. A. Nowak. Pre-existence and emergence of drug resistance in HIV-1 infection. *Proc. R. Soc. B*, 264(1382):631–637, 1997.

[34] N. Bonneuil. Attractors and confiners in demography. *Ann. Oper. Res.*, 37: 17–32, 1992. doi: 10.1007/BF02071046.

[35] T. Boulikas. Evolutionary consequences of nonrandom damage and repair of chromatin domains. *J. Mol. Evol.*, 35:156–180, 1992.

[36] M. van Boven, D. Klinkenberg, I. Pen, F. J. Weissing, and H. Heesterbeek. Self-interest versus group-interest in antiviral control. *PLoS ONE*, 3: e1558, 2008.

[37] M. Brandstrom and H. Ellegren. The genomic landscape of short insertion and deletion polymorphisms in the chicken (*Gallus gallus*) genome: a high frequency of deletions in tandem duplicates. *Genetics,* 176:1691–1701, 2007.

[38] R. Breban, R. Vardavas, and S. Blower. Mean-field analysis of an inductive reasoning game: Application to influenza vaccination. *Phys. Rev. E,* 76: 031127, 2007.

[39] A. Buckling, F. Harrison, M. Vos, M. A. Brockhurst, A. Gardner, S. A. West, and A. Griffin. Siderophore-mediated cooperation and virulence in Pseudomonas aeruginosa. *FEMS Microbiol. Ecol.,* 62(2):135–141, 2007.

[40] L. W. Buss and J. B. C. Jackson. Competitive networks: nontransitive competitive relationships in cryptic coral reef environments. *Am. Nat.,* 113(2):223–234, 1979.

[41] M. J. Capolongo, M. D. DiBonaventura, and G. B. Chapman. Physician vaccinate thyself: why influenza vaccination rates are higher among clinicians than among nonclinicians. *Ann. Behav. Med.,* 31:288–296, 2006.

[42] T. J. Case and M. E. Gilpin. Interference competition and niche theory. *Proc. Natl. Acad. Sci. USA,* 71(8):3073, 1974.

[43] CDC, 2009. URL http://www.cdc.gov/vaccines/.

[44] S. H. Cheeseman, D. Havlir, M. M. McLaughlin, T. C. Greenough, J. L. Sullivan, D. Hall, S. E. Hattox, S. A. Spector, D. S. Stein, and M. Myers. Phase I/II evaluation of nevirapine alone and in combination with zidovudine for infection with human immunodeficiency virus. *J. Acquir. Immune Defic. Syndr.,* 8(2):141–151, 1995.

[45] F. Chen, W.-Q. Liu, A. Eisenstark, R. Johnston, G.-R. Liu, and S.-L. Liu. Multiple genetic switches spontaneously modulating bacterial mutability. *BMC Evol. Biol.,* 10(1):277, 2010. doi: 10.1186/1471-2148-10-277.

[46] J. Q. Chen, Y. Wu, H. Yang, J. Bergelson, M. Kreitman, and D. Tian. Variation in the ratio of nucleotide substitution and indel rates across genomes in mammals and bacteria. *Mol. Biol. Evol.,* 26:1523–1531, 2009.

[47] C. C. Chicone. *Ordinary Differential Equations with Applications.* Springer Verlag, 2006.

[48] T.-C. Chou. Derivation and properties of Michaelis-Menten type and hill type equations for reference ligands. *J. Theor. Biol.*, 59(2):253–276, 1976. doi: 16/0022-5193(76)90169-7.

[49] S. K. Choudhary and D. M. Margolis. Curing HIV: pharmacologic approaches to target HIV-1 latency. *Annu. Rev. Pharmacol. Toxicol.*, 51(1): 397–418, 2011. doi: 10.1146/annurev-pharmtox-010510-100237.

[50] T. W. Chun, D. Finzi, J. Margolick, K. Chadwick, D. Schwartz, and R. F. Siliciano. In vivo fate of HIV-1-infected t cells: quantitative analysis of the transition to stable latency. *Nat. Med.*, 1(12):1284–1290, 1995.

[51] T. W. Chun, L. Carruth, D. Finzi, X. Shen, J. A DiGiuseppe, H. Taylor, M. Hermankova, K. Chadwick, J. Margolick, T. C. Quinn, et al. Quantification of latent tissue reservoirs and total body viral load in HIV-1 infection. *Nature*, 387(6629):183–188, 1997.

[52] T. W. Chun, L. Stuyver, S. B. Mizell, L. A. Ehler, J. A. Mican, M. Baseler, A. L. Lloyd, M. A. Nowak, and A. S. Fauci. Presence of an inducible HIV-1 latent reservoir during highly active antiretroviral therapy. *Proc. Natl. Acad. Sci. USA*, 94(24):13193–13197, 1997.

[53] T. W. Chun, D. Engel, M. M. Berrey, T. Shea, L. Corey, and A. S. Fauci. Early establishment of a pool of latently infected, resting CD4(+) t cells during primary HIV-1 infection. *Proc. Natl. Acad. Sci. USA*, 95(15): 8869–8873, 1998.

[54] A. Cillo. Only a small fraction of HIV-1 proviruses in resting CD4+ t cells can be induced to produce virions ex vivo with anti-CD3/CD28 or vorinostat. In *20th Conference on Retroviruses and Opportunistic Infections*, Atlanta, 2013.

[55] J. M. Coffin. HIV population dynamics in vivo: implications for genetic variation, pathogenesis, and therapy. *Science*, 267(5197):483–489, 1995.

[56] M. G. Cojocaru. Dynamic equilibria of group vaccination strategies in a heterogeneous population. *J. Glob. Optim.*, 40:51–63, 2008.

[57] J. Colgrave. *State of Immunity: The Politics of Vaccination in Twentieth-Century America*. Univ. California Press, Berkeley, 2006.

[58] J. H. Condra. Resistance to HIV protease inhibitors. *Haemophilia*, 4: 610–615, 1998.

[59] J. M. Conway and D. Coombs. A stochastic model of latently infected cell reactivation and viral blip generation in treated HIV patients. *PLoS Comput. Biol.*, 7(4):e1002033, 2011.

[60] M. H. Cortez and S. P. Ellner. Understanding rapid evolution in predator-prey interactions using the theory of fast-slow dynamical systems. *Am. Nat.*, 176(5):e109–e127, 2010.

[61] R. Cressman. *Evolutionary dynamics and extensive form games.* the MIT Press, 2003.

[62] E. Dam, R. Quercia, B. Glass, D. Descamps, O. Launay, X. Duval, H.-G. Kräusslich, A. J. Hance, F. Clavel, and ANRS 109 Study Group. Gag mutations strongly contribute to HIV-1 resistance to protease inhibitors in highly Drug-Experienced patients besides compensating for fitness loss. *PLoS Pathog.*, 5(3):e1000345, 2009. doi: 10.1371/journal.ppat.1000345.

[63] R. T. Davey, Jr., N. Bhat, C. Yoder, T.-W. Chun, J. A. Metcalf, R. Dewar, V. Natarajan, R. A. Lempicki, J. W. Adelsberger, K. D. Miller, J. A. Kovacs, M. A. Polis, R. E. Walker, J. Falloon, H. Masur, D. Gee, M. Baseler, D. S. Dimitrov, A. S. Fauci, and H. C. Lane. HIV-1 and T cell dynamics after interruption of highly active antiretroviral therapy (haart) in patients with a history of sustained viral suppression. *Proc. Natl. Acad. Sci. USA*, 96 (26):15109–15114, 1999.

[64] R. J. De Boer, R. M. Ribeiro, and A. S. Perelson. Current estimates for HIV-1 production imply rapid viral clearance in lymphoid tissues. *PLoS Comput. Biol.*, 6(9):e1000906, 2010. doi: 10.1371/journal.pcbi.1000906.

[65] E. Denamur, G. Lecointre, P. Darlu, O. Tenaillon, C. Acquaviva, C. Sayada, I. Sunjevaric, R. Rothstein, J. Elion, F. Taddei, M. Radman, and I. Matic. Evolutionary implications of the frequent horizontal transfer of mismatch repair genes. *Cell*, 103(5):711–721, 2000.

[66] F. Dercole. *Evolutionary Dynamics Through Bifurcation Analysis: Methods and Applications.* PhD thesis, Department of Electronics and Information, Politecnico di Milano„ Milano, Italy, 2002.

[67] F. Dercole and S. Rinaldi. *Analysis of Evolutionary Processes: the Adaptive Dynamics Approach and its Applications.* Princeton University Press, 2008.

[68] F. Dercole, R. Ferrière, A. Gragnani, and S. Rinaldi. Coevolution of slow-fast populations: evolutionary sliding, evolutionary pseudo-equilibria and complex red queen dynamics. *Proc. R. Soc. B*, 273 (1589):983–990, 2006. doi: 10.1098/rspb.2005.3398.

[69] M. M. Desai and D. S. Fisher. The balance between mutators and nonmutators in asexual populations. *Genetics*, 188(4):997–1014, 2011.

[70] U. Dieckmann and R. Law. The dynamical theory of coevolution: a derivation from stochastic ecological processes. *J. Math. Biol.*, 34(5): 579–612, 1996.

[71] U. Dieckmann, P. Marrow, and R. Law. Evolutionary cycling in predator-prey interactions: population dynamics and the red queen. *J. Theor. Biol.*, 176(1):91–102, 1995. doi: 10.1006/jtbi.1995.0179.

[72] O. Diekmann and J. A. P. Heesterbeek. *Mathematical epidemiology of infectious diseases: model building, analysis, and interpretation.* John Wiley & Sons, Chichester, UK, 2000.

[73] S.P. Diggle, A.S. Griffin, G.S. Campbell, and S.A. West. Cooperation and conflict in quorum-sensing bacterial populations. *Nature*, 450(7168): 411–414, 2007.

[74] J. B. Dinoso, S. Y. Kim, A. M. Wiegand, S. E. Palmer, S. J. Gange, L. Cranmer, A. O'Shea, M. Callender, A. Spivak, T. Brennan, et al. Treatment intensification does not reduce residual HIV-1 viremia in patients on highly active antiretroviral therapy. *Proc. Natl. Acad. Sci. USA*, 106(23):9403, 2009.

[75] M. Doebeli and J. C. Koella. Evolution of simple population dynamics. *Proc. R. Soc. B*, 260(1358):119–125, 1995.

[76] Michael Doebeli. Evolutionary predictions from invariant physical measures of dynamic processes. *J. Theor. Biol.*, 173(4):377—387, 1995. doi: 10.1006/jtbi.1995.0070.

[77] A. Donald and V. Muthu. MMR links with autism and inflammatory bowel disease. *Clin. Evid.*, 7:331–340, 2002.

[78] J. W. Drake, B Charlesworth, D Charlesworth, and J. F. Crow. Rates of spontaneous mutation. *Genetics*, 148:1667–1686, 1998.

[79] K. Drlica. The mutant selection window and antimicrobial resistance. *J. Antimicrob. Chemother.*, 52(1):11–17, 2003. doi: 10.1093/jac/dkg269.

[80] K. Drlica and X. Zhao. Mutant selection window hypothesis updated. *Clin. Infect. Dis.*, 44(5):681 –688, 2007. doi: 10.1086/511642.

[81] C. M. Durand, J. N. Blankson, and R. F. Siliciano. Developing strategies for HIV-1 eradication. *Trends Immunol.*, 33(11):554–562, 2012. doi: 10.1016/j.it.2012.07.001.

[82] L. Duret. Mutation patterns in the human genome: more variable than expected. *PLoS Biol.*, 7:e1000028, 2009.

[83] R. Durrett and S. Levin. The importance of being discrete (and spatial). *Theor. Pop. Biol.*, 1994.

[84] K. T. D. Eames. Networks of influence and infection: parental choices and childhood disease. *J. R. Soc. Interface*, 6:811–814, 2009.

[85] D. J. Earl and M. W. Deem. Evolvability is a selectable trait. *Proc. Natl. Acad. Sci. USA*, 101(32):11531–11536, 2004.

[86] J. P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.*, 57(3):617–656, 1985. doi: 10.1103/RevModPhys.57.617.

[87] N. Elango, S. H. Kim, NISC Comparative Sequencing Program, E. Vigoda, and V. Y. Soojin. Mutations of different molecular origins exhibit contrasting patterns of regional substitution rate variation. *PLoS Comp. Biol.*, 4:e1000015, 2008.

[88] P. & Rényi A. Erdős. On random graphs. *Publ. Math Debrecen*, 6:290–297, 1959.

[89] S. Eriksson, E. H. Graf, V. Dahl, M. C. Strain, S. A. Yukl, E. S. Lysenko, R. J. Bosch, J. Lai, S. Chioma, F. Emad, M. Abdel-Mohsen, R. Hoh, F. Hecht, P. Hunt, M. Somsouk, J. Wong, R. Johnston, R. F. Siliciano, D. D. Richman, U. O'Doherty, S. Palmer, S. G. Deeks, and J. D. Siliciano. Comparative analysis of measures of viral reservoirs in HIV-1 eradication

studies. *PLoS Pathog*, 9(2):e1003174, 2013. doi: 10.1371/journal.ppat.1003174.

[90] W. J. Ewens. *Mathematical Population Genetics*. I. Theoretical Introduction. Springer Verlag, 2004.

[91] P. Fine and J. Clarkson. Individual versus public priorities in the determination of optimal vaccination policies. *Am. J. Epidemiol.*, 124: 1012–1020, 1986.

[92] D. Finzi, M. Hermankova, T. Pierson, L. M. Carruth, C. Buck, R. E. Chaisson, T. C. Quinn, K. Chadwick, J. Margolick, R. Brookmeyer, J. Gallant, M. Markowitz, D. D. Ho, D. D. Richman, and R. F. Siliciano. Identification of a reservoir for HIV-1 in patients on highly active antiretroviral therapy. *Science*, 278(5341):1295 –1300, 1997. doi: 10.1126/science.278.5341.1295.

[93] S. D. Frost, M. Nijhuis, R. Schuurman, C. A. Boucher, and A. J. Leigh Brown. Evolution of lamivudine resistance in human immunodeficiency virus type 1-infected individuals: the relative roles of drift and selection. *J. Virol.*, 74(14):6262–6268, 2000.

[94] H. R. Fryer, J. Frater, A. Duda, M. G. Roberts, SPARTAC Trial Investigators, R. E. Phillips, and A. R. McLean. Modelling the evolution and spread of HIV immune escape mutants. *PLoS Pathog.*, 6(11): e1001196, 2010.

[95] K. J. Fryxell and W. J. Moon. CpG mutation rates in the human genome are highly dependent on local GC content. *Mol. Biol. Evol.*, 22:650–658, 2005.

[96] F. Fu, D. I. Rosenbloom, L. Wang, and M. A. Nowak. Imitation dynamics of vaccination behaviour on social networks. *Proc. R. Soc. B*, 278(1702): 42–49, 2011.

[97] G. A. Funk, M. Fischer, B. Joos, M. Opravil, H. F. Günthard, B. Ledergerber, and S. Bonhoeffer. Quantification of in vivo replicative capacity of HIV-1 in different compartments of infected cells. *J. Acquir. Immune Defic. Syndr.*, 26(5):397–404, 2001.

[98] A. P. Galvani, T. C. Reluga, and G. B. Chapman. Long-standing influenza vaccination policy is in accord with individual self-interest but not with the utilitarian optimum. *Proc. Natl. Acad. Sci. USA*, 104:5692–5697, 2007.

[99] E. M. Gardner, S. Sharma, G. Peng, K. H. Hullsiek, W. J. Burman, R. D. Macarthur, M. Chesney, E. E. Telzak, G. Friedland, and S. B. Mannheimer. Differential adherence to combination antiretroviral therapy is associated with virological failure with resistance. *AIDS*, 22(1):75–82, 2008. doi: 10.1097/QAD.0b013e3282f366ff.

[100] E. M. Gardner, W. J. Burman, J. F. Steiner, P. L. Anderson, and D. R. Bangsberg. Antiretroviral medication adherence and the development of class-specific antiretroviral resistance. *AIDS*, 23(9):1035–1046, 2009. doi: 10.1097/QAD.0b013e32832ba8ec.

[101] A. Gaunersdorfer. Time averages for heteroclinic attractors. *SIAM J. Appl. Math.*, 52(5):pp. 1476–1489, 1992.

[102] S. A. H. Geritz. Resident-invader dynamics and the coexistence of similar strategies. *J. Math. Biol.*, 50(1):67–82, 2005.

[103] S. A. H. Geritz, E. Kisdi, G. Meszeńa, and J. A. J. Metz. Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evol. Ecol.*, 12(1):35–57, 1997.

[104] S. A. H. Geritz, M. Gyllenberg, F. J. A. Jacobs, and K. Parvinen. Invasion dynamics and attractor inheritance. *J. Math. Biol.*, 44:548–560, 2002. doi: 10.1007/s002850100136.

[105] P. J. Gerrish, A. Colato, A. S. Perelson, and P. D. Sniegowski. Complete genetic linkage can subvert natural selection. *Proc. Natl. Acad. Sci. USA*, 104(15):6266–6271, 2007.

[106] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81:2340–2361, 1977.

[107] A. Giraud, I. Matic, O. Tenaillon, A. Clara, M. Radman, M. Fons, and F. Taddei. Costs and benefits of high mutation rates: adaptive evolution of bacteria in the mouse gut. *Science*, 291(5513):2606–2608, 2001.

[108] J. Gore, H. Youk, and A. Van Oudenaarden. Snowdrift game dynamics and facultative cheating in yeast. *Nature*, 458(7244):253–256, 2009.

[109] A.S. Griffin, S.A. West, and A. Buckling. Cooperation and competition in pathogenic bacteria. *Nature*, 430(7003):1024–1027, 2004.

[110] E. Gullberg, S. Cao, O. G. Berg, C. Ilbäck, L. Sandegren, D. Hughes, and D. I. Andersson. Selection of resistant bacteria at very low antibiotic concentrations. *PLoS Pathog.*, 7(7):e1002158, 2011. doi: 10.1371/journal.ppat.1002158.

[111] R. K. Gupta, A. Hill, A. W. Sawyer, A. Cozzi-Lepri, V. von Wyl, S. Yerly, V. D. Lima, H. F. Günthard, C. Gilks, and D. Pillay. Virological monitoring and resistance to first-line highly active antiretroviral therapy in adults infected with HIV-1 treated under who guidelines: a systematic review and meta-analysis. *Lancet Infect. Dis.*, 9:409–417, 2009.

[112] R. K. Gupta, A. Kohli, A. L. McCormick, G. J. Towers, D. Pillay, and C. M. Parry. Full length HIV-1 gag determines protease inhibitor susceptibility within in vitro assays. *AIDS*, 24(11):1651–1655, 2010.

[113] M. Gyllenberg, A. Osipov, and G. Söderbacka. Bifurcation analysis of a metapopulation model with sources and sinks. *J. Nonlin. Sci.*, 6:329–366, 1996. doi: 10.1007/BF02433474.

[114] K. P. Hadeler. Stable polymorphisms in a selection model with mutation. *SIAM J. Appl. Math.*, 41(1):1–7, 1981.

[115] J. Haigh. The accumulation of deleterious genes in a population. *Theor. Pop. Biol.*, 14(2):251–267, 1978.

[116] G. Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1968.

[117] P. R. Harrigan, R. S. Hogg, W. W. Y. Dong, B. Yip, B. Wynhoven, J. Woodward, C. J. Brumme, Z. L. Brumme, T. Mo, C. S. Alexander, and J. S. G. Montaner. Predictors of HIV drug-resistance mutations in a large antiretroviral-naive cohort initiating triple antiretroviral therapy. *J. Infect. Dis.*, 191(3):339–347, 2005. doi: 10.1086/427192.

[118] D. V. Havlir, N. S. Hellmann, C. J. Petropoulos, J. M. Whitcomb, A. C. Collier, M. S. Hirsch, P. Tebas, J.-P. Sommadossi, and D. D. Richman. Drug susceptibility in HIV infection after viral rebound in patients receiving Indinavir-Containing regimens. *J. Am. Med. Ass.*, 283(2):229–234, 2000. doi: 10.1001/jama.283.2.229.

[119] J. M. Heffernan, R. J. Smith, and L. M. Wahl. Perspectives on the basic reproductive ratio. *J. R. Soc. Interface*, 2:281–293, 2005.

[120] M. Heino, J. A. J. Metz, and V. Kaitala. The enigma of frequency-dependent selection. *Trends Ecol. Evol.*, 13(9):367–370, 1998.

[121] T. J Henrich, Z. Hu, J. Z. Li, G. Sciaranghella, M. P Busch, S. M. Keating, S. Gallien, N. H. Lin, F. F. Giguel, L. Lavoie, V. T. Ho, P. Armand, R. J. Soiffer, M. Sagar, A. S. Lacasce, and D. R. Kuritzkes. Long-Term Reduction in Peripheral Blood HIV-1 Reservoirs Following Reduced-Intensity Conditioning Allogeneic Stem Cell Transplantation. *J. Inf. Dis.*, 2013.

[122] A. V. Herz, S. Bonhoeffer, R. M. Anderson, R. M. May, and M. A. Nowak. Viral dynamics in vivo: limitations on estimates of intracellular delay and virus decay. *Proc. Natl. Acad. Sci. USA*, 93(14):7247–7251, 1996.

[123] A. L. Hill, D. I. S. Rosenbloom, and M. A. Nowak. Evolutionary dynamics of HIV at multiple spatial and temporal scales. *J. Mol. Med.*, 2012. doi: 10.1007/s00109-012-0892-1.

[124] M.S. Hirsch, F. Brun-Vézinet, R.T. D'Aquila, S.M. Hammer, V.A. Johnson, D.R. Kuritzkes, C. Loveday, J.W. Mellors, B. Clotet, B. Conway, et al. Antiretroviral drug resistance testing in adult HIV-1 infection: recommendations of an International AIDS Society–USA panel. *J. Am. Med. Ass.*, 283:2417–2426, 2000.

[125] Y.-C. Ho, L. Shan, J. Wang, N. Hosmane, J. N. Blankson, and R. F. Siliciano. Characterization of non--induced HIV--1 proviruses dampens the hope for HIV--1 eradication. In *20th Conference on Retroviruses and Opportunistic Infections*, Atlanta, GA, 2013.

[126] R. D. Hockett, J. M. Kilby, C. A. Derdeyn, M. S. Saag, M. Sillers, K. Squires, S. Chiz, M. A. Nowak, G. M. Shaw, and R. P. Bucy. Constant mean viral copy number per infected cell in tissues regardless of high, low, or undetectable plasma HIV RNA. *J. Exp. Med.*, 189(10):1545 –1554, 1999. doi: 10.1084/jem.189.10.1545.

[127] A. Hodgkinson, E. Ladoukakis, and A. Eyre-Walker. Cryptic variation in the human mutation rate. *PLoS Biol.*, 7:e1000027, 2009.

[128] J. Hofbauer. The selection mutation equation. *J. Math. Biol.*, 23(1):41–53, 1985.

[129] J. Hofbauer. Heteroclinic cycles in ecological differential equations. *Tatra Mountains Math. Publ.*, 4:105–116, 1994.

[130] J. Hofbauer. Heteroclinic cycles in ecological differential equations. *Tatra Mt. Math. Publ.*, 4:105–116, 1994.

[131] J. Hofbauer and K. Sigmund. Adaptive dynamics and evolutionary stability. *Appl. Math. Lett.*, 3(4):75–79, 1990.

[132] J. Hofbauer and K. Sigmund. *Evolutionary Games & Population Dynamics.* Cambridge University Press, Cambridge, UK, 1998.

[133] J. Hofbauer and K. Sigmund. *Evolutionary games and population dynamics.* Cambridge Univ. Press, Cambridge, 1998.

[134] J. Hofbauer and K. Sigmund. Evolutionary game dynamics. *B. Am. Math. Soc.*, 40(4):479–520, 2003.

[135] J. Hofbauer, P. Schuster, and K. Sigmund. A note on evolutionary stable strategies and game dynamics. *J. Theor. Biol.*, 81(3):609–612, 1979.

[136] J. Hofbauer, J. Mallet-Paret, and H. L. Smith. Stable periodic solutions for the hypercycle system. *J. Dyn. Differ. Equ.*, 3(3):423–436, 1991.

[137] C. J. Hoffmann, S. Charalambous, J. Sim, J. Ledwaba, G. Schwikkard, R. E. Chaisson, K. L. Fielding, G. J. Churchyard, L. Morris, and A. D. Grant. Viremia, resuppression, and time to resistance in human immunodeficiency virus (HIV) subtype c during first-line antiretroviral therapy in south africa. *Clin. Infect. Dis.*, 49:1928–1935, 2009.

[138] R. A. Horn and C. R. Johnson. *Matrix Analysis.* Cambridge Univ. Press, Cambridge, 1990.

[139] G. Hütter, D. Nowak, M. Mossner, S. Ganepola, A. Müßig, K. Allers, T. Schneider, J. Hofmann, C. Kücherer, O. Blau, I. W. Blau, W. K. Hofmann, and E. Thiel. Long-term control of HIV by CCR5 Delta32/Delta32 stem-cell transplantation. *N. Engl. J. Med.*, 360(7): 692–698, 2009.

[140] J. P. A. Ioannidis, D. V. Havlir, P. Tebas, M. S. Hirsch, A. C. Collier, and D. D. Richman. Dynamics of HIV-1 viral load rebound among patients with previous suppression of viral replication. *AIDS*, 14:1481–1488, 2000.

[141] K. Ishii, H. Matsuda, Y. Iwasa, and A. Sasaki. Evolutionarily stable mutation rate in a periodically changing environment. *Genetics*, 121: 163–174, 1989.

[142] V. A. Jansen, N. Stollenwerk, and H. J. Jensen. Measles outbreaks in a population with declining vaccine uptake. *Science*, 301:804, 2003.

[143] M. I. Jensen-Seaman, T. S. Furey, B. A. Payseur, Y. Lu, K. M. Roskin, C.-F. Chen, M. A. Thomas, D. Haussler, and H. J. Jacob. Comparative recombination rates in the rat, mouse, and human genomes. *Genome Res.*, 14:528–538, 2004.

[144] B. L. Jilek, M. Zarr, M. E. S. Sampah, S. A. Rabi, C. K. Bullen, J. Lai, L. Shen, and R. F. Siliciano. A quantitative basis for antiretroviral therapy for HIV-1 infection. *Nat. Med.*, 18(3):446–451, 2012. doi: 10.1038/nm.2649.

[145] E. J. Johnson and A. Tversky. Affect, generalization, and the perception of risk. *J Pers Soc Psychol*, 45:20–31, 1983.

[146] T. Johnson. The approach to mutation-selection balance in an infinite asexual population, and the evolution of mutation rates. *Proc. R. Soc. B*, 266(1436):2389–2397, 1999.

[147] T. Johnson. Beneficial mutations, hitchhiking and the evolution of mutation rates in sexual populations. *Genetics*, 151(4):1621–1631, 1999.

[148] T. Johnson. The approach to mutation-selection balance in an infinite asexual population, and the evolution of mutation rates. *Proc. R. Soc. B*, 266:2389–2397, 1999.

[149] G. Jourdain, T. A. Wagner, N. Ngo-Giang-Huong, W. Sirirungsi, V. Klinbuayaem, F. Fregonese, I. Nantasen, M. Techapornroong, G. Halue, A. Nilmanat, P. Wittayapraparat, V. Chalermpolprapa, P. Pathipvanich, P. Yuthavisuthi, L. M. Frenkel, M. Lallemant, and PHPT Study Group. Association between detection of HIV-1 DNA resistance mutations by a sensitive assay at initiation of antiretroviral therapy and virologic failure. *Clin. Infect. Dis.*, 50:1397–1404, 2010.

[150] C. Kamp, C. O. Wilke, C. Adami, and S. Bornholdt. Viral evolution under the pressure of an adaptive immune sytem: Optimal mutation rates for viral escape. *Complexity*, 8(2):28–33, 2003.

[151] E. L. Kaplan and P. Meier. Nonparametric estimation from incomplete observations. *J. Amer. Statist. Assoc.*, 53(282):457–481, 1958. doi: 10.2307/2281868.

[152] S. Karlin and H. E Taylor. *A First Course in Stochastic Processes*. Academic Press, San Diego, 1975.

[153] H. Kastrissios, J. R. Suárez, D. Katzenstein, P. Girard, L. B. Sheiner, and T.F. Blaschke. Characterizing patterns of drug-taking behavior with a multiple drug regimen in an AIDS clinical trial. *AIDS*, 12(17):2295, 1998.

[154] M. J. Keeling and K. T. D. Eames. Networks and epidemic models. *J. R. Soc. Interface*, 2(4):295–307, 2005. doi: 10.1098/rsif.2005.0051.

[155] D. J. Kempf, M. S. King, B. Bernstein, P. Cernohous, E. Bauer, J. Moseley, K. Gu, A. Hsu, S. Brun, and E. Sun. Incidence of resistance in a double-blind study comparing lopinavir/ritonavir plus stavudine and lamivudine to nelfinavir plus stavudine and lamivudine. *J. Infect. Dis.*, 189: 51–60, 2004.

[156] T. B. Kepler and A. S. Perelson. Drug concentration heterogeneity facilitates the evolution of drug resistance. *Proc. Natl. Acad. Sci. USA*, 95 (20):11514–11519, 1998. doi: 10.1073/pnas.95.20.11514.

[157] B. Kerr, M. A. Riley, M. W. Feldman, and B. J. M. Bohannan. Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature*, 418(6894):171–174, 2002.

[158] D. A. Kessler and H. Levine. Mutator dynamics on a smooth evolutionary landscape. *Phys. Rev. Lett.*, 80:2012–2015, 1998.

[159] A. I. Khibnik and A. S. Kondrashov. Three mechanisms of red queen dynamics. *Proc. R. Soc. B*, 264(1384):1049–1056, 1997. doi: 10.1098/rspb.1997.0145.

[160] P. Kiepiela, A.J. Leslie, I. Honeyborne, D. Ramduth, C. Thobakgale, S. Chetty, P. Rathnavalu, C. Moore, K.J. Pfafferott, L. Hilton, et al.

Dominant influence of HLA-B in mediating the potential co-evolution of HIV and HLA. *Nature*, 432(7018):769–775, 2004.

[161] H. Kim and A. S. Perelson. Viral and latent reservoir persistence in HIV-1–Infected patients on therapy. *PLoS Comput Biol*, 2(10):e135, 2006. doi: 10.1371/journal.pcbi.0020135.

[162] M. Kimura. Optimum mutation rate and degree of dominance as determined by the principle of minimum genetic load. *J. Genet.*, 57:21–34, 1960.

[163] M. Kimura. On the evolutionary adjustment of spontaneous mutation rates. *Genet. Res.*, 9:23–34, 1967.

[164] D. G King, M. Soller, and Y. Kashi. Evolutionary tuning knobs. *Endeavour*, 21:36–40, 1997.

[165] B. C. Kirkup and M. A. Riley. Antibiotic-mediated antagonism leads to a bacterial game of rock-paper-scissors in vivo. *Nature*, 428(6981): 412–414, 2004.

[166] M. Kirschner and J. Gerhart. Evolvability. *Proc. Natl. Acad. Sci. USA*, 95 (15):8420–8427, 1998.

[167] R. D. Kouyos, C. L. Althaus, and S. Bonhoeffer. Stochastic or deterministic: what is the effective population size of HIV-1? *Trends Microbiol.*, 14(12):507–511, 2006.

[168] O. Krakovska and L. M. Wahl. Costs versus benefits: best possible and best practical treatment regimens for HIV. *J. Math. Biol.*, 54:385–406, 2007.

[169] V. Krivan and R. Cressman. On evolutionary stability in prey-predator models with fast behavioral dynamics. *Evol. Ecol. Res.*, 11:227–251, 2009.

[170] C. Kuiken, B. Foley, T. Leitner, C. Apetrei, B. Hahn, I. Mizrachi, J Mullins, A. Rambaut, S. Wolinsky, and B. Korber, editors. *HIV Sequence Compendium 2010*. Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, NM, 2010.

[171] R. J. Kulathinal, S. M. Bennett, C. L. Fitzpatrick, and M. A. F. Noor. Fine-scale mapping of recombination rate in *Drosophila* refines its correlation to diversity and divergence. *Proc. Natl. Acad. Sci. USA*, 105: 10051–10056, 2008.

[172] G. Kunzel, L. Korzun, and A. Meister. Cytologically integrated physical restriction fragment length polymorphism maps for the barley genome based on translocation breakpoints. *Genetics*, 154:397–412, 2000.

[173] G. I. Lang. Mutation rate variation in the yeast, *Saccharomyces cerevisiae*. Harvard University, 2007. PhD Thesis.

[174] E. G. Leigh. Natural selection and mutability. *Am. Nat.*, 104:301–305, 1970.

[175] E. G. Leigh. The evolution of mutation rates. *Genetics*, 73:1–18, 1973.

[176] B. R. Levin, M. Lipsitch, and Bonhoeffer S. Population biology, evolution, and infectious disease: convergence and synthesis. *Science*, 283:806–809, 1999.

[177] S. A. Levin. Community equilibria and stability, and an extension of the competitive exclusion principle. *Am. Nat.*, 104(939):413–423, 1970.

[178] G. Levinson and G. A. Gutman. High frequencies of short frameshifts in poly-CA/TG tandem repeats borne by bacteriophage M13 in *Escherichia coli* K-12. *Nucleic Acids Res.*, 15:5323–5338, 1987.

[179] G. Levinson and G. A. Gutman. Slipped-strand mispairing: a major mechanism for DNA evolution. *Mol. Biol. Evol.*, 4:203–221, 1987.

[180] V. D. Lima, R. Harrigan, M. Murray, D. M. Moore, E. Wood, R. S. Hogg, and J. S. Montaner. Differential impact of adherence on long-term treatment response among naive HIV-infected individuals. *AIDS*, 22: 2371–2380, 2008.

[181] S. J. Little, A. R. McLean, C. A. Spina, D. D. Richman, and D. V. Havlir. Viral dynamics of acute HIV-1 infection. *J. Exp. Med.*, 190(6):841–850, 1999. doi: 10.1084/jem.190.6.841.

[182] H. Liu, L. G. Miller, C. E. Golin, R. D. Hays, T. Wu, N. S. Wenger, and
A. H. Kaplan. Repeated measures analyses of dose timing of antiretroviral
medication and its relationship to HIV virologic outcomes. *Stat. Med.*, 26
(5):991–1007, 2007. doi: 10.1002/sim.2592.

[183] M. Lynch. *The Origins of Genome Architecture*. Sinauer Associates, Inc.,
2007.

[184] M. Lynch, R. Bürger, D. Butcher, and W. Gabriel. The mutational
meltdown in asexual populations. *J. Hered.*, 84(5):339–344, 1993.

[185] F. Maggiolo, M. Airoldi, H. D. Kleinloog, A. Callegaro, V. Ravasio,
C. Arici, E. Bombana, and F. Suter. Effect of adherence to HAART on
virologic outcome and on the selection of resistance-conferring mutations
in NNRTI- or PI-treated patients. *HIV Clin. Trials*, 8(5):282–292, 2007.
doi: 10.1310/hct0805-282.

[186] J. R. Magnus and H. Neudecker. *Matrix Differential Calculus with
Applications in Statistics and Econometrics*. Wiley, New York, 1988.

[187] L. M. Mansky and H. M. Temin. Lower in vivo mutation rate of human
immunodeficiency virus type 1 than that predicted from the fidelity of
purified reverse transcriptase. *J. Virol.*, 69(8):5087–5094, 1995.

[188] M. Markowitz, M. Louie, A. Hurley, E. Sun, M. Di Mascio, A. S. Perelson,
and D. D. Ho. A novel antiviral intervention results in more accurate
assessment of human immunodeficiency virus type 1 replication
dynamics and T-Cell decay in vivo. *J. Virol.*, 77(8):5037–5038, 2003. doi:
10.1128/JVI.77.8.5037-5038.2003.

[189] G. van Marle, M. J. Gill, D. Kolodka, L. McManus, T. Grant, and D. L.
Church. Compartmentalization of the gut viral reservoir in HIV-1
infected patients. *Retrovirology*, 4:87, 2007.

[190] P. Marrow, R. Law, and C. Cannings. The coevolution of predator–prey
interactions: Esss and red queen dynamics. *Proc. R. Soc. B*, 250(1328):
133–141, 1992. doi: 10.1098/rspb.1992.0141.

[191] I. Matic, M. Radman, F. Taddei, B. Picard, C. Doit, E. Bingen, E. Denamur,
and J. Elion. Highly variable mutation rates in commensal and pathogenic
Escherichia coli. *Science*, 277(5333):1833–1834, 1997.

[192] R. M. May. Limit cycles in predator-prey communities. *Science*, 177 (4052):900–902, 1972. doi: 10.1126/science.177.4052.900.

[193] R. M. May. *Stability and Complexity in Model Ecosystems*. Princeton University Press, Princeton, NJ, USA, 2001.

[194] R. M. May and W. J. Leonard. Nonlinear aspects of competition between three species. *SIAM J. Appl. Math.*, 29(2):243–253, 1975.

[195] J. Maynard Smith. *Evolution and the theory of games*. Cambridge Univ. Press, Cambridge, 1982.

[196] J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246(5427):15–18, 1973.

[197] R. Mehla, S. Bivalkar-Mehla, R. Zhang, I. Handy, H. Albrecht, S. Giri, P. Nagarkatti, M. Nagarkatti, and A. Chauhan. Bryostatin modulates latent HIV-1 infection via PKC and AMPK signaling but inhibits acute infection in a receptor independent manner. *PloS ONE*, 5(6):e11160, 2010. doi: 10.1371/journal.pone.0011160.

[198] J. A. J. Metz, R. M. Nisbet, and S. A. H. Geritz. How should we define 'fitness' for general ecological scenarios? *Trends Ecol. Evol.*, 7(6):198–202, 1992. doi: 10.1016/0169-5347(92)90073-K.

[199] J. A. J. Metz, S. A. H. Geritz, G. Meszéna, F.J. A. Jacobs, and J. S. van Heerwaarden. Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. In S. J. van Strien and S. M. Verduyn Lunel, editors, *Stochastic and Spatial Structures of Dynamical Systems*, pages pp 183–231. KNAW Verhandelingen, Afd., Amsterdam, 1996.

[200] C. Mezard. Meiotic recombination hotspots in plants. *Biochem. Soc. T.*, 34:531–534, 2006.

[201] L. K. M'Gonigle, J. J. Shen, and S. P. Otto. Mutating away from your enemies: the evolution of mutation rate in a host–parasite system. *Theor. Pop. Biol.*, 2009.

[202] J.-B. Michel, P. J. Yeh, R. Chait, R. C. Moellering, and R. Kishony. Drug interactions modulate the potential for evolution of resistance. *Proc. Natl. Acad. Sci. USA*, 105:14918–14923, 2008.

[203] M. Milinski. Tit for tat in sticklebacks and the evolution of cooperation. *Nature*, 325(6103):433–435, 1987.

[204] E. R. Moxon, P. B. Rainey, M. A. Nowak, and R. E. Lenski. Adaptive evolution of highly mutable loci in pathogenic bacteria. *Curr. Biol.*, 4(1): 24–33, 1994.

[205] G. L. Murphy, T. D. Connell, D. S. Barritt, M. Koomey, and J. G. Cannon. Phase variation of gonococcal protein II: regulation of gene expression by slipped-strand mispairing of a repetitive DNA sequence. *Cell*, 56: 539–547, 1989.

[206] R. A. Neher and T. Leitner. Recombination rate and selection strength in HIV intra-patient evolution. *PLoS Comput. Biol.*, 6(1):e1000660, 2010.

[207] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

[208] A. Nicoll, D. Elliman, and E. Ross. MMR vaccination and autism 1998: Déjà vu–pertussis and brain damage 1974? *Brit. Med. J.*, 316:715–716, 1998.

[209] M. Nijhuis, N. M. van Maarseveen, S. Lastere, P. Schipper, E. Coakley, B. Glass, M. Rovenska, D. de Jong, C. Chappey, I. W. Goedegebuure, G. Heilek-Snyder, D. Dulude, N. Cammack, L. Brakier-Gingras, J. Konvalinka, N. Parkin, H.-G. Kräusslich, F. Brun-Vezinet, and C. A. B. Boucher. A novel substrate-based HIV-1 protease inhibitor drug resistance mechanism. *PLoS Med.*, 4(1):e36, 2007. doi: 10.1371/journal.pmed.0040036.

[210] A. Noe, J. Plum, and C. Verhofstede. The latent HIV-1 reservoir in patients undergoing HAART: an archive of pre-HAART drug resistance. *J. Antimicrob. Chemother.*, 55:410–412, 2005.

[211] M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314: 1560–1563, 2006.

[212] M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314 (5805):1560–1563, 2006.

[213] M. A. Nowak. *Evolutionary Dynamics: Exploring the equations of life*. Harvard Univ. Press, Cambridge, MA, 2006.

[214] M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, 1992.

[215] M. A. Nowak and R. M. C. May. *Virus Dynamics: Mathematical principles of immunology and virology*. Oxford University Press, USA, 2000.

[216] M. A. Nowak and K. Sigmund. The evolution of stochastic strategies in the prisoner's dilemma. *Acta Appl. Math.*, 20(3):247–265, 1990.

[217] M. A. Nowak and K. Sigmund. Evolutionary dynamics of biological games. *Science*, 303:793–799, 2004.

[218] M. A. Nowak and K. Sigmund. Evolutionary dynamics of biological games. *Science*, 303(5659):793–799, 2004. doi: 10.1126/science.1093411.

[219] M. A. Nowak, R. M. Anderson, A. R. McLean, T. F. W. Wolfs, J. Goudsmit, and R. M. May. Antigenic diversity thresholds and the development of AIDS. *Science*, 254:963–969, 1991.

[220] M A Nowak, R. M. May, R E Phillips, S. Rowland-Jones, D. G. Lalloo, S. McAdam, P. Klenerman, B. Köppe, K. Sigmund, and C. R. Bangham. Antigenic oscillations and shifting immunodominance in HIV-1 infections. *Nature*, 375(6532):606–611, 1995.

[221] M. A. Nowak, N. L. Komarova, and P. Niyogi. Evolution of universal grammar. *Science*, 291(5501):114–118, 2001. doi: 10.1126/science.291.5501.114.

[222] M. A. Nowak, C. E. Tarnita, and T. Antal. Evolutionary dynamics in structured populations. *Phil. Trans. R. Soc. Lond. B*, 365:19–30, 2010.

[223] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441: 502–505, 2006.

[224] A. Oliver, R. Cantón, P. Campo, F. Baquero, and J. Blázquez. High frequency of hypermutable Pseudomonas aeruginosa in cystic fibrosis lung infection. *Science*, 288(5469):1251, 2000.

[225] Panel on Antiretroviral Guidelines for Adults and Adolescents. *Guidelines for the use of antiretroviral agents in HIV-1-infected adults and adolescents*. U.S. Department of Health and Human Services, 2011. URL

http://aidsinfo.nih.gov/contentfiles/
AdultandAdolescentGL.pdf.

[226] H. A. Orr. The rate of adaptation in asexuals. *Genetics*, 155(2):961–968, 2000.

[227] C. Pal, M. D. Maciá, A. Oliver, I. Schachar, and A. Buckling. Coevolution with viruses drives the evolution of bacterial mutation rates. *Nature*, 450: 1079–1081, 2007.

[228] R. Palekar, A. Pettifor, F. Behets, and C. MacPhail. Association between knowing someone who died of AIDS and behavior change among South African youth. *AIDS Behav.*, 12:903–912, 2008.

[229] F. J. Palella, K. M. Delaney, A. C. Moorman, M. O. Loveless, J. Fuhrer, G. A. Satten, D. J. Aschman, and S. D. Holmberg. Declining morbidity and mortality among patients with advanced human immunodeficiency virus infection. HIV Outpatient Study Investigators. *N. Engl. J. Med.*, 338: 853–860, 1998.

[230] T. R. Palfrey and H. Rosenthal. Participation and the provision of discrete public goods: A strategic analysis. *J. Pub. Econ.*, 24:171–193, 1984.

[231] C. E. Paquin and J. Adams. Relative fitness can decrease in evolving asexual populations of s. cerevisiae. *Nature*, pages 368–371, 1983. doi: 10.1038/306368a0.

[232] R. Paredes, C. M. Lalama, H. J. Ribaudo, B. R. Schackman, C. Shikuma, F. Giguel, W. A. Meyer III, V. A. Johnson, S. A. Fiscus, R. T. D'Aquilla, R. M. Gulick, D. R. Kuritzkes, and ACTG A5095 Study Team. Pre-existing minority drug-resistant HIV-1 variants, adherence, and risk of antiretroviral treatment failure. *J. Infect. Dis.*, 201:662–671, 2010.

[233] J.-J. Parienti, M. Das-Douglas, V. Massari, D. Guzman, S. G. Deeks, R. Verdon, and D. R. Bangsberg. Not all missed doses are the same: sustained NNRTI treatment interruptions predict HIV rebound at low-to-moderate adherence levels. *PloS ONE*, 3(7):e2783, 2008. doi: 10.1371/journal.pone.0002783.

[234] J.-J. Parienti, K. Ragland, F. Lucht, A. de la Blanchardiére, S. Dargére, Y. Yazdanpanah, J.-J. Dutheil, P. Perré, R. Verdon, and D. R. Bangsberg.

Average adherence to boosted protease inhibitor therapy, rather than the pattern of missed doses, as a predictor of HIV RNA replication. *Clin. Infect. Dis.*, 50(8):1192–1197, 2010. doi: 10.1086/651419.

[235] C. M Parry, A. Kohli, C. J. Boinett, G. J. Towers, A. L. McCormick, and D. Pillay. Gag determinants of fitness and drug susceptibility in protease Inhibitor-Resistant human immunodeficiency virus type 1. *J. Virol.*, 83 (18):9094–9101, 2009. doi: 10.1128/JVI.02356-08.

[236] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.*, 86(14):3200–3203, 2001. doi: 10.1103/PhysRevLett.86.3200.

[237] R. Pastor-Satorras and A. Vespignani. Immunization of complex networks. *Phys. Rev. E*, 65:036104, 2002.

[238] P. S. Pennings. Standing genetic variation and the evolution of drug resistance in HIV. *PLoS Comput Biol*, 8(6):e1002527, 2012. doi: 10.1371/journal.pcbi.1002527.

[239] A. S. Perelson, A. U. Neumann, M. Markowitz, J. M. Leonard, and D. D. Ho. HIV-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time. *Science*, 271(5255):1582–1586, 1996.

[240] A. S. Perelson, P. Essunger, Y. Cao, M. Vesanen, A. Hurley, K. Saksela, M. Markowitz, and D. D. Ho. Decay characteristics of HIV-1-infected compartments during combination therapy. *Nature*, 387(6629):188–191, 1997. doi: 10.1038/387188a0.

[241] I. Perez-Valero and J. R. Arribas. Protease inhibitor monotherapy. *Curr. Opin. Infect. Dis.*, 24:7–11, 2011.

[242] A. Perisic and C. T. Bauch. Social contact networks and disease eradicability under voluntary vaccination. *PLoS Comp. Biol.*, 5:e1000280, 2008.

[243] D. Persaud, P. E. Palumbo, C. Ziemniak, M. D. Hughes, C. G. Alvero, K. Luzuriaga, R. Yogev, E. V. Capparelli, and E. G. Chadwick. Dynamics of the resting CD4(+) T-cell latent HIV reservoir in infants initiating HAART less than 6 months of age. *AIDS*, 26(12):1483–1490, 2012.

[244] M. Pigliucci and P. Box. Is evolvability evolvable? *Nat. Rev. Genet.*, 9: 75–82, 2008.

[245] T. W. Pike, J. R. Kendal, L. E. Rendell, and K. N. Laland. Learning by proportional observation in a species of fish. *Behav. Ecol.*, 21(3):570–575, 2010.

[246] James G. D. Prendergast, Harry Campbell, Nick Gilbert, Malcolm G. Dunlop, Wendy A. Bickmore, and Colin A. M. Semple. Chromatin structure and evolution in the human genome. *BMC Evol. Biol.*, 7:72, 2007.

[247] F. Pulido, J. Arribas, A. Hill, and C. Moecklinghoff. No evidence for evolution of genotypic resistance after three years of treatment with darunavir/ritonavir, with or without nucleoside analogues. *AIDS Res. Hum. Retroviruses*, 2012. doi: 10.1089/AID.2011.0256.

[248] H. Putter, S. H. Heisterkamp, J. M. A. Lange, and F. de Wolf. A Bayesian approach to parameter estimation in HIV dynamical models. *Statist. Med.*, 21(15):2199–2214, 2002.

[249] M. Pérez, A. G. de Vinuesa, G. Sanchez-Duffhues, N. Marquez, M. L. Bellido, M. A. Muñoz-Fernandez, S. Moreno, T. P. Castor, M. A. Calzado, and E. Muñoz. Bryostatin-1 synergizes with histone deacetylase inhibitors to reactivate HIV-1 from latency. *Curr. HIV Res.*, 8(6):418–429, 2010.

[250] M. Radman, I. Matic, and F. Taddei. Evolution of Evolvability. *Ann. N. Y. Acad.*, 870(1):146–155, 1999.

[251] P. B. Rainey and K. Rainey. Evolution of cooperation and conflict in experimental bacterial populations. *Nature*, 425(6953):72–74, 2003.

[252] D. A. Rand, H. B. Wilson, and J. M. McGlade. Dynamics and evolution: Evolutionarily stable attractors, invasion exponents and phenotype dynamics. *Philos. Trans. R. Soc. B*, 343(1305):261–283, 1994. doi: 10.1098/rstb.1994.0025.

[253] T. C. Reluga, C. T. Bauch, and A. P. Galvani. Evolving public perceptions and stability in vaccine uptake. *Math. Biosci.*, 204:185–198, 2006.

[254] R. M. Ribeiro. In vivo dynamics of t cell activation, proliferation, and death in HIV-1 infection: Why are CD4+ but not CD8+ t cells depleted? *Proc. Natl. Acad. Sci. USA*, 99(24):15572–15577, 2002. doi: 10.1073/pnas.242358099.

[255] R. M. Ribeiro, S. Bonhoeffer, and M. A. Nowak. The frequency of resistant mutant virus before antiviral therapy. *AIDS*, 12(5):461, 1998.

[256] R. M. Ribeiro, L. Qin, L. L. Chavez, D. Li, S. G. Self, and A. S. Perelson. Estimation of the initial viral growth rate and basic reproductive number during acute HIV-1 infection. *J. Virol.*, 84(12):6096 –6102, 2010. doi: 10.1128/JVI.00127-10.

[257] L. Rong and A. S. Perelson. Modeling latently infected cell activation: Viral and latent reservoir persistence, and viral blips in HIV-infected patients on potent therapy. *PLoS Comput Biol*, 5(10):e1000533, 2009. doi: 10.1371/journal.pcbi.1000533.

[258] L. Rong and A. S Perelson. Modeling HIV persistence, the latent reservoir, and viral blips. *J. Theor. Biol.*, 2009. doi: 10.1016/j.jtbi.2009.06.011.

[259] L. Rong, Z. Feng, and A. S. Perelson. Emergence of HIV-1 drug resistance during antiretroviral treatment. *Bull. Math. Biol.*, 69(6):2027–2060, 2007. doi: 10.1007/s11538-007-9203-3.

[260] W. Rosche, P. Foster, and J. Cairns. The role of transient hypermutators in adaptive mutation in Escherichia coli. *Proc. Natl. Acad. Sci. USA*, 96: 6862–6867, 1999.

[261] D. I. S. Rosenbloom, A. L. Hill, S. A. Rabi, R. F. Siliciano, and M. A. Nowak. Antiretroviral dynamics determines HIV evolution and predicts therapy outcome. *Nat. Med.*, 18(9):1378–1385, 2012. doi: 10.1038/nm.2892.

[262] D. Ruelle. *Chaotic Evolution and Strange Attractors*. Cambridge University Press, Cambridge, UK, 1989.

[263] L. Ruiz, J. Martinez-Picado, J. Romeu, R. Paredes, M. K Zayat, S. Marfil, E. Negredo, G. Sirera, C. Tural, and B. Clotet. Structured treatment interruption in chronically HIV-1 infected patients after long-term viral suppression. *AIDS*, 14(4):397, 2000.

[264] K. Ruxrungtham, R. J. Pedro, G. H. Latiff, F. Conradie, P. Domingo, S. Lupo, W. Pumpradit, J. H. Vingerhoets, M. Peeters, I. Peeters, T. N. Kakuda, G. De Smedt, and B. Woodfall. Impact of reverse transcriptase resistance on the efficacy of TMC125 (etravirine) with two nucleoside reverse transcriptase inhibitors in protease inhibitor-naïve, nonnucleoside reverse transcriptase inhibitor-experienced patients: study TMC125-C227. *HIV Med.*, 9(10):883–896, 2008. doi: 10.1111/j.1468-1293.2008.00644.x.

[265] M. E. S. Sampah, L. Shen, B. L. Jilek, and R. F. Siliciano. Dose-response curve slope is a missing dimension in the analysis of HIV-1 drug resistance. *Proc. Natl Acad. Sci. USA*, 108(18):7613–7618, 2011. doi: 10.1073/pnas.1018360108.

[266] L. Samuelson. *Evolutionary games and equilibrium selection*. The MIT Press, 1998.

[267] K. H. Schlag. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J. Econ. Theory*, 78:130–156, 1998.

[268] W. Schnabl, P. F. Stadler, C. Forst, and P. Schuster. Full characterization of a strange attractor. *Physica D*, 48(1):65–90, 1991.

[269] G. Schnell, R. W. Price, R. Swanstrom, and S. Spudich. Compartmentalization and clonal amplification of HIV-1 variants in the cerebrospinal fluid during primary infection. *J. Virol.*, 84:2395–2407, 2010.

[270] P. Schuster and K. Sigmund. Replicator dynamics. *J. Theor. Biol.*, 100 (533):8, 1983.

[271] P. Schuster and K. Sigmund. Replicator dynamics. *J. Theor. Biol.*, 100: 533–538, 1983.

[272] A. Sedaghat, R. F. Siliciano, and C. O. Wilke. Low-level HIV-1 replication and the dynamics of the resting CD4+ t cell reservoir for HIV-1 in the setting of HAART. *BMC Inf. Dis.*, 8(1):2, 2008. doi: 10.1186/1471-2334-8-2.

[273] A. R Sedaghat, J. B Dinoso, L. Shen, C. O Wilke, and R. F Siliciano. Decay dynamics of HIV-1 depend on the inhibited stages of the viral life cycle. *Proc. Natl. Acad. Sci. USA*, 105(12):4832, 2008.

[274] J. Seger and J. Antonovics. Dynamics of some simple host-parasite models with more than two genotypes in each species [and discussion]. *Philos. Trans. R. Soc. B*, 319(1196):541–555, 1988.

[275] L. Shan, K. Deng, N. S. Shroff, C. M. Durand, S. A. Rabi, H.-C. Yang, H. Zhang, J. B. Margolick, J. N. Blankson, and R. F. Siliciano. Stimulation of HIV-1-Specific cytolytic t lymphocytes facilitates elimination of latent viral reservoir after virus reactivation. *Immunity*, 36(3):491–501, 2012. doi: 10.1016/j.immuni.2012.01.014.

[276] L. Shan et al. *In prep.*, 2013.

[277] R. Shankarappa, J. B. Margolick, S. J. Gange, A. G. Rodrigo, D. Upchurch, H. Farzadegan, P. Gupta, C. R. Rinaldo, G. H. Learn, and X. He. Consistent viral evolutionary changes associated with the progression of human immunodeficiency virus type 1 infection. *J. Virol.*, 73(12): 10489–10502, 1999.

[278] A. C. Shaver and P. D. Sniegowski. Spontaneously arising *mutL* mutators in evolving *Escherichia coli* populations are the result of changes in repeat length. *J. Bacteriol.*, 185(20):6076–6079, 2003.

[279] L. Shen, S. Peterson, A. R. Sedaghat, M. A. McMahon, M. Callender, H. Zhang, Y. Zhou, E. Pitt, K. S. Anderson, E. P. Acosta, et al. Dose-response curve slope sets class-specific limits on inhibitory potential of anti-HIV drugs. *Nat. Med.*, 14(7):762–766, 2008.

[280] L. Shen, S. A. Rabi, A. R. Sedaghat, L. Shan, J. Lai, S. Xing, and R. F. Siliciano. A critical subset model provides a conceptual basis for the high antiviral activity of major HIV drugs. *Sci. Transl. Med.*, 3(91):91ra63, 2011. doi: 10.1126/scitranslmed.3002304.

[281] A. Sigal, J. T. Kim, A. B. Balazs, E. Dekel, A. Mayo, R. Mili, and D. Baltimore. Cell-to-cell spread of HIV permits ongoing replication despite antiretroviral therapy. *Nature*, 477:95–98, 2011. doi: 10.1038/nature10347.

[282] K. Sigmund. Time averages for unpredictable orbits of deterministic systems. *Ann. Oper. Res.*, 37:217–228, 1992.

[283] J. D. Siliciano and R. F. Siliciano. Enhanced culture assay for detection and quantitation of latently infected, resting CD4+ t-cells carrying replication-competent virus in HIV-1-infected individuals. *Methods Mol. Biol.*, 304:3–15, 2005. doi: 10.1385/1-59259-907-9:003.

[284] J. D. Siliciano, J. Kajdas, D. Finzi, T. C. Quinn, K. Chadwick, J. B. Margolick, C. Kovacs, S. J. Gange, and R. F. Siliciano. Long-term follow-up studies confirm the stability of the latent reservoir for HIV-1 in resting CD4+ t cells. *Nat. Med.*, 9(6):727–728, 2003. doi: 10.1038/nm880.

[285] J. D. Siliciano, J. Lai, M. Callender, E. Pitt, H. Zhang, J. B. Margolick, J. E. Gallant, J. Cofrancesco, R. D. Moore, S. J. Gange, and R. F. Siliciano. Stability of the latent reservoir for HIV-1 in patients receiving valproic acid. *J. Inf. Dis.*, 195(6):833–836, 2007. doi: 10.1086/511823.

[286] J. C. Silva and A. S. Kondrashov. Patterns in spontaneous mutation revealed by human-baboon sequence comparison. *Trends Genet.*, 18: 544–547, 2002.

[287] B. B. Simen, J. F. Simons, K. H. Hullsiek, R. M. Novak, R. D. Macarthur, J. D. Baxter, C. Huang, C. Lubeski, G. S. Turenchalk, M. S. Braverman, B. Desany, J. M. Rothberg, M. Egholm, and M. J. Kozal. Low-abundance drug-resistant viral variants in chronically HIV-infected, antiretroviral treatment-naive patients significantly impact treatment outcomes. *J. Infect. Dis.*, 199:693–701, 2009.

[288] B. Sinervo and R. Calsbeek. The developmental, physiological, neural, and genetical causes and consequences of frequency-dependent selection in the wild. *Annu. Rev. Ecol. Evol. Syst.*, 37:581–610, 2006.

[289] B Sinervo and C. M. Lively. The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380:240–243, 1996.

[290] B. Sinervo, D. B. Miles, W. A. Frankino, M. Klukowski, and D. F. DeNardo. Testosterone, endurance, and Darwinian fitness: natural and sexual selection on the physiological bases of alternative male behaviors in side-blotched lizards. *Horm. Behav.*, 38:222–233, 1996.

[291] B. Sinervo, A. Chaine, J. Clober, R. Calsbeek, L. Hazard, L. Lancaster, A. G. McAdam, S. Alonzo, G. Corrigan, and M. E. Hochberg.

Self-recognition, color signals, and cycles of greenbeard mutualism and altruism. *Proc. Natl. Acad. Sci. USA*, 103:7372–7377, 2006.

[292] B. Sinervo, B. Heulin, Y. Surget-Groba, J. Clobert, D. B. Miles, A. Corl, A. Chaine, and A. Davis. Models of density-dependent genic selection and a new rock-paper-scissors social system. *Am. Nat.*, 170(5):663–680, 2007.

[293] R. J. Smith. Adherence to antiretroviral HIV drugs: how many doses can you miss before resistance emerges? *Proc. R. Soc. B*, 273(1586):617–624, 2006. doi: 10.1098/rspb.2005.3352.

[294] P. D. Sniegowski, P. J. Gerrish, and R. E. Lenski. Evolution of high mutation rates in experimental populations of *Escherichia coli*. *Nature*, 387 (6634):703–705, 1997.

[295] P. D. Sniegowski, P. J. Gerrish, T. Johnson, and A. Shaver. The evolution of mutation rates: separating causes from consequences. *Bioessays*, 22(12): 1057–1066, 2000.

[296] M. C. Strain, S. J. Little, E. S. Daar, D. V Havlir, H. F Günthard, R. Y. Lam, O. A. Daly, J. Nguyen, C. C. Ignacio, C. A. Spina, D. D. Richman, and J. K. Wong. Effect of treatment, during primary infection, on establishment and clearance of cellular reservoirs of HIV-1. *J. Inf. Dis.*, 191(9): 1410–1418, 2005.

[297] G. Szabó and C. Hauert. Phase transitions and volunteering in spatial public goods games. *Phys. Rev. Lett.*, 89:118101, 2002.

[298] G. Szabó and C. Tőke. Evolutionary prisoner's dilemma game on a square lattice. *Phys. Rev. E*, 58:69–73, 1998.

[299] F. Taddei, M. Radman, J. Maynard Smith, B. Toupance, P. H. Gouyon, and B. Godelle. Role of mutator alleles in adaptive evolution. *Nature*, 387 (6634):700–702, 1997.

[300] B. Taiwo, L. Zheng, S. Gallien, R. M. Matining, D. R. Kuritzkes, C. C. Wilson, B. I. Berzins, E. P. Acosta, B. Bastow, P. S. Kim, and J. J. Eron, Jr. Efficacy of a nucleoside-sparing regimen of darunavir/ritonavir plus raltegravir in treatment-naive HIV-1-infected patients (ACTG a5262). *AIDS*, 25(17):2113–2122, 2011. doi: 10.1097/QAD.0b013e32834bbaa9.

[301] F. Takens. On the numerical determination of the dimension of an attractor. In Boele Braaksma, Hendrik Broer, and Floris Takens, editors, *Dynamical Systems and Bifurcations*, volume 1125 of *Lecture Notes in Mathematics*, pages 99–106. Springer Berlin / Heidelberg, Germany, 1985. doi: 10.1007/BFb0075637.

[302] M. M. Tanaka, C. T. Bergstrom, and B. R. Levin. The evolution of mutator genes in bacterial populations: the roles of environmental change and timing. *Genetics*, 164:843–854, 2003.

[303] C. E. Tarnita, T. Antal, H. Ohtsuki, and M. A. Nowak. Evolutionary dynamics in set structured populations. *Proc. Natl. Acad. Sci. USA*, 106: 8601–8604, 2009.

[304] C. E. Tarnita, H. Ohtsuki, T. Antal, F. Fu, and M. A. Nowak. Strategy selection in structured populations. *J. Theor. Biol.*, 259:570–581, 2009.

[305] P. D. Taylor and L. B. Jonker. Evolutionary stable strategies and game dynamics. *Math. Biosci.*, 40(1-2):145–156, 1978.

[306] P. D. Taylor and L. B. Jonker. Evolutionarily stable strategies and game dynamics. *Math. Biosci.*, 40:145–156, 1978.

[307] O. Tenaillon, B. Toupance, H. Le Nagard, F. Taddei, and B. Godelle. Mutators, population size, adaptive landscape and the adaptation of asexual populations of bacteria. *Genetics*, 152(2):485–493, 1999.

[308] O. Tenaillon, H. Le Nagard, B. Godelle, and F. Taddei. Mutators and sex in bacteria: conflict between adaptive strategies. *Proc. Natl. Acad. Sci. USA*, 97(19):10465–10470, 2000.

[309] D. Tian, Q. Wang, P. Zhang, H. Araki, S. Yang, M. Kreitman, T. Nagylaki, R. Hudson, J. Bergelson, and J. Q. Chen. Single-nucleotide mutation rate increases close to insertions/deletions in eukaryotes. *Nature*, 455: 105–108, 2008.

[310] A. Traulsen, J. C. Claussen, and C. Hauert. Coevolutionary dynamics: from finite to infinite populations. *Phys. Rev. Lett.*, 95:238701, 2005.

[311] A. Traulsen, J. M. Pacheco, and M. A. Nowak. Pairwise comparison and selection temperature in evolutionary game dynamics. *J. Theor. Biol.*, 246: 522–529, 2007.

[312] A. Traulsen, D. Semmann, R.D. Sommerfeld, H.-J. Krambeck, and M. Milinski. Human strategy updating in evolutionary games. *Proc. Natl. Acad. Sci. USA*, 107:2962–2966, 2010.

[313] J M J Travis and E R Travis. Mutator dynamics in fluctuating environments. *Proc. R. Soc. B*, 269(1491):591–597, 2002. doi: 10.1098/rspb.2001.1902.

[314] A. Tversky and D. Kahneman. Availability: A heuristic for judging frequency and probability. *Cogn. Psych.*, 5:207–232, 1973.

[315] G. F. Vanhove, J. M. Schapiro, M. A. Winters, T. C. Merigan, and T. F. Blaschke. Patient compliance and drug failure in protease inhibitor monotherapy. *J. Am. Med. Ass*, 276(24):1955–1956, 1996. doi: 10.1001/jama.1996.03540240033024.

[316] R. Vardavas, R. Breban, and S. Blower. Can influenza epidemics be prevented by voluntary vaccination? *PLoS Comp. Biol.*, 3:e85, 2007.

[317] A. Wagner. Robustness and evolvability: a paradox resolved. *Proc. R. Soc. B*, 275(1630):91, 2008.

[318] L. M. Wahl and M. A. Nowak. Adherence and drug resistance: predictions for therapy outcome. *Proc. R. Soc. B*, 267(1445):835 –843, 2000. doi: 10.1098/rspb.2000.1079.

[319] M. Weber. Evolutionary plasticity in prokaryotes: a Panglossian view. *Biol. Philos.*, 11:67–88, 1996.

[320] J. W. Weibull. *Evolutionary game theory*. The MIT press, 1997.

[321] S. Wielgoss, J. E. Barrick, O. Tenaillon, M. J. Wiser, W. J. Dittmar, S. Cruveiller, B. Chane-Woon-Ming, C. Médigue, R. E. Lenski, and D. Schneider. Mutation rate dynamics in a bacterial population reflect tension between adaptation and genetic load. *Proc. Natl. Acad. Sci. USA*, 110(1):222–227, 2013.

[322] W. Winckler, S. R. Myers, D. J. Richter, R. C. Onofrio, G. J. McDonald, R. E. Bontrop, G. A. T. McVean, S. B. Gabriel, D. Reich, P. Donnelly, and D. Altshuler. Comparison of fine-scale recombination rates in humans and chimpanzees. *Science*, 308:107–111, 2005.

[323] J. K. Wong, M. Hezareh, H. F. Günthard, D. V. Havlir, C. C. Ignacio, C. A. Spina, and D. D. Richman. Recovery of replication-competent HIV despite prolonged suppression of plasma viremia. *Science*, 278(5341): 1291–1295, 1997. doi: 10.1126/science.278.5341.1291.

[324] R. J. Woods, J. E. Barrick, R. F. Cooper, U. Shrestha, M. R. Kauth, and R. E. Lenski. Second-order selection for evolvability in a large escherichia coli population. *Science*, 331(6023):1433–1436, 2011. doi: 10.1126/science.1198914.

[325] H. Wu, Y. Huang, E. P. Acosta, S. L. Rosenkranz, D. R. Kuritzkes, J. J. Eron, A. S. Perelson, and J. G. Gerber. Modeling long-term HIV dynamics and antiretroviral response: effects of drug potency, pharmacokinetics, adherence, and drug resistance. *J. Acquir. Immune Defic. Syndr.*, 39(3): 272–283, 2005.

[326] C. S. Wylie, C. Ghim, D. Kessler, and H. Levine. The fixation probability of rare mutators in finite asexual populations. *Genetics*, 181:1595–1612, 2009.

[327] S. Xing, C. K. Bullen, N. S. Shroff, L. Shan, H.-C. Yang, J. L. Manucci, S. Bhat, H. Zhang, J. B. Margolick, T. C. Quinn, D. M. Margolis, J. D. Siliciano, and R. F. Siliciano. Disulfiram reactivates latent HIV-1 in a bcl-2-transduced primary CD4+ t cell model without inducing global t cell activation. *J. Virol.*, 85(12):6060–6064, 2011. doi: 10.1128/JVI.02033-10.

[328] S. Xing, S. Bhat, N. S. Shroff, H. Zhang, J. A. Lopez, J. B. Margolick, J. O. Liu, and R. F. Siliciano. Novel structurally related compounds reactivate latent HIV-1 in a bcl-2-transduced primary CD4+ t cell model without inducing global t cell activation. *J. Antimicrob. Chemother.*, 67(2):398–403, 2012. doi: 10.1093/jac/dkr496.

[329] X. Zhao and K. Drlica. Restricting the selection of Antibiotic-Resistant mutant bacteria: Measurement and potential use of the mutant selection window. *J. Infect. Dis.*, 185(4):561–565, 2002. doi: 10.1086/338571.

[330] Z. Zhao and C. Jiang. Methylation-dependent transition rates are dependent on local sequence lengths and genomic regions. *Mol. Biol. Evol.*, 24:23–25, 2007.

Colophon

THIS THESIS WAS TYPESET
using LaTeX, originally developed by
Leslie Lamport and based on Donald
Knuth's TeX. The body text is set in 11 point
Arno Pro, designed by Robert Slimbach in
the style of book types from the Aldine Press
in Venice, and issued by Adobe in 2007. A
template, which can be used to format a PhD
thesis with this look and feel, has been
released under the permissive MIT (X11)
license, and can be found online at
github.com/suchow/ or from the author at
suchow@post.harvard.edu.