Synchronization of Viral Lifecycle Length to Antiviral Drug Dosage Schedules and the Emergence of "Cryptic Resistance"

Synchronization of viral lifecycle length to antiviral drug dosage schedules and the emergence of "cryptic resistance"

Mark C. Freeman

Applied Mathematics Senior Thesis
Advisor: Professor Martin Nowak
Harvard University, Cambridge, MA

April 1, 2015

Abstract

Viral infections, such as HIV, are often treated with orally administered antiviral medications that are dosed at particular intervals, leading to periodic drug levels and hence periodic inhibition of viral replication. These drugs generally bind to viral proteins and inhibit particular steps in the viral lifecycle, and resistance often evolves due to point mutations in the virus that prevent the drug from binding its target. However, it has been proposed (Wahl & Nowak, Proc Roy Soc B, 2000) that a completely different "cryptic" mechanism for resistance could exist: the virus population may evolve towards synchronizing its lifecycle with the pattern of drug treatment. If the lifecycle of the virus is a multiple of the dosing interval, it is possible that over time the bulk of the virus population will replicate during trough concentrations of the drug. In this thesis, we use stochastic mathematical models of viral dynamics to demonstrate that cryptic resistance could plausibly provide a powerful fitness advantage to a wide variety of viral strains whose expected lifecycle times are slightly less than the expected time between doses of an antiviral drug, allowing them to survive drug regimes that would otherwise drive infected cell populations to extinction. This in turn suggests that continuously-administered antiviral drug treatments may be significantly more effective than periodically-administered treatments in combatting viral infections.

Contents

1 Introduction

Viral infections are a major cause of human morbidity and mortality. While vaccines to prevent viral illnesses have existed for over a century, it is only in the past several decades that drug directly targeting viral replication have been developed. Antiviral drugs now exist for pathogens including the human immunodeficiency virus (HIV), hepatitis B and C viruses, influenza A and B viruses, herpes simplex viruses, cytomegalovirus, Epstein Barr virus, and varicella zoster virus (chickenpox virus). These drugs each target a specific phase of the virus's lifecycle, and by binding to a viral protein or otherwise interfering with a critical step in viral replication are able to reduce the virus's growth rate. Examples of viral functions targeted by antiviral drugs include binding of the viral particle to the target cell membrane, transcription of the viral genome, integration of the virus in the host cell genome, or post-translational cleavage of viral proteins.

Antiviral treatments that are initially successful at reducing viral loads may eventually be rendered ineffective, in individual patients or in entire populations, by the emergence of drug-resistant strains. Drug resistance occurs when a viral strain gains a mutation that allows it to replicate efficiently despite the presence of the drug, and this strain subsequently outcompetes the wildtype strain to fix in the viral population. Resistance is generally accepted to be conferred by mutations that interfere with the ability of the drug molecule to inhibit the intended viral target. For example, for antivirals that block the fusion of viral particles with the target cell membrane, the mutations which confer resistance can alter the shapes and chemical properties of viral proteins to either prevent drug binding or allow the virus to enter the cell despite the presence of the drug.

There may, however, be another mechanism by which resistance can develop in viral populations. In a 2000 paper, Wahl and Nowak (1) hypothesized that a heretofore unobserved effect, which they termed "cryptic resistance" may prevent a viral population from being suppressed by an antiviral drug, without it needing to evolve the ability to alter drug binding or even complete its lifecycle in the drug's presence. This insight was motivated by the realization that drug levels are not constant during the course of viral treatment, and hence the viral fitness in the presence of the drug is also time-varying. Like most medications, antiviral drugs are administered in discrete doses of constant size separated by approximately equal time intervals. Shortly after a dose, the drug is absorbed and drug levels are high, and the relevant stage of the viral lifecycle is maximally inhibited. However between doses, drug levels decay, and may eventually reach low levels where it is not longer suppressive. This pattern repeats in a periodic pattern over the course of treatment. These authors hypothesized that the virus could avoid the effects of the drug by always completing

the targeted phase of it's lifecycle at near the time when drug concentrations is at a minimum. They suggested that if the length of the viral lifecycle was a mutable trait, the selection pressure of treatment would push it to evolve to become approximately equal to the duration of time between successive drug doses. The virus population would then become synchronized with the pattern of drug levels. In this manner, the virus could sustain itself indefinitely by "hiding" from the highest concentrations of the drug, even though it has no means of counteracting the effects of drug molecules when they are present.

The goal of this project is use mathematical models to explore whether cryptic resistance could plausibly arise in a viral population subjected to antiviral treatment. We start by augmenting well-established models for viral infection dynamics to account for distinct phases of the viral lifecycle. This model includes a maturation rate, which can be varied to change the viral lifecycle length. Fluctuating drug concentrations are incorporated as a periodic time-dependent infectivity of the virus. We then use this model to understand how the length of the viral lifecycle determines the fitness of a viral strain in the presence of period drug levels. Deterministic differential equation models are used initially, and later replaced with a stochastic formulation. We use two methods to evaluate the potential success of a cryptic resistance strategy. First, we examine viral fitness in a single-strain infection by determining the growth rate and eventual equilibrium level of the virus, and look for the lifecycle length that optimizes these values. Secondly, we simulate the more evolutionarily-relevant scenario where competition between multiple strains occurs within a single infected individual. In this case we look for the strain that outcompetes the others and is able to persist despite drug treatment. In each case, we are interested in whether strains that are synchronized to the drug dose schedule - that is, have lifecycle lengths that are an integer multiple of the time between drug doses - are able to take over the viral population, and thereby confer cryptic resistance.

2 The basic viral dynamics model

The dynamics of viral populations within infected individuals have been studied mathematically for decades, with applications to diseases such as HIV, Hepatitis B, and Hepatitis C (2, 3, 4, 5). The simplest viral dynamics models divide the relevant cells populations in the body into compartments based on their state of infection, and use a system of differential equations to describe the change in these populations over time. The size of the population of healthy cells is described by state variable x, the population of infected cells by y, and a population of free viruses v. Healthy

cells are produced at a rate λ, die with rate constant d_x, and are infected at a rate proportional to the product of v and an infection rate constant b. Infected cells die with rate constant d_y and produce free virus particles at a rate k. Free viruses are cleared from the system at rate c. A schematic diagram of this model is shown in Figure 1.

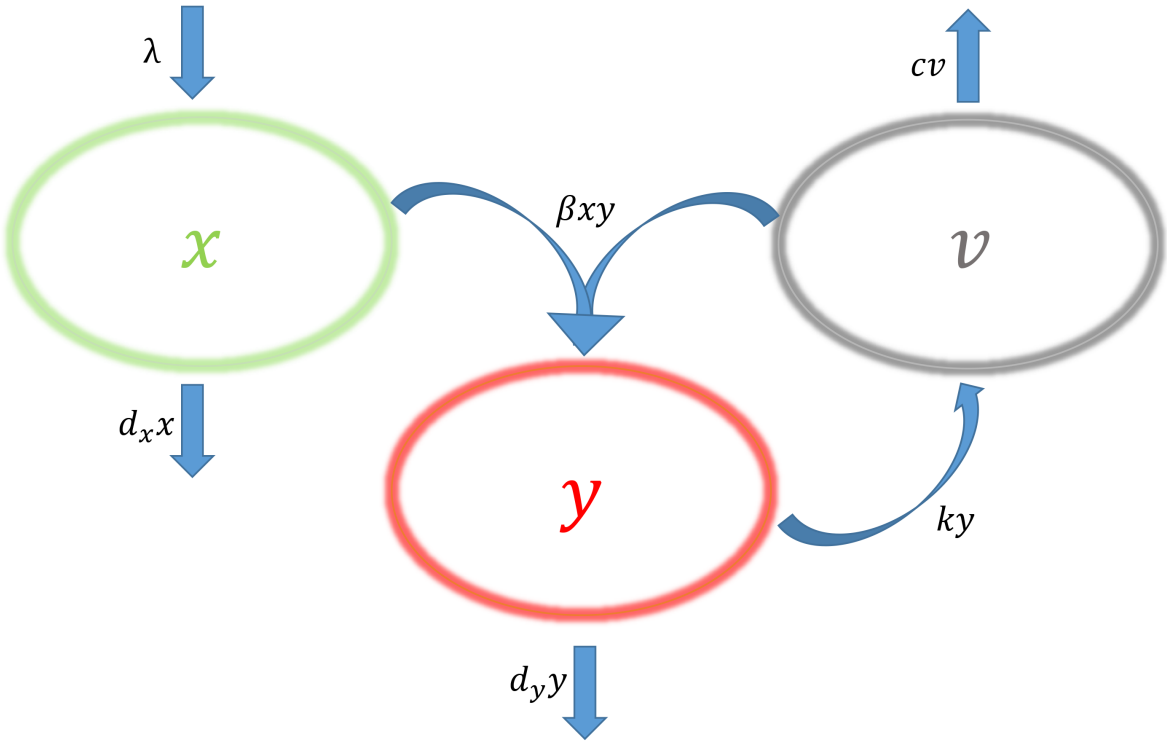

Figure 1: The basic viral dynamics model. A flow diagram conceptualizing the basic viral dynamics model. Ellipses represent populations, and arrows represent changes in population. Uninfected cells (x) are produced at a rate λ and die at a rate d_x. They become infected upon contact with free virus (v) at a rate b. Infected cells (y) produce free virus at a rate k, and die at a rate d_y. Free virions are cleared from the system at a rate c.

The overall model is represented by the following system of equations:

$$\begin{aligned} \frac{dx}{dt} &= \lambda - bxv - d_x x \\ \frac{dy}{dt} &= bxv - d_y y \\ \frac{dv}{dt} &= ky - cv \end{aligned} \tag{2.0.1}$$

This model has two equilibria, which can be determined by setting $\frac{dx}{dt} = \frac{dw}{dt} = \frac{dy}{dt} = 0$. At the

"uninfected" equilibrium, the infection is not able to take hold:

$$\begin{aligned} x^* &= \frac{\lambda}{d_x} \\ y^* &= 0 \\ v^* &= 0 \end{aligned} \tag{2.0.2}$$

At the "infected" equilibrium, the infection persists:

$$\begin{aligned} x^* &= \frac{cd_y}{bk} \\ y^* &= \frac{\lambda}{c} - \frac{d_x d_y}{bk} \\ v^* &= \frac{d_y}{c} \end{aligned} \tag{2.0.3}$$

Standard stability analysis shows that there is always one and only one stable equilibria for this system, and the choice of equilibria is determined by the value of a parameter composition called the basic reproductive ratio, $R_0 = \frac{b\lambda k}{d_x d_y c}$. This quantity describes the average number of new cells that would be infected by a single infected cell introduced into a population of otherwise susceptible cells. This can be understood by taking the product of all the steps in this process: the average lifetime of an infected cell ($1/d_y$), the rate of virion production k, the average lifetime of those virions ($1/c$), the rate at which virions infect new cells, per time and target cell level (b), and the target cell level before infection (λ/d_x). If $R_0 > 1$, the population of infected cells will grow to reach the infected equilibrium, Eq. (2.0.3), while if this $R_0 < 1$, the population of infected cells will immediately decrease towards zero, Eq. (2.0.2).

In addition to determining stability, larger R_0 entails larger infected cell and free virus populations and smaller healthy cell populations at equilibrium. This can be seen by rewriting Eq. (2.0.3) as

$$\begin{aligned} x^* = x &= \frac{\lambda}{d_x}\left(\frac{1}{R_0}\right) \\ y^* &= \frac{\lambda}{d_y}\left(1 - \frac{1}{R_0}\right) \\ v^* &= \frac{\lambda}{c}\left(1 - \frac{1}{R_0}\right). \end{aligned} \tag{2.0.4}$$

The basic viral dynamics model can be further simplified by using the fact that for the vast majority of infections, the parameters governing the dynamics of free virions tend to be much higher than those governing the dynamics of cells. Virus is produced in large quantities and is rapidly cleared *in vivo*. This implies the viral population tends to reach a "quasi-steady state" level with respect to the level of infected cells, and suggests that a separation of timescales can be be applied to the system (2, 6). By setting $\frac{dv}{dt} = 0$ we get the relation $v = \frac{k}{c}y$, which then allows us to replace v in the system of Eq. (2.0.1), and obtain the reduced system

$$\begin{aligned} \frac{dx}{dt} &= \lambda - bxv - d_x x \\ \frac{dy}{dt} &= bxv - d_y y \end{aligned} \tag{2.0.5}$$

where b in this model equals bk/d_y in the previous one. One could repeat the stability analysis for this system, and would find after incorporating the parameter substation for b, the equilibria are equivalent to Equations (2.0.2) and (2.0.2), and R_0 again determines the stable equilibrium. Throughout the rest of the sections of the paper analyzing deterministic models, we will therefore simply model uninfected and infected cells, and assume viral loads are proportional to levels of virus-producing cells.

This basic model can be used to understand viral evolution. In the simplest analysis, one can simply consider two viral strains and track cells infected with either by y_1 and y_2, each with unique parameters (b_1, d_{y1}) and (b_2, d_{y2}) leading to a three-equation system. Stability analysis of this augmented system (2) leads to the insight that there is competitive exclusion between the strains, and only the one with the highest R_0 value will remain at the infected equilibrium.

3 Deterministic viral dynamics models of cryptic resistance

3.1 Incorporating maturation time into the basic viral dynamics model

The standard viral dynamics model makes the simplifying assumption that infected cells produce new virus particles as soon as they are infected. In reality, a virus must complete many stages of its lifecycle prior before new virions are created. These stages many include uncoating the viral particle, transcription and translation of the viral genome, copying of the viral genome, assembly of viral proteins, or even cell-cycle dependent events. The exact steps depend on the particular mode

of replication of the virus, of which there are many variations in nature. However, the common effect is that there is a time-lag in virus production, a necessary pre-requisite for the emergence of "cryptic resistance".

We created the simplest possible model that captures the finite length of the viral lifecycle and the fact that any particular drug acts only on a particular phase of the lifecycle. To do this, we subdivided the infected cell population into two subpopulations, "immature" (w) and "mature" infected cells (y) (Figure 2). When healthy cells were initially infected, they were added to the recently infected cell population. These cells transition to a mature state sate with rate constant m, in which they can go on to produce virus and infected other cells at rate b. It is this latter process that we will assume the drug inhibits, as discussed in the next section. Recently infected cells die without maturing with rate constant d_w. Other aspects of the model are the same, and it is described the system of equations

$$\begin{aligned} \frac{dx}{dt} &= \lambda - bxy - d_x x \\ \frac{dw}{dt} &= bxy - mw - d_w w \\ \frac{dy}{dt} &= mw - d_y y \end{aligned} \tag{3.1.1}$$

Similar models to this have been used previously to study multiple stages of the viruses life-cycle (7). This system's equilibria can again be found by setting $\frac{dx}{dt} = \frac{dw}{dt} = \frac{dy}{dt} = 0$. Like the basic viral dynamics model, it has an infected and an uninfected equilibrium. At the "uninfected" equilibrium, the infection is not able to take hold:

$$\begin{aligned} x^* &= \frac{\lambda}{d_x} \\ w^* &= 0 \\ y^* &= 0 \end{aligned} \tag{3.1.2}$$

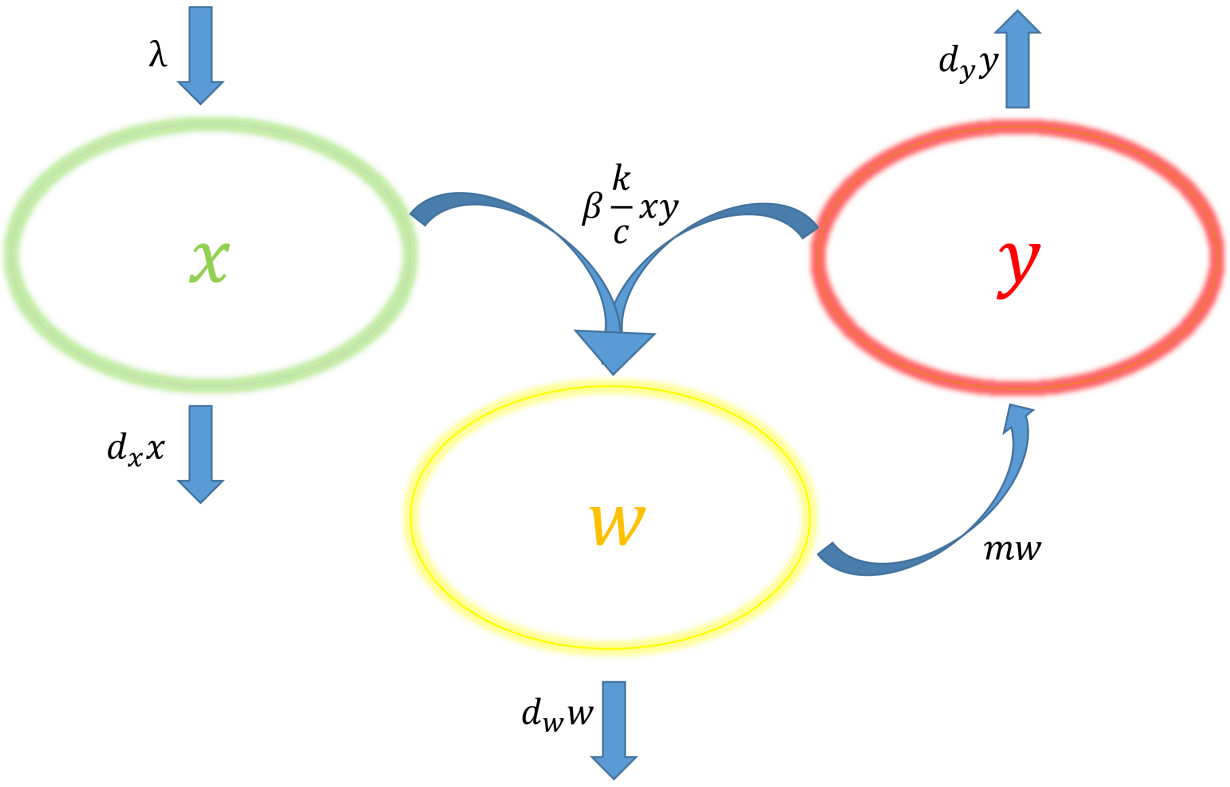

Figure 2: The basic viral dynamics model with infected cell maturation time. A flow diagram for the viral dynamics model that incorporates two stages of the viral lifecycle, with a time lag between them. Ellipses represent populations, and arrows represent changes in population. Uninfected cells (x) are produced at a rate λ and die at a rate d_x. They become infected upon contact with free virus (v) at a rate b, first entering the immature stage (w). Immature cells die at a rate d_w and progress to a mature stage (y) at rate m. Mature infected cells can produce virus and lead to infection of other cells at rate b, or die at a rate d_y.

At the "infected" equilibrium, the infection persists:

$$\begin{aligned} x^* &= \frac{d_y}{b}\left(\frac{m+d_w}{m}\right) = \frac{\lambda}{d_x}\left(\frac{1}{R_0}\right) \\ w^* &= \frac{\lambda}{m+d_w} - \frac{d_x d_y}{bm} = \frac{\lambda}{m+d_w}\left(1-\frac{1}{R_0}\right) \\ y^* &= \frac{\lambda}{d_y} - \frac{d_x}{b}\left(\frac{m+d_w}{m}\right) = \frac{\lambda}{d_y}\left(1-\frac{1}{R_0}\right) \end{aligned} \tag{3.1.3}$$

where the formula for the basic reproductive ratio is now

$$R_0 = \frac{\lambda b m}{d_x d_y (m + d_w)}. \tag{3.1.4}$$

As in the basic model, the stability of each equilibrium hinges on the value of R_0. When the basic reproductive ratio exceeds unity, the infected equilibrium is the only stable equilibrium of the system, and when the ratio is less than one, the uninfected equilibrium is the only stable equilibrium. Incorporating maturation time into the model reduces the rate at which the viral population grows, and R_0 is reduced by a factor equal to the probability that a immature cell survives to mature, $m/(d_w + m)$.

Again analogous the basic viral dynamics model, stability analysis of a multi-strain version of this model would show that the strain with the highest R_0 value would outcompete the other and take over the whole viral population at equilibrium. Since R_0 is monotonically increasing with m, this shows that in agreement with intuition, it is always better for a virus to mature as quickly as possible. Higher maturation rates also result in lower levels of uninfected host cells, lower levels of immature cells, and higher levels of mature cells. If immature cells do not die ($d_w = 0$), then R_0, x^*, and y^* do not depend on m.

In order for there to be any potential benefits to a lag in virus production (low m), the inherent reproductive advantage to a fast viral generation time would need to be outmatched by an advantage to delaying viral assembly time. In the subsequent sections, we will demonstrate how in the presence of a periodically fluctuating drug, reducing the maturation time can indeed provide such an advantage.

3.2 Incorporating periodic drug concentrations into the viral dynamics model

Could fluctuations of antiviral drug concentration favor the synchronization of viral lifecycle to the period of the drug dosing, therefore leading to "cryptic resistance"? This was the hypothesis provided by Wahl and Nowak (1) which we set out to test. To do so, we incorporated periodic antiviral drug concentrations into the viral dynamics model. We assumed that drugs prevent new infections from occurring, by reducing the b term. This heuristically captures many drug actions that may occur intra- or extra-cellularly, such as preventing viral budding, creating defective virions, binding and deactivating free virions, or preventing viral entry.

Antiviral drug concentration was first incorporated into the model in the simplest way: as an

on-off switch. The drug was assumed to be taken at doses separated by an interval T, and to totally inhibit new infections for a time fT thereafter. For the last $(1-f)T$ before each dose, the drug was assumed to be inactive. This simple switch was meant to model the time-dependent drug concentrations, and therefore viral fitness levels, expected from periodic dosing schedules. We also expected that the dichotomy between the fitness of viruses whose lifecycles were synchronous versus those who were asynchronous would be maximized with this drug scheme.

The model becomes

$$\begin{aligned}
\frac{dx}{dt} &= \lambda - bC(t)xy - d_x x \\
\frac{dw}{dt} &= bC(t)xy - mw - d_w w \\
\frac{dy}{dt} &= mw - d_y y \\
C(t) &= \begin{cases} 0 & t \bmod T \leq f \\ 1 & \text{otherwise} \end{cases}
\end{aligned} \tag{3.2.1}$$

Standard stability analysis reveals that the only true equilibrium of the system is the uninfected equilibrium of the model without drug dynamics, Eq.(3.1.2). Due to periodic fluctuations in C, the system cannot settle on a single equilibrium with infected cells present. However, we suspect that due to the periodic nature of the differential equations, a "stable periodic equilibrium" might exist, with nonzero infected cell populations and the property that $(x(t), w(t), y(t)) = (x(t+T), w(t+T), y(t+T))$.

As long as we are not in a parameter regime where the dynamics of Eq. (3.2.1) are vastly faster than the period of drug switching (which we expect to be true for most infections), we can approximate R_0 for this system to be the weighted average the value without drug, Eq. (3.1.4), and what it would be if the drug always suppressed all infection, i.e. $R_0 = 0$. This leads to $R_0 = (1-f)\frac{\lambda bm}{d_x d_y (m+d_w)}$. However, we do not expect for this system that R_0 alone determines the ability of the infection to persist. This R_0 calculation assumes that cells are equally like to be at any phase of their lifecycle at any point of time in the drug cycle, where in reality, synchronization with the dosing pattern may occur.

Due to the difficulty of analyzing systems with time-dependent coefficients analytically, we instead chose to examine these dynamics numerically. Our goal was to determine the dynamics

Table 1: Parameter values used in numerically solving the viral dynamics model with maturation time and periodic drug concentration dynamics. The units of cell concentration ("cell conc.") are arbitrary but could be a common medical unit such as cells per microliter.

Parameter	Description	Value	Units
λ	Input rate of uninfected cells	10	cell conc. $\cdot$ days^{-1}
d_x	Death rate of uninfected cells	10^{-2}	days^{-1}
b	Infection rate	0.01	(cell conc.)$^{-1} \cdot$ days^{-1}
d_w	Death rate of immature cells	0	days^{-1}
d_y	Death rate of mature cells	1	days^{-1}
T	Time between drug doses	2	days
f	Fraction of the time that drug is active	0.85	N/A

of viral populations under this treatment regime, as a function of the length of the viral lifecycle. Parameters for this analysis are given in Table 1 and were chosen according the following logic. In the absence of the drug, we required $R_0 > 1$ and chose $R_0 = 10$. We required that the virus was pathogenic, such that $d_y < d_x$. Within these constraints, we chose values of d_x, d_y, and λ that are consistent with HIV infection (e.g.(8, 9)). We don't think that our results are particularly sensitive to the exact choices of these parameters within these constraints.

We allowed m to vary such that $\frac{1}{m} \in [0.5, 5.5]$.We chose a drug efficacy f so that $R_0(f)$ was dramatically reduced but still greater than 1, allowing the infection to persist at low levels so that we could compare the infection level cross different lifecycle lengths. We also suspected that reducing the fraction of the time for which the reproductive window was open, $(1 - f)T$, would ensure that high-level growth would require a maturation rate such that mT was approximately an integer; in other words, synchronization. In order to assess the effect of m on viral fitness independent of its direct value on R_0, we set $d_w = 0$, which removes the dependence of $R_0(f)$ on m. This assumption is relaxed in later sections.

For the numerical solution, the drug treatment was assumed to begin after the infection had already been established, so the initial conditions were those characterizing the infected equilibrium of the model without drug treatment, (3.1.2). The equations were integrated from time $t = 0$ until the average size of each cell population over the last 20 days differed by no more than 1% from the average population over the previous 20 days. Results were then reported at averages over the previous 10 days.

We examined three different metrics of viral reproductive success as a function of m:

1. The level of mature infected cells (y) at equilibrium.

2. The time until equilibrium was obtained.

3. The maximum growth rate of mature infected cells ($\frac{dy}{dt}$) over the course of the simulation.

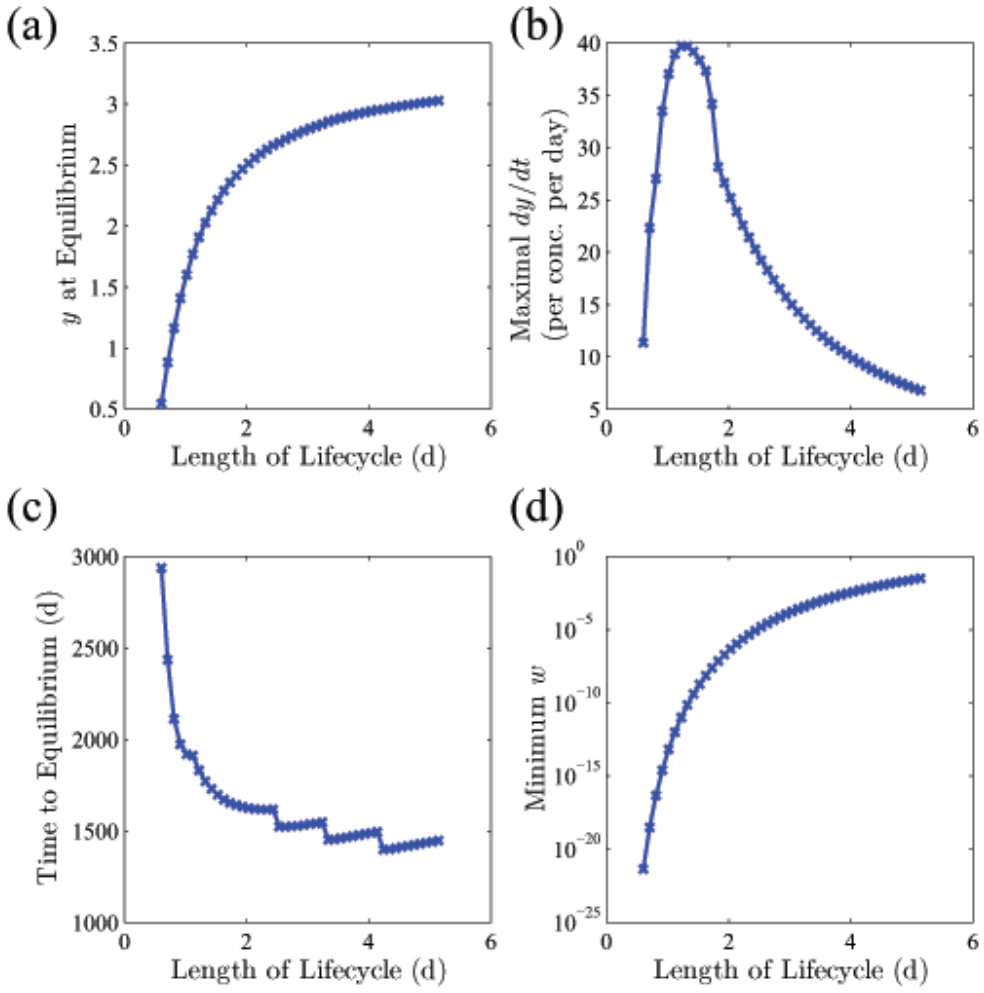

Figure 3: Outcomes of the viral dynamics model in the presence of on-off drug levels. Four outputs from the numerical integration of Eq. (3.2.1) are plotted against the average viral lifecycle length. $\frac{1}{m} + \frac{1}{bx}$, where x is taken at equilibrium. a) Average mature infected cell population (y) at equilibrium. b) Maximal rate of change in the level of infected cells (dy/dt). c) Time required to reach equilibrium population levels . d) Minimum level of immature cells (w) achieved over the course of the simulation. Crosses represent data points and lines denote interpolation between them.

These results are presented in Figure 3a)-c). Instead of plotting results as a function of the average maturation time $1/m$, which only represents one step of the complete lifecycle of the

virus, we incorporated the average time for an infected cell to infect another cell, $\frac{1}{bx}$, were x is the level of uninfected cells. Since we report results at equilibrium, x is the equilibrium value. Therefore the average lifecycle length is given by $\frac{1}{m} + \frac{1}{bx}$.

The infected cell level (which will be proportional the viral load), increased monotonically with the length of the viral lifecycle. Therefore in this model, it is the strain with the slowest maturation time, or the longest lifecycle length, that has the highest fitness by this measure. This was contrary to our expectations that strains that were synchronized to the drug level (i.e. had lifecycle lengths of 2, 4, 6, etc days), would have optimal fitness. To examine this further, we decided to look at other measures of fitness that might gives strains an advantage. If viral strains were to be in a situation where they were directly competing, we hypothesized that equilibrium levels may not be the only important fitness measure, because early growth may be more important for infecting many host cells before another strain. We looked at the two measures of the growth rate of the infection: the maximal rate of increase in mature infected cells over the course of treatment (dy/dt), and the time until the equilibrium of all populations was reached. For the latter, we found that longer viral life cycles were always better because they sped up the time to reach equilibrium and achieved higher population sizes at equilibrium. This was believed to reflect the fact that in this model as in the previous one, the viral strains with long maturation times would tend to have larger immature infected cell populations at equilibrium, and thus could support higher mature infected cell populations in the periods when the drug was inactive.

When looking at the maximal growth rate, we did find that strains with a lifecycle close to the expected optimum did have a clear advantage. However, this was a temporary advantage that tended to be most prominent at the beginning of the time course of the solution, and did not translate to a larger population at equilibrium. It was believed to reflect the result of a tradeoff between the rapid maturation of the immature infected cell population of strains with low maturation times after the drug vanishes from the compartment and the high immature infected cell populations sustained by the strains with longer maturation times.

Re-running the simulation for other values of f ranging from 0.5 to 0.95, we were unable to find any sets of parameters in which viral strains whose lifecycle lengths approximated integer multiples of the time between successive drug doses had higher equilibrium populations or attained those populations more quickly than other viral strains.

Upon more closely examining the output of this model, we discovered an effect that suggested that a differential equation based approach may not be ideal for examining the emergence of cryptic

resistance. When we looked at the minimum value for the level of immature infected cells (w) over the course of the simulation, we noticed that these values were becoming extremely low during times when the drug was on, to the point where in any finite model they would be zero (as low as 10^{-20}). This is expected, as no new infections can occur when the drug is on, since $b \to 0$. In all viral strains, these levels quickly rebounded when the drug was turned off and the few immature cells that survived began maturing at constant rates, infecting new cells and increasing populations of infected cells. This ability for immature cells to mature at a constant rate allows all strains to have infected cells mature when the drug is off and thereby replenish their infected cell populations. However, the reproductive advantage enabling cryptic resistance was believed to be that strains exhibiting cryptic resistance are able to have a much higher proportion of their immature infected cells mature at opportune times than are strains not exhibiting cryptic resistance. Thus, the constant rate of maturation of immature infected cells was hypothesized to nullify the advantage of cryptic resistance, explaining its lack of reproductive success under ths model.

If this hypothesis was correct, then cryptic resistance should be observed if the model conditions are changed so that immature infected cells were discrete and matured a fixed amount of time after they were infected. This would prevent immature infected cells from maturing at the opportune moment as the drug vanished from the compartment unless they had been infected at the appropriate time, which would prevent all strains except those exhibiting cryptic resistance from experiencing significant population rebounds as the drug vanishes from the compartment.

In accordance with this hypothesis, we introduce a fully stochastic, discretized model of viral dynamics in the next section, which allows us to explicitly specify any choice of distribution for the time between events. It also allows for the possibility that viral populations could get so low in the presence of the drug that the infection could go extinct. We also consider direct competition between viral strains with different maturation times that co-infect the same host, which side-steps the need to determine which measure of viral fitness is optimal. Instead, we can determine which strain eventually fixes in a population.

4 Stochastic viral dynamics models of cryptic resistance

4.1 Model formulation

Because of the shortcomings of the ordinary differential equation (ODE) model discussed above, we converted the model to an analogous discretized, probabilistic one in which the expected rate of each biological process was equal to its rate in the ODE model. The reactions that can occur in this model are described in Table 2. For most transitions, we use the standard assumption of a memoryless process, so that the probability density function for the time to the next event is exponentially distributed. However, because of our interest the maturation time lag, we allow this waiting time to be an arbitrary distribution, and use some examples with tunable variance in subsequent sections.

Table 2: The events that can occur in the stochastic version of the two-stage viral dynamics model. X, Y, W represent individual cells and arrows represent reactions that can occur and allow cells to change state. The notation $t \sim P(a)$ means that t is drawn from an unspecified probability distribution with mean a, and $Exp(t)$ refers to the exponential distribution. The time-dependent infectivity $b(t)$ can take a variety of forms.

Event	Description	Waiting time	Average rate
$\varnothing \to X$	A healthy cell enters in the system	$t \sim \text{Exp}(\lambda)$	λ
$X \to \varnothing$	A healthy cell dies	$t \sim \text{Exp}(d_x)$	d_x
$X + Y \to W + Y$	A healthy cell is infected	$t \sim \text{Exp}(b(t))$, varied	varies
$W \to \varnothing$	An immature infected cell dies	$t \sim \text{Exp}(d_w)$	d_w
$W \to Y$	An immature infected cell matures	$t \sim P(m)$, varied	m
$Y \to \varnothing$	A mature infected cell dies	$t \sim \text{Exp}(d_y)$	d_y

Simulations were run both for single viral strains alone (Section 4.2) and for multiple viral strains infecting the same host (Section 4.2). In each case, we again varied the average maturation time between simulations to look for synchronization. Parameter values were the same as reported in Table 1, with the exceptions that λ was scaled by a factor of 10 to increase the size of the compartment and thereby reduce the stochasticity and extinction probability in simulations, and d_x was scaled by a factor of 10 to keep $R_0 \approx (1-f)\frac{\lambda b m}{d_x d_y (m+d_w)}$ the same as for the ODE model. The drug dynamics, which determine the time-dependent infectivity, $b(t)$, were modeled using both on-off switches and more realistic pharmacologic functions. The drug period was kept constant between simulations.

4.2 Single strain dynamics

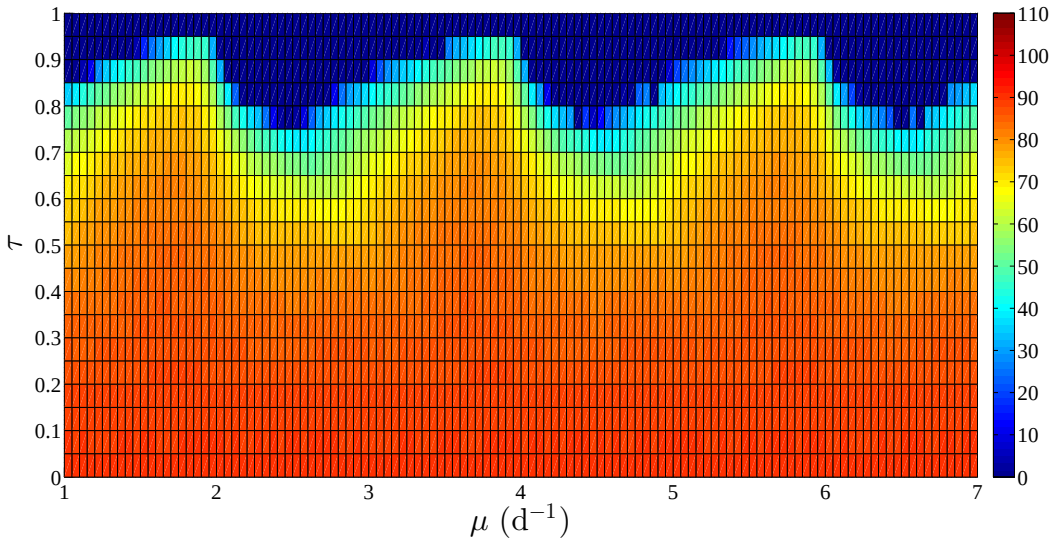

Figure 4: Average equilibrium infection level for the single-strain stochastic model. The color of each point represents the equilibrium number of mature infected cells (y) averaged over 100 simulations, as a function of the maturation time ($1/m$) and the fraction of the time that the drug was on (f). There was no variation in maturation time between cells.

We first investigated the fitness of individual viral strains in the presence of period drug levels, as a function of the maturation time. We chose to model a precisely timed maturation phase, so that each cell spent exactly $\frac{1}{m}$ time in the immature state before maturing, if it did not die before that. We set $d_w = 0$, so that maturation was guaranteed. We used the same on-off drug switch as in Eq. (3.2.1), so that $b(t) \to C(t)b$. We also varied the drug efficacy, measured by the fraction of time the drug is active (f).

The initial cell populations were those of the infected equilibrium of the drugless ODE model. Fitness was measured as the mature infected cell level (y) at equilibrium, averaged over 100 runs of the simulation. Equilibrium was said to be achieved when the average cell population levels over the last 10 days equal to those over the previous 10 days, and this value was returned.

Results of these simulations are presented in Figure 4. We observed three regimes for the behavior of the average equilibrium infected cell level as a function of the drug efficacy f and the maturation time m. In the "low drug effectiveness regime" where f was less than 0.5, the equilibrium infection level was largely independent of the maturation time. In the "high-drug-effectiveness regime" where f exceeded 0.95, the infection was eradicated at the end of every simulation run.

At f values in a "Goldilocks regime" situated between these two extremes, equilibrium infection levels exhibited a strong dependence on the maturation time, peaking at m such that the expected maturation time was 0.25 d less than the time between an integer number of consecutive drug doses, i.e. ($\frac{1}{m} \bmod T \approx 1.75$) for all f values. It is in this regime that synchronization, and hence cryptic resistance, appears to emerge.

We hypothesized that the observed division of the results into three regimes stemmed from the differential success of the synchronization strategy under different levels of effectiveness of the drug. In the **low-drug-effectiveness regime** (low f), the drug was is likely so inefficacious in controlling the infection that it has negligible effect on population dynamics. As a result, the system could settle into an equilibrium almost identical to the infected equilibrium of the drugless, deterministic model, Eq. 3.1.3. Since the equilibrium level of mature infected cells (y^*) in this model is independent of the maturation time when immature cell don't die, this is what we observe in this regime.

In the intermediate **Goldilocks regime** where $f \in (0.5, 0.95)$, the drug is effective enough of the time that the viral population's chances of long-term survival are hampered by the drug's, unless it has a particular lifecycle length. Each generation of immature infected cells must be able to produce mature, virus-producing cells at a time when the drug is inactive. In this way the infection can continue to grow, unhindered by the drug's presence. Since these "windows of reproductive opportunity" are periodic with period T, only maturation times that are equal to integer multiples of T allow for successful replication. Consequently, the equilibrium infection levels of such populations are considerably higher than those of populations that cannot synchronize in this manner. This is the regime in which cryptic resistance is a successful strategy for the virus to pursue.

Interestingly, the maximal infection level in this regime is observed in the case of "imperfect synchronization" where $\frac{1}{m}$ is slightly less than, rather than equal to, an integer multiple of T. The difference between $\frac{1}{m}$ and the next-largest multiple of T is greater than the average contribution to lifecycle length of the other stages in the viral lifecycle (the $1/bx$ contribution discussed §3.2). This in turn suggests that the optimal synchronization between viral lifecycle length and drug dosage interval is indeed imperfect, with the viral lifecycle being slightly shorter than the drug dosage interval.

We hypothesize that this difference reflects the concept of "off-window utilization". The off-window refers to the time when the drug is inactive, and infection is possible. An immature cell that is newly infected in the middle of the off-window will mature in time to output free virus during

the entire subsequent window period if $\frac{1}{m}$ is slightly less than T, but will not be able to make full use of the off-window if $\frac{1}{m}$ is greater than or equal to T. However, the less $\frac{1}{m}$ is, the greater the likelihood that the infected cell matures *before* the off-window opens, increasing the time in which it is exposed to possible death and no reproductive gain. These two forces combined cause the optimal lifecycle length to be less than, but only slightly less than, that needed to allow perfect synchronization. Note that even though the maturation time is a fixed value in this simulation, the time for a mature infected cell to produce virus and infect another cell is still exponentially distributed at a rate that depends on b and $x(t)$.

The results also exhibit an intriguing periodicity: At equilibrium, $\forall m, y(m, f) \approx y(\frac{1}{m} + T, f) \approx y(\frac{1}{m} + 2T, f)$. This suggests that the factor determining the degree of synchronization is $\frac{1}{m}$ mod T, as opposed to $\frac{1}{m}$ itself, i.e., there is no benefit to rapid reproduction in the intermediate regime. This likely occurs because we have assumed that immature infected cells do not die ($d_w = 0$), so there is not fitness loss to waiting longer to mature. In later sections we relax this assumption.

In the **high-drug-effectiveness regime** ($f > 0.95$), the off-window is very small. The probability that a cell which was infected during one off-window matures at the exact time to coincide with the next off-window is low enough that even cryptic resistance cannot insulate the viral population from the effects of the drug. Thus, the infected cell populations go to zero in this regime, and the healthy cell populations go to their pre-infection level, regardless of the lifecycle length.

Overall, these results suggest that viral strains with maturation times that were slightly "under-synchronized" to the drug period ($\frac{1}{m}$ mod T ≈ 1.75) reached higher than expected equilibrium populations despite the presence of the drug. Because in the absence of the drug $R_0 = 10$, if the drug is turned on more than 90% of the time($f < 0.9$), complete eradication of the infection is expected for cells that are unsynchronized and equally likely to mature at any time. However, synchronization allowed some strains to persist even at these high drug efficacies. Synchronized strains can match their phase to the drug dosing schedule such that they are much more likely to mature during the off-window than at other times.

These results give strong support to the idea that cryptic resistance could compromise antiviral treatment. However, while synchronized strains enjoyed greater long-term success than other strains when each strain was examined individually, these simulations shed no light on whether these strains were advantageous enough that they could invade and eventually dominate a population originally containing mostly non-synchronized strains. In particular, these results strongly

suggest that synchronized strains could outcompete others in the longer term, but perhaps this may not allow them to grow up from small numbers. When drug treatment first starts, many faster-reproducing, non-synchronized strains will still exist in high numbers. They could pose enough competition that synchronized strains would have a high probability of extinction due to stochastic events before they could get established. Additionally, it is not clear from the single-strain results which of the many possible synchronized strains would be favored. These strains may be able to invade each other, and given synchronization, it may be best to also reproduce as fast as possible. To answer these questions, a multi-strain model of direct competition between strains with different maturation rates in the presence of period drug levels is needed.

4.3 Multi-strain competition dynamics

To better determine whether cryptic resistance could emerge from a heterogeneous viral population and lead to treatment failure, we built a many-strain analogue of the stochastic model described in §4.3. For each strain, the reactions that could occur were those described in Table 2. Immature and mature infected cells were tracked for each strain, and each strain infected the same pool of uninfected cells. Since the purpose of the model was to explore the differences in fitness caused by differences in maturation time, all parameters except maturation time were held constant. Maturations times were varied linearly between 1 and 7 days to create 121 unique viral strains.

At the start of simulation, healthy cell populations were set to the infected equilibrium of the ODE model without drug (100 cells under the given parameter values), and the sum of immature and mature infected cells of all strains were also set to the respective equilibrium value ($y_i = 3$ and $w_i = 1$ for each strain). The simulation was run until a single strain remained ("fixed"), and then continued until each cell populations for this strain reached equilibrium. Equilibrium was said to be achieved when the average cell population levels over the last 10 days equal to those over the previous 10 days, and this value was returned. The simulation was run 500 times at each of 101 f-values ranging from 0 to 1.

We simulated this model for a variety of different parameters, drug-patterns, and maturation-time distributions. The results are presented in the subsequent sections.

4.3.1 Baseline Case

We first simulated the multi-strain model using all the same parameters and distributions as for the single strain model in Figure 4. These results are presented in Figure 5. As for the single-strain model, the fitness benefit of synchronized strains was observed clearly, manifesting itself as a high probability that these strains would outcompete others strains, fix in the population (Fig.5a) and reach levels approximating that when the drug was absent (Fig. 5b). We again observed that the fixation probability landscape could be divided into three regions corresponding to low, intermediate, or high drug efficacy. It was only in the intermediate "Goldilocks" regime when cryptic resistance was observed.

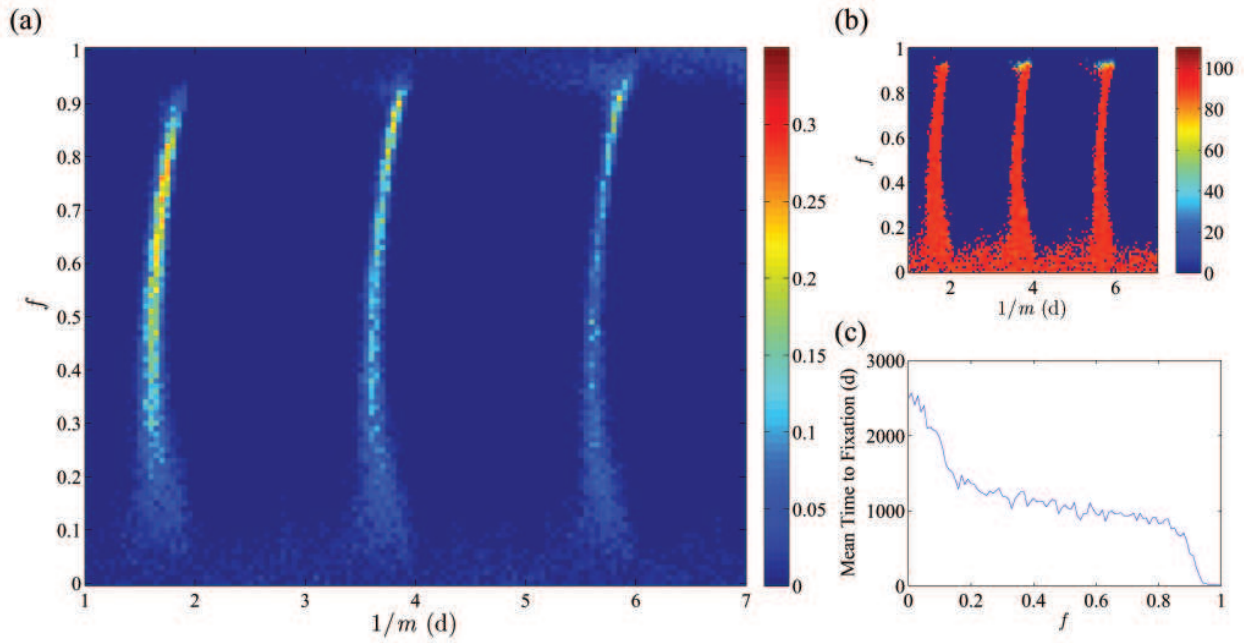

Figure 5: Results of the baseline multi-strain model. For each drug level f, strains with maturation times between 1 and 7 days were started at equal levels in the population and simulated until only a single strain remained. f is the fraction of time the drug is on in the on-off switch model of drug dynamics. a) The fixation probability (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. b) The average equilibrium level of mature infected cells (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. c) The mean time to fixation of a single strain as a function of the drug efficacy.

Motivated by the striking patterns observed in the fixation probability landscape, and their recurrence in varieties of the model presented later, we decided to define some terminology to describe and compare them. A schematic of the results is shown in Figure 6.

When visualized on a heatmap of maturation rate (m) vs drug on-time f, the advantage of syn-

chronized strains manifested itself in the Goldilocks regime in the form of **critical strips**. These are regions of relatively constant m-value which spanned the f-values of the Goldilocks regime between $f \approx 0.06$ and 0.97. These strips contained the vast majority of the fixation probability density in this regime. The critical strips are clustered around m-values whose reciprocals approximate integer multiples of the drug period T. The strips located near $mT = 1, 2, 3$ were respectively numbered I, II, and III.

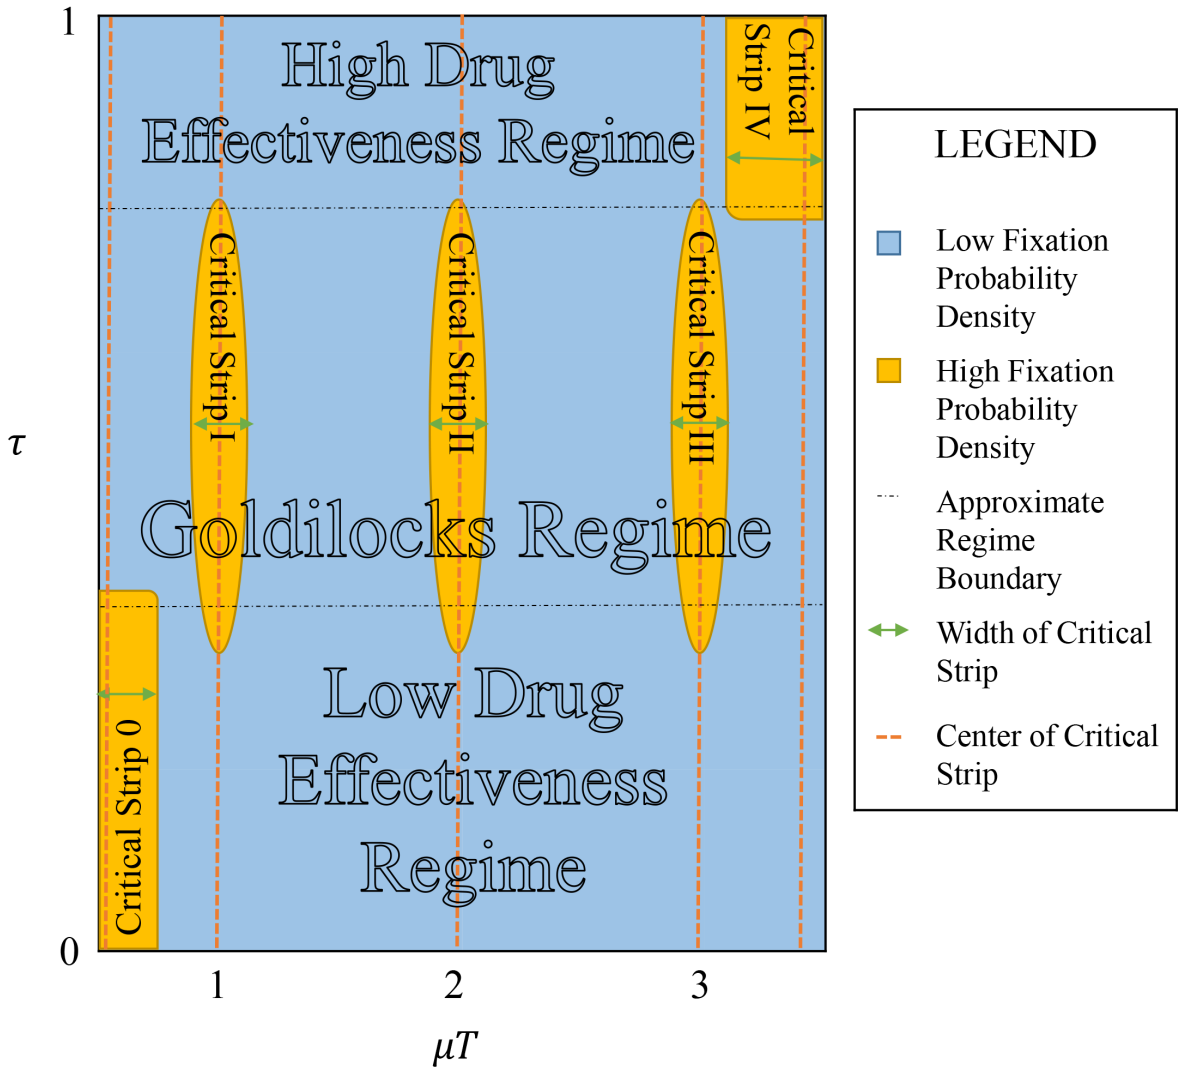

Figure 6: Terminology relevant to interpreting the results of the multi-strain stochastic model. A stylized fixation probability landscape in space outlining the definitions of various features of these graphs relevant to analysis of the system dynamics.

A "critical strip 0" was also defined as a similarly-shaped region spanning the low-drug-

effectiveness regime near the minimum maturation time. Although this regime was not observed for the baseline model, its importance will become apparent in later sections. We also defined a somewhat broader and fainter region spanning the high-drug-effectiveness regime near the maximum maturation times, which we termed the **high-f critical region**. Its utility will also become apparent later. The **width** of a critical strip was defined as the difference between the largest and smallest $\frac{1}{m}$-values contained within its cross-section at the average f-value of the regime it spanned, and the strip was said to be "centered on" the maturation time at which the fixation probability density was maximized. We now can use this terminology to further discuss the results.

The low-drug-effectiveness regime with $f \lessapprox 0.06$ featured a relatively uniform distribution of fixation probability density. In this regime, the drug is present so little of the time that a cell with any lifecycle length will have a good chance of maturing and producing virus at a time when the drug is absent. Thus, there is little selective pressure for synchronization, and fixation is governed almost purely by random chance. In this regime, fixation took approximately 2500 days to occur, which was far longer than in any other regime, suggesting intense competition between strains. The the equilibrium infection level of the strain that eventually fixed was independent of the maturation time, as expected, and approximately equal that which it would be in the drug-free scenario.

The high-drug-effectiveness regime existed for $f \gtrapprox 0.97$ and was characterized by a broad high-f critical region. In this inhospitable regime, the drug prevented infection of healthy cells almost 100% of the time. The average time to fixation was under three weeks in all cases, and even the strain that fixed went extinct shortly thereafter. This suggests that in the high-f regime, the increased fixation probability density in the high-f critical region did not reflect higher fitness of the strains therein. No strain was able to reproduce at these high drug doses, but the strains with the highest longest maturation times were simply able to wait longest in the immature state, without any risk of dying, before transitioning to the mature infected state and meeting their inevitable fate there.

The vast majority of f-values, however, fell within the Goldilocks regime, where cryptic resistance was observed to emerge. In this regime, the vast majority of the strains that fixed lay within critical strip I, II, or III, and for $f > 0.15$, every fixing strain in every simulation without exception lay in critical strip I, II, or III. Note that in reality there are an infinite number of these parallel critical strips, repeating whenever the maturation time is an integer multiple of the drug period. The finite number observed is simply because of the maximum m value we chose to simulation.

As in the single-strain model, we observed "under-synchronization". The reproductive success

seemed to be highest for strains whose lifecycle lengths are slightly less than integer multiples of the length of the interval between drug doses, with the optimal offset a decreasing function of f. As explained in §4.2, we hypothesize that this offset is due to the benefit of maturing near the start of an window in time where the drug is off. Since these windows open earlier for lower f-values, the expected optimal offset is larger there. However, if that were the sole constraint on optimal lifecycle length, then the relationship between optimal offset and f would be expected to be linear, which it is not. One possible explanation is that at sufficiently low f, the off-window is wide enough compared to the expected lifetime of a mature infected cell that the cell has a good chance of spending all of this phase of its lifecycle in an off-window even if it had matured in the middle rather than at the beginning of one, and this reduces the selective pressure to mature near the beginning of the window. The observed sublinear relationship between optimal amount of "under-synchronization" and f may thus reflect the combined effects of these two selective pressures.

Interestingly, although strains in critical strip I fix far more often than those in strip II or III strains over much of the Goldilocks regime, they do not have higher equilibrium infection levels when they fix. The fixation probability difference is thus hypothesized to reflect the difference in the speed of growth of different strains. Since critical strip I strains have the shortest lifecycle time, they can build up their populations most quickly, and are therefore expected to be the least vulnerable to stochastic extinction early on in treatment. This increased reproductive speed comes at the price of having more viral generations in a given length of time that must be precisely timed to mature near off-windows. Consequently, when the off-window narrows for higher drug efficacies,the advantage of rapid generation time becomes an Achilles heel, and it becomes less easy to build up the large population numbers. This is consistent with the observation that as f increases in the Goldilocks regime, first strip II and then strip III strains overtake strip I in strains in cumulative fixation probability.

We also observed distinct differences in the average fixation times between the different regimes. In the low-f regime, time to fixation averaged approximately 2500 days, in the Goldilocks regime a typical value was 1500 days, while no strain survived even three weeks in the high-f regime. These patterns, and the abrupt transitions in average fixation time between them, support the use of a tripartite classification of simulation outcomes.

4.3.2 Model variations

The results of the baseline model demonstrate that synchronization of the viral lifecycle with the drug dosing schedule was an extremely powerful competitive advantage when multiple strains are struggling to fix in a host. This synchronization lead the emergence of cryptic resistance, whereby infection levels could grow to near pre-treatment values even in the presence of ordinarily suppressive drug treatments. To further probe the biological realism of these results, we decided to relax some of the assumptions that we hypothesized to have contributed to the emergence of cryptic resistance, but were known to be less exact in a real human patient than allowed for in the model. We investigated five alternate formulation of the model parameters, both individually and combined. These variations are summarized in Table 3.

Table 3: Variations of the model used to examine the emergence of cryptic resistance under more biologically-relevant scenarios

Case	Description	Details
1	Pharmacologically-realistic drug dynamics	$b(t) \rightarrow b/(1 + (D(t)/IC_{50})^M)$, $D(t) = D_{\text{eq}}e^{kt}$
2	Immature infected cells may die	$d_w = 0.2$
3	Time to cell maturation can vary	$t \sim \text{Gamma}(\frac{1}{m_i}, \sigma)$, $\sigma = 0.30d$, $\forall$ strains i
4	Time between drug doses can vary	$t \sim \text{Gamma}(T, \sigma)$, $\sigma = 0.30d$
5	Patient imperfectly adherent to treatment	$P(\text{non-adherance}) = 0.3$

Firstly, the assumption that a drug is only "on" or "off", and that suppression is completely when it's on and absent when it's off, is obviously not realistic. In reality, the combination of pharmacokinetics (changes in drug levels in the body over time after a dose is taken) and pharmacodynamics (the concentration-dependent effect that drug has on viral replication) lead more continuous changes in drug levels. We investigate these dynamics in Case 1. Secondly, the drug dynamics may not be perfectly periodic, because drug doses may not be taken at the exact same time each day. In Case 4, we allow the time between doses to vary. Thirdly, for many drugs it is well-established that patients are not perfectly adherent to treatment. In Case 5 we allow for patients to skip doses. Fourthly, the assumption of immortality for immature infected cells is not realistic. At minimum, even if early stage infection was completely benign, the death rate should be at least equal to that of uninfected cells. In Case 2 we allowed for $d_w > 0$. Finally, it is unreasonable to believe that immature cells mature after an exactly fixed time interval. In Case 3, we allow for variance in the time to maturation for a given strain.

In the subsequent sections, we determine if cryptic resistance can still emerge in the presence of these more realistic dynamics.

4.3.3 Case 1: Drug dynamics follow realistic pharmacologic functions

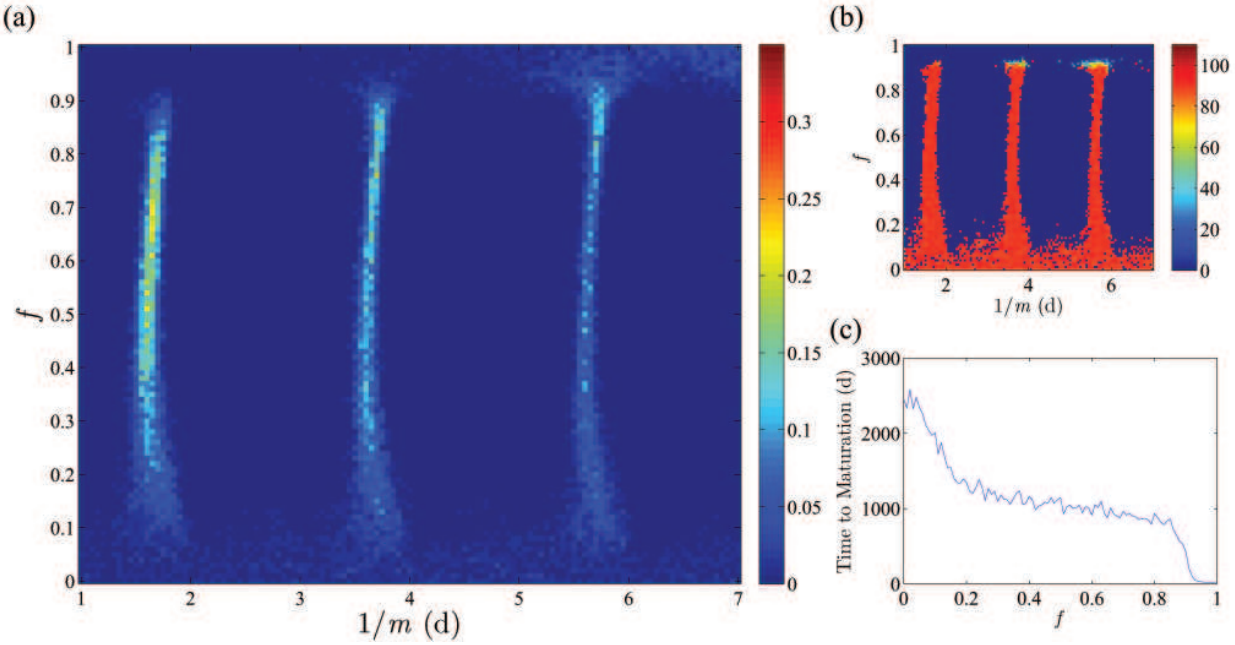

Figure 7: Case 1: Results of the multi-strain model with realistic drug dynamics. Drug dynamics were modeled with peak-and-exponential-decline pharmacokinetics and Hill-curve pharmacodynamics. For each drug level f, strains with maturation times between 1 and 7 days were started at equal levels in the population and simulated until only a single strain remained. f is the fraction of time the drug is on in the on-off switch model of drug dynamics. a) The fixation probability (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. b) The average equilibrium level of mature infected cells (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. c) The mean time to fixation of a single strain as a function of the drug efficacy.

We sought to construct a more realistic description of how drug levels change over time, and how this affects viral fitness. Because of the importance of "on" and "off" windows of drug in the simpler switch model, we hypothesized that the drug dynamics may play an important role in promoting or inhibiting the emergence of cryptic resistance.

To model drug dynamics, we chose a simple and commonly used description. We assumed that drug concentrations increased to their maximum value immediately after a dose was taken, and then decayed exponentially until the next dose. If each dose increases the drug concentration by an amount C_0, then eventually the dynamics reach a periodic steady state in which the amount of drug

eliminated from the body between doses exactly equals the amount added with each subsequent dose. In this steady state, $D(t) = D_{\text{eq}}e^{-k(t-t^*)}$, where $D_{\text{eq}} = \frac{D_0}{1-e^{-kT}}$ and t^* is the time of the most recent drug dose. The decay parameter k is related to the drug's half-life by the expression $k = \ln(2)/t_{1/2}$.

We also need a function to describe how the infection rate depends on the drug levels. For most drugs, the action can be described by a sigmoidal "Hill-curve" equation, $b(t) = \frac{b_0}{1+(\frac{D(t)}{IC_{50}})^M}$. Here b_0 is the infection rate in the absence of the drug, $D(t)$ is the concentration of the drug at a particular time, IC_{50} is the drug concentration at which the infection rate is halved, and m is a parameter describing the steepness of the curve.

The model was identical to the baseline model except for the drug dynamics which gave $b(t)$. For parameter values, we chose $D_0 = 50,000 IC_{50}$ and $m = 1$. For each f-value tested in the baseline model, a corresponding k-value to be tested was defined by solving for k in the equation

$$\int_0^T \frac{1}{1+\left(\frac{C_{eq}e^{-kt}}{IC_{50}}\right)^M}\mathrm{d}t = \int_{fT}^T (1-f)T\mathrm{d}t$$

so that the average rate of infection in corresponding simulations would be equal and thus results would be comparable.

The results of Case 1 were largely identical to those of the baseline model. We also exhibited a three regions of dynamics as a function of k, directly analogous to that exhibited in the baseline model a function of f: A high effective drug concentration regime prevailed for k corresponding to $f \geq 0.94$, a Goldilocks regime extended over the range of k-values corresponding to $f \in [0.08, 0.93]$, and a low effective drug concentration regime at k-values corresponding to the lowest f-values of the baseline model.

The one noticeable effect of the Hill curve drug dynamics on the simulation results seems to be an increased width of critical strips I-III relative to that seen in the baseline model. Critical strips average 0.20 days in width, which is significantly wider than the 0.14 d seen before. This extra thickness is consistent with the belief that the realistic drug dynamics decrease the sharpness of fitness fluctuations in between doses, thereby lessening the precision of timing required to make cryptic resistance work and consequently expanding the range of m-values that are able to sustain it. Re-running this case with $\frac{D_0}{IC_{50}}$ taking on different values between 10^3 and 10^7 appeared con-

firmed this hypothesis. The width of the critical strips was a decreasing function of $\frac{D_0}{IC_{50}}$, and the larger $\frac{D_0}{IC_{50}}$, the more "step-function-like" $C(t)$ became. As in the baseline model, however, there was no difference in any of these case 1 simulations between strains' equilibrium mature infected cell populations at fixation. Thus, the only notable effect of adding realistic drug dynamics to the model seems to have been an increase in the range of maturation times that could lead to cryptic resistance. However, this affected solely the fixation probability and not the equilibrium infection level. Overall, cryptic resistance could still emerge despite the introduction of realistic drug dynamics.

4.3.4 Case 2: Immature infected cells can die before maturing

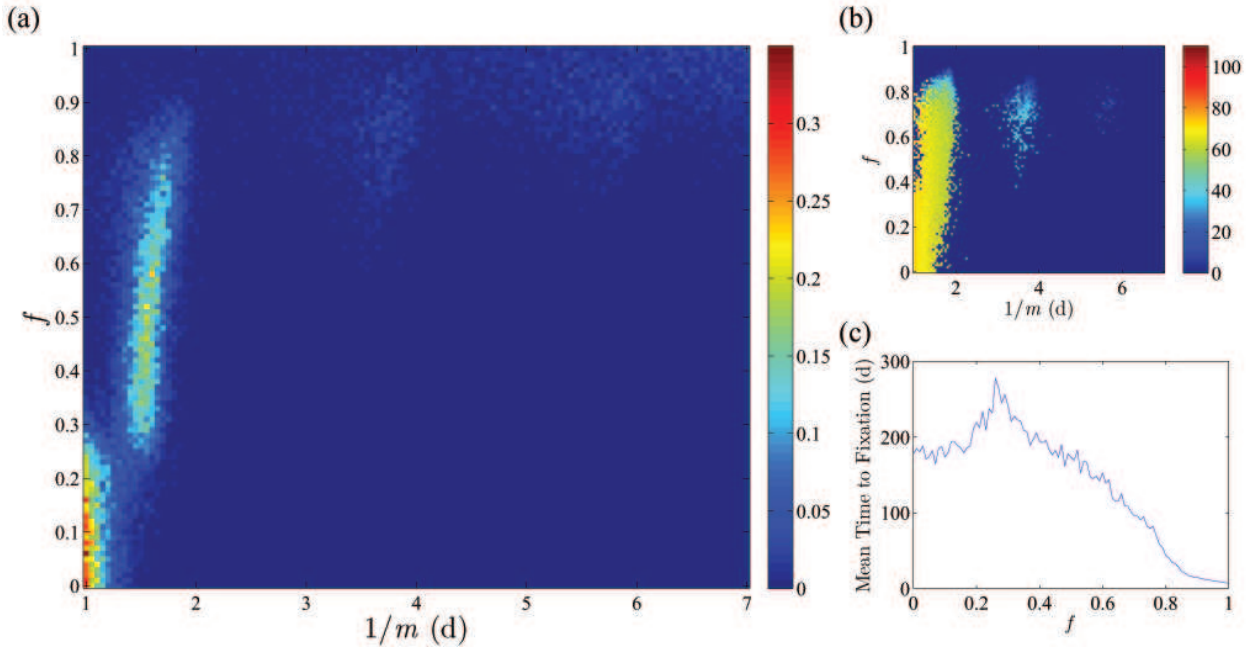

Figure 8: Case 2: Results of multi-strain model with death of immature cells. In this figure,$d_w = 0.2\text{d}^{-1}$ and immature infected cells died with time to next immature infected cell death of strain i exponentially distributed with mean $\frac{1}{d_w w}$. For each drug level f, strains with maturation times between 1 and 7 days were started at equal levels in the population and simulated until only a single strain remained. f is the fraction of time the drug is on in the on-off switch model of drug dynamics. a) The fixation probability (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. b) The average equilibrium level of mature infected cells (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. c) The mean time to fixation of a single strain as a function of the drug efficacy.

The only difference between Case 2 and the baseline model is that in Case 2, immature infected cells can die, and the time to the next death of an immature infected cell of strain i is exponentially

distributed with mean $\frac{1}{d_w w_i} = \frac{1}{0.2 w_i}$. However, the fixation probability landscape, mean time to fixation, and equilibrium infected cell population landscapes of Case 2 are vastly different from those of the baseline model. While the behavior of Case 2l can still be broadly classified into a low-drug-effectiveness regime ($f < 0.27$), a Goldilocks regime ($f \in (0.265, 0.9)$), and a high-drug-effectiveness regime ($f > 0.9$), the Goldilocks regime in Case 2 covers a far smaller range of f-values in Case 2 than in the baseline model (Figure 8a). This suggests that the possibility of infected cell death in the early-infected stage has a significant adverse impact on the competitiveness of strains that synchronize lifecycle length with the dosage cycle strains because they have to spend more time in the immature stage of the viral lifecycle than faster-reproducing strains do. This fitness difference is also reflected in the fact that in the absence of a drug, $R_0 = \frac{\lambda k b m}{d_x d_y (m + d_w) c}$ is a decreasing function of lifecycle length when all other parameter values are held constant. Consequently, the fitness reduction imposed by the drug needs to be higher in Case 2 than in the baseline model for the benefits of drug-virus synchronization to outweigh the fitness cost of increased attrition of early-stage infected cells imposed by the longer lifecycle.

This analysis of the dynamics of Case 2 accords well with the observed differences in the shapes and sizes of the critical strips in Case 2 and the baseline model. In stark contrast to the baseline model, strip I strains are the only ones to fix in any of the simulations conducted over the majority of the Goldilocks regime $f \in (0.30, 0.70)$. Strip II and III strains are entirely absent over most of this range, fixing with (small) nonzero probability only above $f = 0.70$, and even there, strip I strains fix the vast majority of the time. This dominance of strip I is almost certainly a testament to the major fitness impact of early-stage infected cell death on cryptic resistance-capable viral strains. Strip II and III strains, which have to spend over 2 and 3 times as long, respectively, as strip I strains in the early-infected stage of the viral lifecycle, experience so much more attrition in these lifecycle phases than do strip I strains that they are outcompeted almost all of the time in the Goldilocks zone. Even when they do fix, their equilibrium population levels are far lower than those of strip I strains (Figure 8a).

Another noteworthy difference between the baseline model and Case 2 lies in the behavior of the model in the low-drug-effectiveness regime. In the baseline model, this regime is characterized by a more-or-less uniform distribution of fixation probabilities across the various strains (Figure 5a), reflecting the fact that all have identical R_0-values in the absence of drug (which is indeed the state of affairs more than 90% of the time in this regime), and thus probably approximately equal fitness as well. In Case 2, however, the low-drug-effectiveness regime is dominated by strains in critical strip 0, which contains the only strains able to fix in any simulations conducted in this

regime. This seems to be yet another reflection of the powerful pressure exerted by d_w in favor of rapid progression of infected cells to maturation.

The addition of a nonzero d_w to the bas appears to have some subtler effects on the shape of critical strip I as well. The strip is centered near $f = 1.6$ in this model, lower than in the baseline model, and at 0.40 d in width is far broader than the critical strips of the baseline model. Like the rest of the differences between this model and the baseline model, these features are believed to be a reflection of the trade-off between fast reproduction and cryptic resistance, which favors fixation of strains which compromise slightly imperfect drug-virus synchronization with faster reproduction times and thereby allows a greater number of more poorly-synchronizing viral strains to have comparable viral fitness to more precisely synchronizing strains, smoothening the fixation probability landscape.

Overall, Case 2, unlike Case 1, had a profound effect on the dynamics of the model. When the fact that (real life) nonzero d_w favors faster viral reproduction is accounted for, the fitness landscape of the Goldilocks zone undergoes what can only be described as a sea change. Its lower reaches in the baseline model become part of the low-concentration zone, and the remainder shifts from slightly favoring strip II and III strains to seeing strip I strains enjoy an overwhelming competitive advantage over all others, and equilibrium population of the fixing strain becomes a decreasing function of $\frac{1}{m}$ and thus of lifecycle length. These changes appear to confirm our intuition that the principal effect of newly-infected cell attrition is to introduce selection pressure in favor of faster viral reproduction. Indeed, when the Case 2 simulations were re-run several times with different values of d_w between 0.05 and 0.40 d^{-1}, the Case 2-specific dynamics of the model were more pronounced for larger d_w.

However, it is a testament to the powerful advantage afforded by cryptic resistance that the selective pressures brought about by changes of model Case 2 are still sufficiently weak even at $d_w = 0.4\text{d}^{-1}$ that CRC strains can outcompete faster-reproducing strains that have greater R_0-values in the absence of the drug.The vast majority of f-values remain part of the Goldilocks zone, and cryptic resistance easily and consistently outcompetes rapid reproduction as a viral reproductive strategy in that zone.

4.3.5 Case 3: The time to maturation is variable

For model Case 3, the distribution chosen to model time to maturation for each strain was a Gamma distribution whose mean was the reciprocal of its m_i-value and whose standard deviation

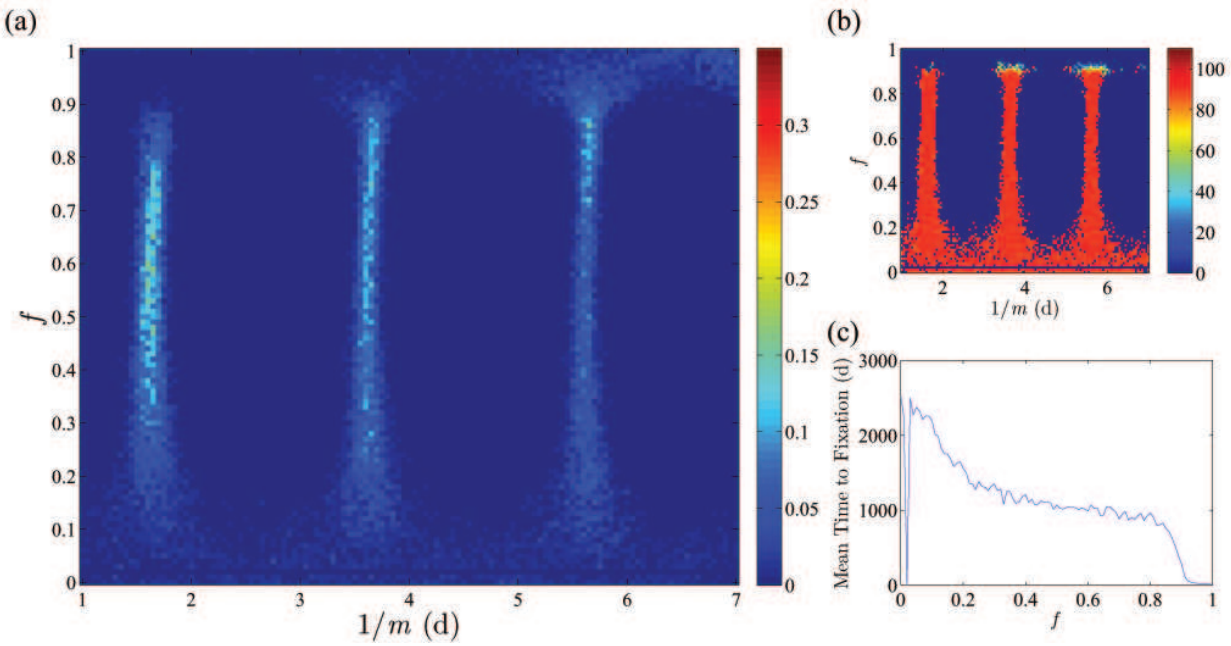

Figure 9: Case 3: Results of multi-strain model with variable time to maturation. Time to maturation of immature infected cells of strain i was assumed to be Gamma-distributed with mean $\frac{1}{m_i}$ and standard deviation 0.30 d. For each drug level f, strains with maturation times between 1 and 7 days were started at equal levels in the population and simulated until only a single strain remained. f is the fraction of time the drug is on in the on-off switch model of drug dynamics. a) The fixation probability (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. b) The average equilibrium level of mature infected cells (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. c) The mean time to fixation of a single strain as a function of the drug efficacy.

was 0.30 d. The Gamma distribution was chosen because it allowed mean to be controlled independently of variance, and the variance of the distribution was held constant at 0.30 d for each strain to ensure that the mean absolute deviation of each strain from perfect synchronization would be the same.

The results of this model were almost identical to those of Case 1 – similar in all main points to the baseline model, but with broader (0.30 d) critical strips believed to stem from the fact that Gamma-distributing the time to maturation, like Hill curve drug dynamics, diminished the effectiveness of cryptic resistance capabilities in ensuring reproductive success and thereby levelled the fitness differential between strains with CR-optimal lifecycle lengths and those with slightly less optimal lifecycle lengths for CR purposes.

The one key difference between the baseline model and Case 3 was that strip I strains were not as heavily favored for fixation in the Goldilocks regime of Case 3 as they were in the baseline

model. This is believed to reflect the fact that the variance of the distribution of time to maturation was held constant between the different simulations. This caused the variance in the sums of the maturation times of three successive generations of a strip I strain to be nine times that of the variance of the maturation time of the single generation of a contemporaneous strip III strain. As a result, the expected "synchronization drift" was much more severe in the faster-reproducing strains than in the slower-reproducing ones, making the strip III strains better able to consistently reap the benefits of cryptic resistance than could those in strips I and II. Furthermore, when this case's simulations were re-run with different σ between 0.10 and 0.50 d, the degree to which strip III strains were favored in the Goldilocks regime was an increasing function of σ.

Overall, therefore, like Case 1 and Case 2, the changes brought about in Case 3 mitigate the advantages of cryptic resistance but not nearly enough to prevent strains capable of cryptic resistance from dominating a wide Goldilocks regime.

4.3.6 Case 4: The time between drug doses is variable

As in Case 3, a Gamma distribution was chosen to model a probabilistic phenomenon – this time, the time until the next drug dose, which was assumed to be Gamma-distributed with mean T and variance 0.30 d in order to make this simulation's changes analogous to those of Case 3. It was suspected that this model, like Case 3, would be dominated by the effects of the modification on maturation timing, and thus show similar trends to Case 3. However, while Case 3 strongly resembles the baseline model and Case 1 in its outcomes, Case 4 looks strikingly different. Strip I is far more strongly favored in Case 4 than in Case 3, to the point where strip III is so faint as to be entirely absent throughout all but the very edges of the Goldilocks regime in parameter space, and the mean time to fixation for simulations conducted under the Goldilocks regime of Case 4 are far lower than those of the analogous simulations in Case 3 and the baseline model (Figure 10).

A careful consideration of these discrepancies reveals that they support, rather than contradict, the hypothesis that the same dynamics of variance dominating Case 3 drive Case 4 as well. By the same reasoning as in Case 3, the variance in the timing between a drug dose and the one after the one after the one after it is nine times as great as that between a drug dose and the succeeding one. However, in Case 3, each infected cell's maturation time is independent of the others', so given a sufficiently large population of infected cells of a given strain, in every generation there will be a non-negligible probability that at least one will mature at an opportune time near the beginning of a window and stave off extinction for the strain. However, in the case of Case 4, the

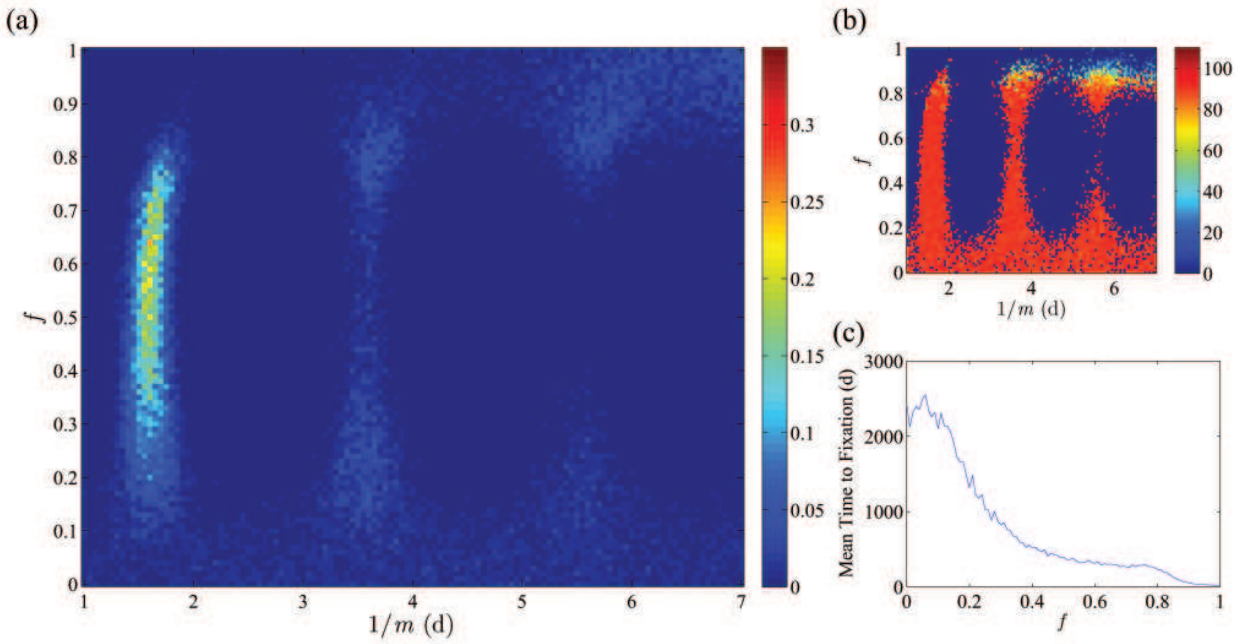

Figure 10: Case 4: Results of multi-strain model with variable times between drug doses. Time to the next drug dose was assumed to be Gamma-distributed, with mean $T = 2$ d and standard deviation 0.30 d. For each drug level f, strains with maturation times between 1 and 7 days were started at equal levels in the population and simulated until only a single strain remained. f is the fraction of time the drug is on in the on-off switch model of drug dynamics. a) The fixation probability (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. b) The average equilibrium level of mature infected cells (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. c) The mean time to fixation of a single strain as a function of the drug efficacy.

same variation in drug timing affects all cells alive in a given generation simultaneously; a single unusually early drug dose will have a high probability of killing off large numbers of cells in a "mass extinction event" against which the best protection is the large population size enjoyed by the strain that ultimately fixes. As a result, the number of strains alive is quickly reduced, sharply cutting mean fixation time relative to Case 3 and the baseline model, but the equilibrium infected cell populations of the surviving strain are approximately equal to what they were in Case 3 and the baseline model. Since strip I strains can increase in population far more rapidly than strip II or III strains, they are far less vulnerable to these mass extinction events, and thus fix far more frequently over the main part of the Goldilocks regime. (The boundary between the Goldilocks regime and the high-drug-effectiveness regime, the only part of the Goldilocks regime where strip III strains fix more frequently than strip I strains, is dominated by a different dynamic mentioned above in the discussion of the baseline model that favors strip III strains over strip I strains in that same region in model the baseline model.) Thus, the net effect of Case 4, like that of Case 2 is to

preclude the development of viable strip II and III CRC strains by exerting selective pressure in favor of rapid reproduction. Unsurprisingly, the severity of this pressure proved to be an increasing function of σ when this case's simulations were re-run with different σ between 0.10 and 0.50 d.

4.3.7 Case 5: Patient adherence to treatment is imperfect

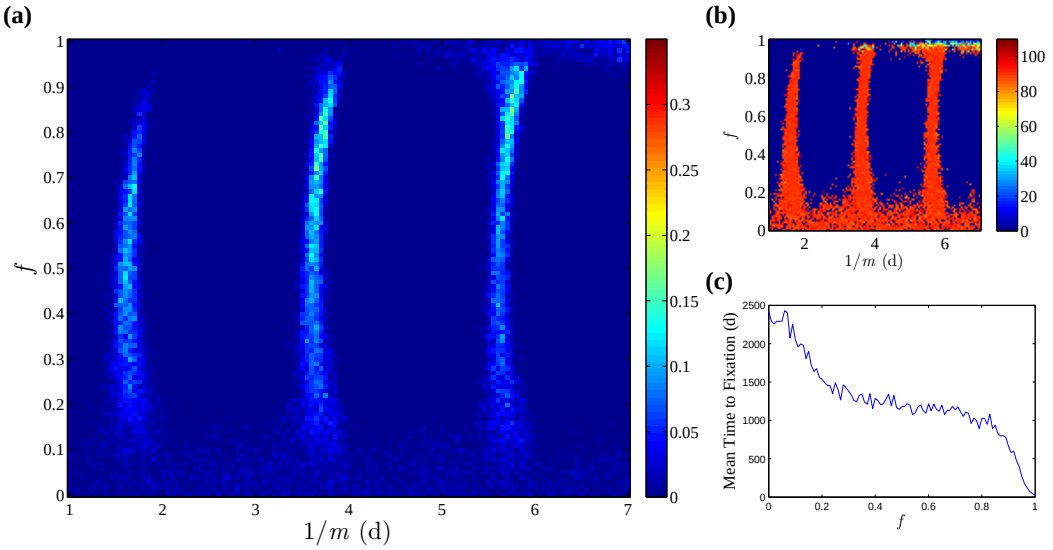

Figure 11: Case 5: Results of model with imperfect patient adherence to treatment. Imperfect adherence was modeled by mandating that each scheduled drug dose was taken as planned with probability 0.7, and skipped with probability 0.3. The adherence or lack thereof to each scheduled drug dose was assumed to be independent of all other events in the simulation, including the adherence or lack thereof to other scheduled drug doses. For each drug level f, strains with maturation times between 1 and 7 days were started at equal levels in the population and simulated until only a single strain remained. f is the fraction of time the drug is on in the on-off switch model of drug dynamics. a) The fixation probability (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. b) The average equilibrium level of mature infected cells (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. c) The mean time to fixation of a single strain as a function of the drug efficacy.

To simulate imperfect patient compliance with the drug treatment, the baseline model was altered by making the time to the next drug dose equal to the product of T with a geometric-distributed random variable of mean $\frac{10}{7}$. With selective pressure for CR temporarily vanishing from time to time, it is easier to build up a significant population of infected cells and thus the slower reproduction of strip III strains is less salient to their fixation probability and they fix more frequently than in the baseline model. However, this is the only difference of any significance between the baseline model and Case 5, and incorporating the 30% noncompliance rate of Case 5 does little to prevent cryptic resistance from remaining a viable reproductive strategy (Figure 11).

Even when the simulation was re-run with non-compliance rate as high as 50%, cryptic resistance was still a dominant strategy over a broad Goldilocks zone.

4.3.8 Combination of Cases 1 – 5

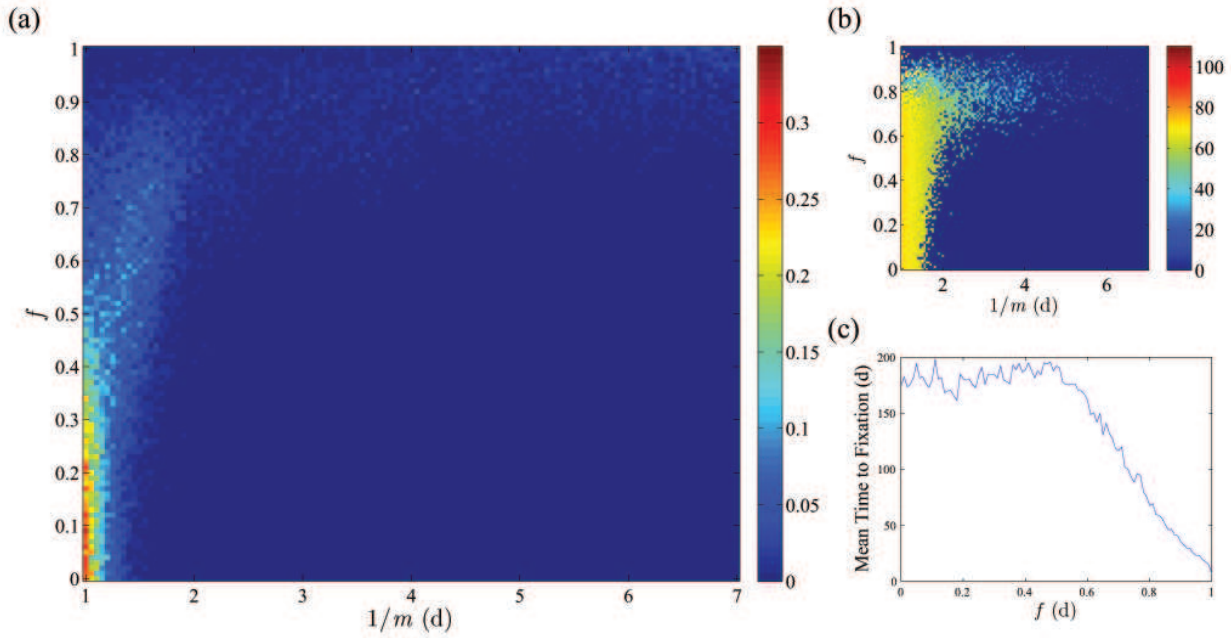

Figure 12: Results of incorporating all the changes of Cases 1-5. For the final simulations combining all cases, parameter values were set to $\frac{D_0}{IC_{50}} = 50000$, $d_w = 0.2\text{d}^{-1}$, $\sigma_{\text{drug dose}} = 0.30d$, $p_{\text{adherence}} = 0.7$, and $= \sigma$maturation time $= 0.30d$ given adherence to the next drug dose. For each drug level f, strains with maturation times between 1 and 7 days were started at equal levels in the population and simulated until only a single strain remained. f is the fraction of time the drug is on in the on-off switch model of drug dynamics. a) The fixation probability (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. b) The average equilibrium level of mature infected cells (red gives highest probability) as a function of the maturation rate (m) for a given drug efficacy. c) The mean time to fixation of a single strain as a function of the drug efficacy.

In a final confirmation of the viability of cryptic resistance as a viral reproductive strategy, all the cases were incorporated into one set of simulations, with $\frac{D_0}{IC_{50}} = 50000$, $d_w = 0.2\text{d}^{-1}$, $\sigma_{\text{drug dose}} = 0.30d$, $p_{\text{adherence}} = 0.7$, and $= \sigma$maturation time $= 0.30d$ given adherence to the next drug dose. Combining cases did not lead to any new effects; at the parameter values tested, each combination essentially resembled one of the models whose changes it incorporated, while showing subtle signs of influence from the other changes. Moreover, there was a clear hierarchy of importance of the different modifications made to the baseline model on the dynamics of the system. Case 2, which selected against long maturation times, was dominant. Any model containing

its modification featured a sharp critical strip 0 in the low-drug-effectiveness regime and a critical strip I but no II or III in the Goldilocks regime. The second most important modification was that of Case 4, whose influence was also characterized by a critical strip I but little or no II or III in the Goldilocks regime and no critical strip 0 in the low-drug-effectiveness regime. Next was the modification of Case 5, which allowed for the presence of all three critical strips in the heart of the Goldilocks regime but somewhat strongly heavily favored III over II over I near the border of the high-drug-effectiveness regime. Case 3, which favored III and II in the Goldilocks regime slightly more than the baseline model and slightly broadened all three critical strips, was only dominant over Case 1, which was not dominant over anything. The only exceptions to this otherwise strict hierarchy were the simulations combining cases (3, 4, 5) and (1, 3, 4, 5), which resembled Case 3 and Case 5 even though Case 4 is dominant over both Case 3 and Case 5. This is believed to be because Case 3 and Case 5 both favor critical strip II and III strains, which Case 4 disfavors, and Case 4 is not sufficiently more powerful than either that it can dominate the dynamics of a model which also incorporates both Case 3 and Case 5.

In accordance with this hierarchy, the most realistic model incorporating all five changes of models Case 1-Case 5 behaved like Case 2 (Figure 12). While critical strip I was extremely wide (0.80 d) and the Goldilocks regime much smaller (k corresponding to $f \in (0.4, 0.7)$) in the former than the latter, this is consistent with the lesser but still non-negligible combined influence of the other modifications, of which (1, 4) tended to broaden critical strips and all of which reduced the effectiveness of CR as an evolutionary strategy. Nevertheless, despite the presence of all 5 modifications, critical strip I strains have a clear evolutionary advantage over all others at drug effectivenesses in the Goldilocks regime, a significant swathe of parameter space, and CR-capable strains are able to maintain an equilibrium population of 60-70 infected cells (cf 100 in model the baseline model in the low-drug-effectiveness regime) of the drugless infected equilibrium value, far too high to consider the virus eradicated.

5 Discussion

The results of the probabilistic simulations strongly suggest that cryptic resistance is a powerful strategy for increasing the fitness of viral strains in the face of a (somewhat) regularly-administered antiviral drug dose. Strains synchronizing their lifecycle lengths with the antiviral drug dosage schedule achieve higher equilibrium populations and fix with far higher probability than all other strains under a wide range of drug effectivenesses in the final, most realistic model. These abilities

proved robust to fluctuations in timing of drug doses and timing of viral maturation, and are therefore likely to be present in actual biological systems. The main constraints on cryptic resistance in biologically realistic regimes appear the limitation on lifecycle length imposed by immature cell mortality and irregularity of drug dose timing, both of which preclude any but strip I strains from viably employing cryptic resistance as a reproductive strategy. Furthermore, because cryptic does not rely on aspects of the viral lifecycle or antiviral drug dosage patterns that are specific to one particular drug or virus, nor is it particularly sensitive to the value of f, it is plausible that cryptic resistance could play a significant role in the emergence of resistance to a wide variety of drug treatments in a wide variety of viral populations.

Any mathematical system used to understand a biological phenomenon such as cryptic resistance must in accurately represent the important features of the underlying biology if it is to be meaningful. The most obvious biological concern about the proposed "cryptic resistance" scenarios is, "How could a virus control the length of its own lifecycle?' Viral genomes are often so simple that it is hard to imagine them being capable of possessing the complexity needed to evolve a biological timekeeping mechanism. Moreover, as no virus has been observed to possess such a mechanism (to our knowledge), the number of mutations likely needed for a viral population to evolve such a mechanism *de novo* over the course of a single infection makes it a highly improbable development.

While the obstacles a viral population would need to overcome to develop a true timekeeping mechanism may be insurmountable, there is a far simpler way a viral mutation could crudely affect the rate at which it completes its lifecycle: by altering the ease with which certain tasks along the lifecycle can be completed. For example, a key step in the viral lifecycle is the hijacking of cellular machinery to translate viral proteins and assemble them into new virus particles. Mutations to promoter regions of viral genes, such as transcriptases, could affect the affinity of cellular translation machinery for them and thereby increase or decrease the length of the time between when viral gene expression starts and when new virions are released from the host cell. We have called this time the "maturation time". Alternatively, mutations to the viral proteins that recognize and bind receptors on target cell membranes may alter the rate of virus entry, thereby altering the expected duration of the "release and re-infect time" comprising the duration of the "free virus" stage of the viral lifecycle. Mechanisms like these could allow simple, feasible point mutations in the viral genome to lead to significant changes in viral lifecycle length, allowing this length to synchronize with the time interval between doses of the antiviral drug.

Throughout this paper, we have assumed that the drug targets either the production and release of viral particles or the infection of new host cells by free virus. Consequently, we have assumed that it was the maturation time of the virus that could be varied and could evolve in the presence of drug treatment. However, the alternative scenario could have been examined, in which the drug acts to prevent maturation of infected cells, and the virus may instead vary the release and re-infect time. Which scenario is more likely may depend on a few factors. One is how and where the drug acts: if it works extracellularly by binding free virus particles, or if it acts intracellularly and binds viral proteins or genes. The second is what type of mutations can alter the lifecycle time without having detrimental fitness costs for the virus. While changes to viral proteins would likely be needed to alter the "release and re-infect time", either viral proteins or viral genetic regulatory elements could be altered to vary the "maturation time". We expect that similar results to hold independent of the choice between these two scenarios, and hope to investigate this question in future work.

Needless to say, further work is needed to confirm the existence of cryptic resistance through *in vivo* or *in vitro* studies. If it exists in nature, its power to disrupt antiviral treatments would be significant, and it may help explain why the rate of failure of antiviral drug treatment is not any lower than it is. Fortunately, even if cryptic resistance is indeed widespread, it should be relatively easy to prevent. Since it thrives on regular fluctuations in drug concentration, administering antiviral drugs continuously via IV instead of periodically in pill form would preclude it from developing entirely, boosting the probability of success of antiviral drug drug treatments significantly without the monetary and temporal cost of developing novel antiviral drugs.

Acknowledgements

The simulations upon which many of the results in this thesis are based would not have been possible without the access to the Orchestra computer cluster afforded to me by Harvard's Program for Evolutionary Dynamics. It also could not have come to fruition without the research advice, mathematical critiques, and biological knowledge of Dr. Alison Hill, to whom I am deeply grateful for helping me to both make sense of the deceptively complicated phenomenon of cryptic resistance, and organize my thoughts and insights on the project into a coherent whole. I also would like to thank the Program for Evolutionary Dynamics's Prof. Martin Nowak, whose wonderful MATH 153 class introduced me to the natural beauty of mathematical biology and ultimately led me to conduct the research that led to this thesis.

References

[1] Wahl, L. M. & Nowak, M. A. Adherence and drug resistance: predictions for therapy outcome. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **267**(1445), 835 – 843, April (2000).

[2] Nowak, M. A. & May, R. M. C. *Virus dynamics: mathematical principles of immunology and virology*. Oxford University Press, USA, (2000).

[3] Neumann, A. U., *et al.* Hepatitis C Viral Dynamics in Vivo and the Antiviral Efficacy of Interferon-? Therapy. *Science* **282**(5386), 103–107, October (1998).

[4] Ciupe, S. M., Ribeiro, R. M., Nelson, P. W., & Perelson, A. S. Modeling the mechanisms of acute hepatitis B virus infection. *Journal of Theoretical Biology* **247**(1), 23–35, July (2007).

[5] Perelson, A. S., Neumann, A. U., Markowitz, M., Leonard, J. M., & Ho, D. D. HIV-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time. *Science (New York, N.Y.)* **271**(5255), 1582–1586, March (1996).

[6] Otto, S. P., Day, T., & Day, T. *A biologist's guide to mathematical modeling in ecology and evolution*. Princeton University Press, (2007).

[7] Sedaghat, A. R., Dinoso, J. B., Shen, L., Wilke, C. O., & Siliciano, R. F. Decay dynamics of HIV-1 depend on the inhibited stages of the viral life cycle. *Proceedings of the National Academy of Sciences* **105**(12), 4832 (2008).

[8] Alexander, H. K. & Bonhoeffer, S. Pre-existence and emergence of drug resistance in a generalized model of intra-host viral dynamics. *Epidemics* **4**(4), 187–202, December (2012).

[9] Ribeiro, R. M., *et al.* Estimation of the Initial Viral Growth Rate and Basic Reproductive Number during Acute HIV-1 Infection. *Journal of Virology* **84**(12), 6096 –6102, June (2010).