J. Math. Biol. (2007) 54:385–406
DOI 10.1007/s00285-006-0059-1

Mathematical Biology

Costs versus benefits: best possible and best practical treatment regimens for HIV

O. Krakovska · L. M. Wahl

Received: 29 November 2005 / Revised: 4 November 2006 /
Published online: 5 January 2007

Abstract Current HIV therapy, although highly effective, may cause very serious side effects, making adherence to the prescribed regimen difficult. Mathematical modeling may be used to evaluate alternative treatment regimens by weighing the positive results of treatment, such as higher levels of helper T cells, against the negative consequences, such as side effects and the possibility of resistance mutations. Although estimating the weights assigned to these factors is difficult, current clinical practice offers insight by defining situations in which therapy is considered "worthwhile". We therefore use clinical practice, along with the probability that a drug-resistant mutation is present at the start of therapy, to suggest methods of rationally estimating these weights. In our underlying model, we use ordinary differential equations to describe the time course of in-host HIV infection, and include populations of both activated $CD4^+$ T cells and $CD8^+$ T cells. We then determine the best possible treatment regimen, assuming that the effectiveness of the drug can be continually adjusted, and the best practical treatment regimen, evaluating all patterns of a block of days "on" therapy followed by a block of days "off" therapy. We find that when the tolerance for drug-resistant mutations is low, high drug concentrations which maintain low infected cell populations are optimal. In contrast, if the tolerance for drug-resistant mutations is fairly high, the optimal treatment involves periods of reduced drug exposure which consequently boost the immune response through increased antigen exposure. We elucidate the dependence of the optimal treatment regimen on the pharmacokinetic parameters of specific antiviral agents.

O. Krakovska (✉) · L. M. Wahl
Department of Applied Mathematics, University of Western Ontario,
London, ON, N6A 5B7 Canada
e-mail: okrakovs@uwo.ca

Keywords Human immunodeficiency virus · Drug treatment · Optimization · HAART · Mathematical model · Side effects · Resistance

1 Introduction

The human immunodeficiency virus (HIV) infects and ultimately leads to the depletion of a critical component of the cellular immune response, the $CD4^+$ T cell population. As HIV infection progresses, the concentration of these helper T cells gradually wanes, eventually crippling the immune response. When the level of helper T cells is sufficiently low, the patient has little resistance to other "opportunistic" infections and a diagnosis of AIDS (acquired immunodeficiency syndrome) is obtained.

Although many currently available antiviral drugs are highly effective, they may also cause severe and even life-threatening side effects. This suggests the possibility of using mathematical models to design and evaluate "drug-sparing" treatment regimens which reduce overall drug exposure, while maximizing helper T cell counts and thus delaying the progression to AIDS.

If patient T cell counts can be regularly observed, drug regimens might be dynamically adjusted to optimize treatment outcome; for recent examples of such feedback control, see [20,40]. Unfortunately, T cell counts are not always readily or regularly available. In the absence of this regular feedback, treatment regimens are usually evaluated mathematically by constructing a target function which balances the costs of side effects against the benefits of increased T cell counts. Prior research has examined target functions that maximize the helper T cell count at each time point [8,15,21,24] or at the final time point [27]; minimize the viral population at each time point [1]; or incorporate the immune response by either maximizing the sum of helper T cells and effector cells ($CD8^+$ T cells) [11]; or minimizing the weighted difference between viral titre and effector cell counts [2]. In each of these target functions, the weights associated with costs or benefits have been estimated by considering the orders of magnitude of the quantities involved.

A third factor which may be weighed by the target function is the emergence of drug-resistant mutations. The effects of drug therapy on viral strains with differing levels of resistance are explored in [27] and [45]. These models, however, assume that any drug-resistant viral strains are present in the initial population, and thus the question of preventing the emergence of resistance *de novo* is not addressed.

A related approach for evaluating treatment regimens has been used in the analysis of structured treatment interruptions (STI). Here a precise schedule of periods "on" and "off" therapy is constructed. The underlying idea is that periods off therapy may boost the immune system while decreasing the drug burden; unfortunately such strategies may also facilitate drug resistance. For a recent review of clinical STI results, see [22].

Mathematical modelling was first used to compare three different STI strategies in the absence of drug resistance [25]. Later models such as [3,7] and [13]

were used to predict the effects of STI strategies with, for example, weekly or monthly interruptions in treatment, including the possibility of drug resistance. The question of whether immunological control of HIV could be established through STI is addressed in [26] and [46]; the clinical conditions under which STI might be optimal are explored in [44].

In Sect. 2 we propose a model of in-host HIV infection which includes drug therapy and the possibility of immune boosting. In Sect. 3 we develop a target function which balances the costs of side effects against the benefits of improved $CD4^+$ T cell counts, and imposes a severe cost on treatment regimens that facilitate the emergence of drug-resistance. An important contribution of this work is that we propose a quantitative means of weighing these effects against each other, rather than using order-of-magnitude estimates alone. Finally, in Sect. 4 we use this target function to predict the best possible treatment regimen, which would be possible if drug concentrations could be continuously adjusted, and to predict the best practical drug regimen, achievable with the known pharmacokinetics of currently available antiviral drugs.

2 Our model

In order to effectively model potential boosts to the immune system through periodic antigen exposure, we propose a system of differential equations which consists of five populations: $CD4^+$ T cells which are naïve x, activated r or productively infected y; and $CD8^+$ T cells which are naïve u or activated z. For each population we model the concentration of cells in a small volume of well mixed plasma as follows:

$$\dot{x} = \lambda_x - \beta(1-\eta)xy - \mathrm{d}_x x - \alpha xy \tag{1}$$

$$\dot{r} = \alpha b xy - \beta(1-\eta)ry - \mathrm{d}_r r \tag{2}$$

$$\dot{y} = \lambda_y + \beta(1-\eta)(r+x)y - \mathrm{d}_y y - \rho zy \tag{3}$$

$$\dot{u} = \lambda_u - \mathrm{d}_u u - \xi u r y \tag{4}$$

$$\dot{z} = \xi g u r y - \mathrm{d}_z z \tag{5}$$

Here we make a standard assumption that cell densities in plasma reflect the dynamics in the lymph.

Naïve $CD4^+$: Naïve $CD4^+$ T cells are produced by the thymus at constant rate λ_x and die at constant per capita rate d_x. They become infected at a rate proportional to the infected cell concentration, βy. The underlying assumption here is that free virus is roughly in equilibrium with the infected cell population [34]. This infection rate is reduced by the overall effectiveness of drug treatment to $\beta(1-\eta)$, with $0 \leq \eta \leq 1$. We discuss η further in the sections to follow, but note that we allow η to vary with time. Naïve $CD4^+$ T cells may also become activated at a rate proportional to the infected cell concentration, αy. Again, the underlying assumption is that the degree of antigen presentation is proportional to the number of infected cells.

Activated CD4$^+$: When activated, a CD4$^+$ T cell may clone itself; newly activated and newly cloned CD4$^+$ T cells are jointly captured through the term αbxy, where $(b-1)$ gives the number of clones produced by a newly activated cell. Activated CD4$^+$ T cells can become infected at rate $\beta(1-\eta)y$ and die with natural death rate d_r.

Infected CD4$^+$: Infected CD4$^+$ T cells are produced from either naïve or activated CD4$^+$ at rate $\beta(1-\eta)(r+x)$. An important simplifying assumption here is that naïve CD4$^+$ T cells, once infected, are productively infected. Infected CD4$^+$ T cells die at rate d_y with $\mathrm{d}_y > \mathrm{d}_r$, or are killed via mass action kinetics by activated CD8$^+$ T cells; ρ describes the effectiveness of this process. Our model does not explicitly include the latently infected T cell population. This limits the scope of our investigation as we cannot address the repopulation of the latent reservoir during STI; we discuss this further in Sect. 5. Our model does, however, implicitly include latently infected cells through the parameter λ_y. This term describes a continuous input of productively infected CD4$^+$ T cells from the latent reservoir [10,36].

Naïve CD8$^+$: Naïve CD8$^+$ T cells are produced by the thymus at constant rate λ_u and die at constant per capita rate d_u. We assume that CD8$^+$ T cells are activated at a rate proportional to the degree of antigen presentation, and thus proportional to the infected cell concentration y[1]. This process of activation, however, also requires an activated CD4$^+$ T cell. Thus, we model the overall rate of CD8$^+$ activation as ξury.

Activated CD8$^+$: Since newly activated CD8$^+$ T cells may also clone themselves, we let the term $\xi gury$ describe both newly activated and newly cloned CD8$^+$ T cells, with $g > 1$. Activated CD8$^+$ T cells die with natural death rate d_z.

Although we omit the details here, in the absence of drug therapy and for reasonable parameter values (see the following section), this system has two physically meaningful equilibria an unstable, uninfected equilibrium and a stable, infected equilibrium. The separation of naïve and activated cell populations in this model requires an extension to the class of well-studied and validated in-host HIV models [34]. These models predict the system behaviour near the uninfected or infected equilibria quite well, but are not typically valid for studying the decline of CD4$^+$ T cells over many years during chronic infection, nor for the course of primary infection. Since our study focusses on treatment intervals of less than 2.5 years, well after primary infection, neither of these limitations are relevant.

[1] Both activation and infection rates in this model are proportional to the degree of antigen presentation, whereas in reality these rates would saturate. However, for the effective treatment strategies we predict here, the degree of antigen presentation remains relatively low. Thus assuming a linear response allows for substantial simplification of the model, without loss of accuracy.

Table 1 Parameter estimates from the literature

Parameter	Value	Reference
d_x	$0.057\,\text{day}^{-1}$	[37]
d_y	$1\,\text{day}^{-1}$	[30]
d_u	$0.037\,\text{day}^{-1}$	[37]
λ_y	$0.0033\,\text{cells}/(\text{day}\,\mu\text{l})$	[10]
b	9	[31]
R_0	3.03	[17,18]

Table 2 Equilibrium cell populations from the literature

Uninfected equilibrium			
x_0	Naïve $CD4^+$ T cells	$1,066\,\text{cells}/\mu\text{l}$	[31]
u_0	Naïve $CD8^+$ T cells	$603\,\text{cells}/\mu\text{l}$	[31]
Infected equilibrium			
$x+r$	Total $CD4^+$ T cells	$388\,\text{cells}/\mu\text{l}$	[31]
$u+z$	Total $CD8^+$ T cells	$816\,\text{cells}/\mu\text{l}$	[31]
y	Infected $CD4^+$ T cells	$50\,\text{cells}/\mu\text{l}$	[28]
f_r	Fraction activated $CD4^+$	0.243	[37]
f_z	Fraction activated $CD8^+$	0.427	[37]

2.1 Model parameters

The model described above, of necessity, simplifies or neglects some potentially relevant aspects of the underlying immunology. Even so, this system includes a total of five populations and 14 free parameters. Fortunately, a number of important recent studies have provided experimental estimates of critical parameters for in-host HIV-1 infection [9,10,12,14,16–18,23,28–31,36,37,47].

Table 1 gives experimentally estimated values for five parameters of the model, along with an estimate for the basic reproductive ratio R_0. The estimate for λ_y is based on the rate of decline of the latent cell reservoir during aggressive drug therapy [10,36]; we assume that cells are primarily lost from this pool due to activation. Our estimate for the number of $CD4^+$ clones, b, agrees well with some experimental estimates [12,23], but a range of values have been reported in other studies [37,47].

In addition, experimental estimates are available for the cell concentrations at both the uninfected equilibrium (two populations) and the infected equilibrium (five populations). These estimates are provided in Table 2. At the infected equilibrium, the total $CD4^+$ and $CD8^+$ T cell concentrations have been estimated experimentally ($x+r$ and $u+z$, respectively), as have the fraction of these totals that are activated, f_r and f_z, respectively. Thus, the concentration of naïve $CD4^+$ T cells, for example, at the infected equilibrium is given by $(1-f_r)(x+r)$, while the concentration of activated $CD4^+$ T cells is $f_r(x+r)$.

These equilibrium values can then be used, along with the other "known" parameters, to algebraically determine unknown parameters. The seven parameters estimated in this way are provided in Table 3. Table 3 also gives the values chosen for the two remaining free parameters, d_r and d_z. We are not aware of experimental estimates for these parameters, therefore we assume that these

Table 3 Estimated parameter values

From the uninfected equilibrium	
λ_x	60.76 cells/(dayμl)
λ_u	22.31 cells/(dayμl)
From the infected equilibrium	
α	0.14×10^{-3} cells^{-1}day^{-1}
β	0.28×10^{-2} cells^{-1}day^{-1}
ρ	0.30×10^{-3} cells^{-1}day^{-1}
ξ	0.23×10^{-5} cells^{-2}day^{-1}
g	3.22
Free parameters	
d_r	0.065 day^{-1}
d_z	0.047 day^{-1}

death rates for activated cells are greater than or equal to the corresponding death rates for naïve cells, d_x and d_u, respectively.

Because of the uncertainty surrounding our estimates of b, d_r and d_z in particular, we performed sensitivity analysis by computing the partial derivative for each algebraically determined parameter (Table 3) with respect to changes in b, d_r and d_z. For example, the value of the CD4$^+$ T cell cloning rate, α, is determined from the infected equilibrium as:

$$\alpha = \frac{(\lambda_x - d_x x + d_r x)r}{xy(bx + r)}.$$

Therefore, the relative change of this value with a change in the death rate d_r is

$$\left(\frac{\partial \alpha}{\partial d_r}\right) \Big/ \left(\frac{\alpha}{d_r}\right) = \frac{d_r x}{\lambda_x - d_x x + d_r x}.$$

We then simply substitute the experimental estimate of the uninfected CD4$^+$ T cell density for x, and numerical estimates for the other parameters, to gauge the sensitivity of our estimate of α to estimates of d_r. Repeating for each parameter in Table 3, we verified that all parameter estimates were fairly insensitive to changes in these three values; no parameter would change by more than 10% for a 10% change in any of these values.

We likewise used partial derivatives to quantify the sensitivity of the estimated parameters to the values assumed for equilibrium populations (Table 2) and other measured parameters (Table 1). This analysis revealed that the parameters (α, β, ρ, ξ and g) were somewhat sensitive to changes in the concentration of infected cells at equilibrium. Since the infected cell concentration is only estimated to within an order of magnitude [9,28], we used the measured value of R_0 [16,18] to constrain our estimate of the equilibrium value of y.

Finally, we observed through this sensitivity analysis that ρ, the rate at which CD8$^+$ T cells kill infected cells, was very sensitive to many of the parameters estimated from the literature. For example, a 1% change in d_y, the death rate of infected cells, would result in a 9.5% change in our estimate of ρ. At the

conclusion of our study we therefore tested the sensitivity of our main results to variations in ρ; this is described in greater detail in Sect. 4.3.

3 The optimal treatment problem

3.1 Formulation of the target function

As described in Sect. 1 our goal is to formulate a target function that balances the costs and benefits of three main effects:

1. *Maximize helper T cells* Our main goal is to maintain the helper T cell population throughout therapy. Specifically, we wish to maximize the naïve $CD4^+$ T cell population, since these cells are critical for combating opportunistic infections. We consider a therapeutic interval from time t_0 to t_f and maximize the integral $\int_{t_0}^{t_f} x(t, \eta(t))\mathrm{d}t$.
2. *Minimize side effects* Side effects depend nonlinearly on the size of the drug intake, and we follow [15] and [24] in assuming that the cost of treatment is proportional to η^2, that is, we minimize a term $q_1 \int_{t_0}^{t_f} (\eta(t))^2 \mathrm{d}t$. The constant q_1 gives the weight with which we can balance the cost of side effects against the benefits of maintaining the $CD4^+$ T cell population. We address this further in Sect. 3.2.
3. *Minimize drug resistance* When drug therapy is intermittent, drug-resistant mutations may occur and virions carrying these mutations may replicate under positive selective pressure [38]. To minimize the probability that new drug-resistant mutations emerge during therapy, an appropriate term must be included in the target function; we justify our choice below.

We assume, conservatively, that the drug-resistant viral strain has complete resistance to the drug, but has a natural selective disadvantage σ. This allows us to compute R_m, the basic reproductive ratio of the drug-resistant viral strain, as the product of the rate at which a single infected cell produces "descendents", and its mean lifetime, $1/(\rho z + \mathrm{d}_y)$:

$$\mathrm{R}_m = \frac{\beta(1-\sigma)(x+r)}{\rho z + \mathrm{d}_y} \leq \frac{\beta(x+r)}{\mathrm{d}_y}. \tag{6}$$

We use this basic reproductive ratio to estimate the probability that each mutant lineage ultimately goes extinct due to stochastic fluctuations in an initially small population. If we assume that each cell infected by the drug-resistant viral strain has descendents regulated by linear birth–death process with mean R_m, then the probability that a single mutant lineage ultimately goes extinct is given by X, $X = \frac{1}{\mathrm{R}_m}$. We use $\pi = 1 - X$ to denote the probability that a single mutant lineage ultimately survives.

We therefore include a term in the target function which minimizes the following probability: that a cell infected by the drug-resistant viral strain is produced through mutation, and that this cell produces a lineage of secondarily

infected cells which ultimately does not go extinct. Since the rate at which such cells are produced by *de novo* mutation is given by $\mu\beta(1-\eta)(x+r)y$, we arrive at $q_2 \int_{t_0}^{t_f}(1-\eta)(x(t,\eta)+r(t,\eta))y(t,\eta)\pi(\eta,t)\mathrm{d}t$, where any constants which do not vary with t or η have been folded into q_2.

Thus the final form of our target function is:

$$T = \int_{t_0}^{t_f} x\mathrm{d}t - q_1 \int_{t_0}^{t_f} \eta^2 \mathrm{d}t - q_2 \int_{t_0}^{t_f} (1-\eta)(x+r)\, y\, \pi\, \mathrm{d}t \to \max_{\eta} \tag{7}$$

To use this target function in selecting optimal treatment regimens, we must balance the benefits of increased CD4$^+$ T cell counts against the costs of side effects and the risks of drug resistance. This is done by determining reasonable values for the weights, q_1 and q_2, as described in the following section.

3.2 Estimation of the weight of side effects

We first evaluate the balance between CD4$^+$ T cell counts and side effects, neglecting the risks of drug resistance. Since patients naturally differ in their tolerance to drugs, we assume that q_1 may take a range of possible values. Note that high drug tolerance implies a low weight q_1, and vice versa. Our estimates for high and low values of q_1 within the clinically relevant range are based on clinical practice: patients are typically introduced to antiviral therapy only after CD4$^+$ T cell counts have fallen to 350 cells/μl or less. We also use conservative estimates of treatment effectiveness, that is, we assume that the therapy prescribed is just sufficient to reduce the basic reproductive ratio of the drug-sensitive virus to a value less than one. We denote this level of drug effectiveness $\eta_{\min}$; for the parameters given, $\eta_{\min} \approx 0.67$.

Upper estimate: The clinical practice outlined above implies that when $x > 350$ cells/μl, the target function is maximized if the patient does not take therapy. Likewise, if $x < 350$ cells/μl, the target function is maximized by therapy. Thus when x is exactly 350 cells/μl, the target function should be equivalent in the two cases. If the patient takes drug therapy, the target function is equal to $\int \left(x_{\max} - q_1\eta_{\min}^2\right)\mathrm{d}t$, where $x_{\max}$ is the maximum T cell count attainable during aggressive therapy. We take $x_{\max}$ to be the uninfected equilibrium level of CD4$^+$ T cells. If the patient is not on therapy, the target function value is $\int 350\,\mathrm{d}t$. Equating the two, we find $q_1 = (x_{max} - 350)/(\eta_{min}^2)$.

Lower estimate: Alternatively, we can consider the benefit of delaying therapy until t_{350}, the time at which T cell counts reach 350 cells/μl. We argue that the side effects avoided by this delay exactly balance the loss of CD4$^+$ T cells experienced between t_0, the start of the infection, and t_{350}. Thus, $\int \left(x_{max} - q_1\eta_{\min}^2\right)\mathrm{d}t = \int x(t)\mathrm{d}t$, where $x(t)$ gives the time course of CD4$^+$ T cell counts in the absence of therapy, and the integrals are taken from t_0 to t_{350}. The pattern of loss in x over time varies markedly between patients, however, for the purposes of

estimating an approximate value of q_2 in the lower range of those clinically reasonable, we assume that the decrease in x is linear over time. This yields $q_1 = (x_{max} - 350)/(2\eta_{\min}^2)$.

3.3 Estimation of the weight of resistance mutations

We follow [6] in comparing two probabilities: the probability that a drug-resistant mutation is present in the infected cell population at the start of therapy, and the probability that such a mutation arises during therapy. Specifically, we find the level of drug effectiveness $\bar{\eta}$, at which these two probabilities are equal. If treatment is applied at this level, drug treatment does not increase the risk of resistance mutations.

The probability of initiating therapy when the mutant is present is given, when mutations are rare, by the ratio of average time that a mutant lineage survives, in the absence of the drug, over the average time between such mutations arising. Since new drug-resistant infected cells, denoted y_μ, arise at rate $\mu\beta(x+r)y$, the average time between such events is clearly $(\mu\beta(x+r)y)^{-1}$.

In the absence of drug therapy, the dynamics of a drug-resistant viral lineage can be approximated as $\dot{y_\mu} = \beta(1-\sigma)(x+r)y_\mu - \mathrm{d}_y y_\mu - \rho z y_\mu$. Thus y_μ decays exponentially, and the average time that a single lineage is present in the population is given by $(-\beta(1-\sigma)(x+r)+\mathrm{d}_y+\rho z)^{-1} = (\beta\sigma(x+r))^{-1}$ again assuming that the patient is at the infected equilibrium before therapy begins. We also assume that λ_y makes a negligible contribution to y compared to new infections through the β term; this is reasonable for patients at the infected equilibrium.

Therefore the probability of starting therapy when drug resistance already exists is given by the ratio of the average time that a single lineage exists, over the average time between such mutations arising, which simplifies directly to $y\mu/\sigma$. This value is less than one since μ is small and y is not too large in one μl. Estimates of the total population of infected cells in an infected individual range from 10^7 to 10^8 [9,28]. Assuming a point mutation rate of 10^{-5} to 10^{-4} [29] and selective disadvantage $\sigma = 10^{-3}$ to 10^{-2} [5], we find for example that the probability that a 3-point mutation already exists in the pool of infected cells at the start of the therapy is in the range 10^{-6} to 10^{-1}.

In contrast, we estimate the probability that a cell infected by a drug-resistant viral strain is produced by *de novo* mutation during treatment as:

$$\int_{t_0}^{t_f} \mu\beta(1-\eta)(x(t,\eta)+r(t,\eta))y(t,\eta)\mathrm{d}t \tag{8}$$

where we made a conservative assumption that $\pi(\eta,t) = 1$

Equating (8) with the expression $y\mu/\sigma$, we solve for $\bar{\eta}$, a treatment level which ensures that the probability of developing drug resistance is not increased during treatment. Since a large number of antiviral drugs are currently available, we assume a treatment time course, on the same drug combination, of two years,

such that $t_0 = 0$ and $t_f = 730$. This gives a numerical estimate of $\bar{\eta} = 0.659$. This estimate is fairly insensitive to the length of the treatment interval; for a four year interval for example we find $\bar{\eta} = 0.7$. Estimating $\bar{\eta}$ allows us to deduce a numerical estimate for q_2, however, we do not use this value to constrain our search for "optimal" $\eta(t)$, below, in any way.

To estimate the weight of resistance mutations, q_2, we evaluate the target function predicted for a constant drug effectiveness $\bar{\eta}$. We equate this value with the target function predicted in the complete absence of treatment; in this case we assume a worst-case scenario of $x = 200$ cells/μl, i.e., the development of AIDS. The underlying assumption is that the net benefit of drug-treatment at level $\bar{\eta}$ must at least be better than this worst-case scenario. We thus find an upper limit for the weight of resistance mutations by solving the following equation for q_2:

$$\int_{t_0}^{t_f} \bar{x}\mathrm{d}t - q_1 \int_{t_0}^{t_f} \bar{\eta}^2 \mathrm{d}t - q_2 \int_{t_0}^{t_f} \bar{y}(1-\bar{\eta})(\bar{x}+\bar{r})\pi(\bar{\eta},t)\mathrm{d}t = \int_{t_0}^{t_f} 200\mathrm{d}t \qquad (9)$$

Here we use $\bar{x}$, $\bar{r}$, and $\bar{y}$ to denote $x(t,\bar{\eta})$, $r(t,\bar{\eta})$, and $y(t,\bar{\eta})$.

In the sections that follow we consider two cases: high tolerance for the risk of drug resistance, $q_2 = 0$, and low tolerance, q_2 as determined by Eq. (9), such that the risk of drug resistance is not increased by treatment. High tolerance might correspond in practice to drug-naïve patients, who can expect to switch to a number of different drugs if resistance develops. Low tolerance might correspond to patients who have already developed resistance to some classes of antivirals and have more limited therapeutic options. Numerical integration is done by the 4th order Runge-Kutta method (GNU scientific library gsl_odeiv_step_rk4).

3.4 Modelling drug time courses

In modelling drug delivery, we consider two quite different approaches. First, we assume that the drug effectiveness can take any constant value between $0 \leq \eta \leq 1$ on any given day, and can be adjusted the following day to a new level. This case is of interest because it explores the theoretical possibility of boosting the immune system during therapy. We assume that cell populations are at the infected equilibrium at the beginning of treatment, and use a simulated annealing method (GNU scientific library `gsl_siman_solver`) to find 100 values of $\eta(t_i)$ over a 100 day treatment interval which are predicted to maximize the target function T. We thus predict the "best possible treatment" which could be obtained over this interval. Note that by using this numerical optimization over a 100-dimensional search space, we can by no means guarantee that a global maximum is found. To test the solutions obtained, we re-ran each optimization using nine different starting conditions; the starting conditions were constant values $\eta(t_i) = \eta_j$, $i = 1$–100, where η_j varied between 0.12 and 0.92. We also tested the degree to which the solution(s) thus obtained were

superior to constant drug therapy, and verified that the "best" regimens we obtained effectively controlled the infection.

Realistically, however, drugs are taken in fixed oral doses at fixed times per day (Table 4). Thus the drug effectiveness, $\eta(t)$, is predetermined to a large degree by pharmacokinetics. In a second approach, we thus constrain $\eta(t)$ using realistic pharmacokinetics for specific antiviral drugs (see Eq. 11 below). In addition, patients are unlikely to adhere to a complex regimen of continually changing doses. We thus additionally constrain the treatment regimen to consist of patterns of a days of full antiviral therapy, followed by a complete treatment break of b days. To predict the "best practical" treatment regimen, we then evaluate *all possible patterns* of a and b up to 100 days in total length. We exhaustively compared the performance of all such patterns repeated over a 1,000 day treatment interval to find the best pattern. Thus in this second case, we obtained a unique, global maximum for each set of pharmacokinetic parameters and weights.

Pharmacokinetics: Table 4 gives relevant pharmacokinetic parameters for three protease inhibitors (PIs) and three non-nucleoside reverse transcriptase inhibitors (NNRTIs); these parameters are from the following sources: [4,19,32,33,35,39,41,42].

Following [43], given the recommended dosing interval τ, time to peak $T_{\max}$, and serum half-life $t_{1/2}$, the concentration of drug is assumed to grow linearly to its maximum concentration $C_{\max}$ during the time to peak, and decay exponentially thereafter:

$$C(t) = \begin{cases} C(0) + \frac{t}{T_{\max}}(C_{\max} - C(0)) & \text{if } 0 \le t \le T_{\max} \\ C_{\max} e^{-w(t-T_{\max})} & \text{if } T_{\max} \le t \le \tau \end{cases} \tag{10}$$

where $w = log(2)/t_{1/2}$.

Table 4 Pharmacokinetic parameters: after a single dose the drug concentration increases to C_{max} in time-to-peak T_{max}, then decays exponentially with half-life $t_{1/2}$

Drug	Dosage	Protein binding(%)	C_{max} μM	T_{max} (h)	$t_{1/2}$ (h)	In vitro (IC_{50} μM)
PIs						
Amprenavir	1,200 mg 12 h	90	20.7	1–2	7.1–10.6	0.012–0.08
Indinavir	800 mg 8 h	60	12.63	0.8	1.5	0.025–0.1
Ritonavir	600 mg 12 h	98–99	11.2	2–4	3–5	0.0038–0.154
NNRTIs						
Delavirdine	400 mg 12 h	98	35	1.3	3.3–8.3	0.001–0.69
Efavirenz	600 mg 24 h	99.5–99.75	12.9	3–5	52–76	0.001–0.11
Nevirapine	200 mg 24 h	60	5.58	4	16.5	0.01–0.1

IC_{50} is the concentration required to inhibit viral replication by 50%

Drug effectiveness depends on drug concentration and the IC_{50} value:

$$\eta(t) = \frac{C(t)}{C(t) + IC_{50}} \tag{11}$$

where IC_{50} is the *intracellular* concentration of drug that is necessary to inhibit the viral replication by 50%. We estimate this value by taking into consideration protein binding of the drug in plasma. Specifically, we estimate the intracellular IC_{50} by dividing in vitro estimates for the IC_{50} (Table 4) by the unbound fraction [4], the fraction of the dose that is not bound to protein in plasma. Also note that in Eq. 7 we have assumed that the cost of side effects is a function of drug effectiveness, η, not drug concentration. The underlying assumption here is that side effects largely scale with the dose to tissues, not with drug bound to plasma proteins.

We note that drug IC_{50} values, unlike the other parameters in Table 4, may range over several orders of magnitude. In the results to follow we therefore consider both upper and lower estimates of the intracellular IC_{50}, which correspond to lower and higher drug effectiveness, respectively. Figure 6 illustrates the time course of η after a single dose for each of the six drugs in Table 4. This figure illustrates the vastly different patterns of drug effectiveness achievable with different drugs.

4 Results

4.1 Best possible treatment

We estimated the "best" possible treatment for six cases, when the tolerance for mutation was either high or low, and when tolerance for side effects was high, intermediate or low. We again note that unique optima cannot be guaranteed, and in fact for the nine different starting conditions, several distinct local optima were often obtained. In particular, very low starting values of $\eta(t_i)$ often converged to suboptimal solutions. In four of the six parameter cases, the five highest starting conditions converged to solutions for which each value of $\eta(t_i)$ differed by less than 0.004; effectively the same solution was reached. This solution also yielded a higher value of the target function than any other predicted solution for these parameter sets. In the two remaining cases ($q_2 = 0$, $q_1 = 1{,}194$, or $1{,}592$), different starting conditions yielded solutions that were distinguishable by eye, but qualitatively identical in every case. We are thus confident that the predicted solutions, if not strictly optimal, are very good solutions to this problem.

Figure 1 shows the optimal treatment strategies predicted by our model when tolerance for the risk of mutations is high, i.e., $q_2 = 0$. We explore three cases: high, intermediate and low tolerance for side effects. Although the optimal drug regimen differs between these three cases, the initial strategy is common to all cases. We note that "boosting" of the immune system is not only possible in this

system, it is optimal when the tolerance for side effects is relatively low. This occurs when the drug regimen is periodically reduced (panels (c) and (e)) such that the immune system is stimulated (note the time course of z in panels (d) and (f)). We note that these "boosting" strategies emerged from our simulated annealing method without any *a priori* assumptions; we searched for a global optimum starting from a constant drug regimen.

The best possible treatment strategy changes if the tolerance for the risk of drug resistance is low, as illustrated in Fig. 2. In this situation the tolerance for side effects has little impact on the optimal strategy. In none of these cases do we observe a reduction in the prescribed drug regimen at later times.

We also note that the optimal level of drug effectiveness is generally less than $\eta_{\min}$, the level required to reduce the value of the basic reproductive ratio R_0 to one. Each of these intriguing results will be discussed further in Sect. 5.

4.2 Best practical treatment

Table 5 gives the best practical treatment strategy for each of the six drugs, depending on the tolerance for side effects or the risk of mutations, and the IC_{50} value assumed. These values represent unique global optima obtained through an exhaustive search. Note that horizontal lines indicate that the best possible strategy is not to take the drug at all. Thus if patient tolerance for side effects is low, treatment with only three of the six drugs would be recommended. For each of these three drugs (ritonavir, delavirdine and efavirenz), protein binding in plasma is very high (see Table 4), and the effective dose to tissues is consequently reduced, particularly when the IC_{50} value is high (Fig. 6, solid lines).

This effect is further elucidated in Fig. 3. This figure plots the overall benefit of therapy versus the weight of side effects, q_1. The benefit is computed by finding the maximum possible value of the target function attainable with any combination of a and b, integrated over the 1,000 day treatment period. The value of the target function in the absence of drug therapy is then subtracted; both values are normalized per day of treatment. Thus negative values indicate that it is better not to treat the patient with that particular drug. Here we see that the drugs which are most effective for patients with high tolerance for side effects also tend to be most sensitive to the drug tolerance; when drug tolerance is very low, these drugs are no longer recommended.

Most of the optimal strategies reported in Table 5 fall into two broad classes. In several cases, a period of drug therapy lasting 2 weeks to a month is followed by a drug holiday of about 1 week. An example of the cell dynamics and drug time course for such a regimen is illustrated in Fig. 4. In these cases, populations of activated $CD4^+$ and $CD8^+$ T cells are maintained at high levels, and are clearly boosted during drug holidays, when the infected cell population increases.

For most other cases, a short period of one to three days of therapy is followed by a similarly short holiday. Figure 5 illustrates the cell dynamics and the

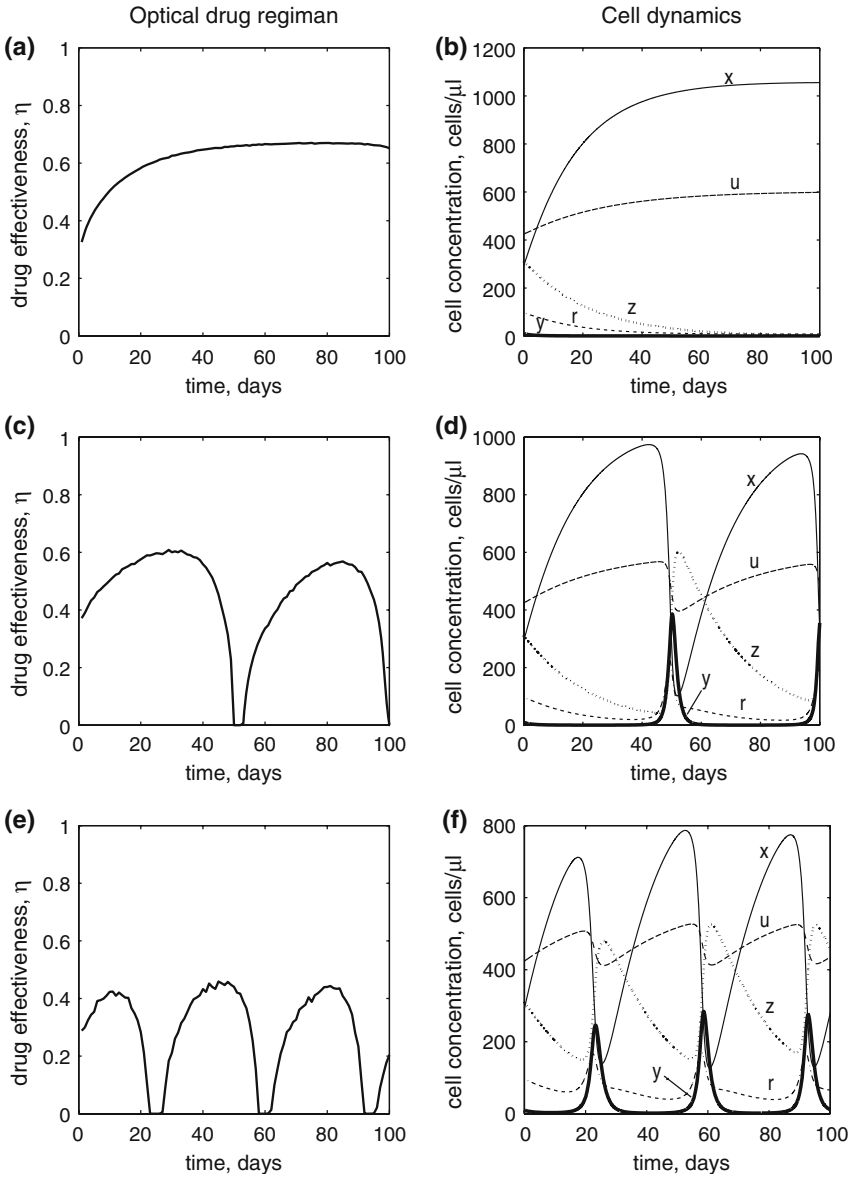

Fig. 1 The best possible drug regimen when the tolerance for drug-resistant mutations is high. The tolerance for side effects is either high (*top row*), intermediate (*middle row*), or low (*bottom row*). Note that when tolerance for side effects is intermediate or low, the best possible regimen includes periods of reduced drug exposure and consequent immune boosting. A numerical solution to Eq. (7) subject to the cell dynamics described by Eqs. (1–5) was determined by simulated annealing for $q_1 = 796$, $q_2 = 0$ (*top row*), $q_1 = 1{,}194$, $q_2 = 0$, (*middle row*), $q_1 = 1{,}594$, $q_2 = 0$ (*bottom row*). The annealing method converged to solutions very similar to those illustrated independent of starting values; for results shown here we started with a constant drug regiment $\eta(t) = 0.52$. *Panels* at left give drug effectiveness $\eta(t)$ versus time. Plots at right give the corresponding cell dynamics. Parameter values as provided in Tables 1, 2 and 3

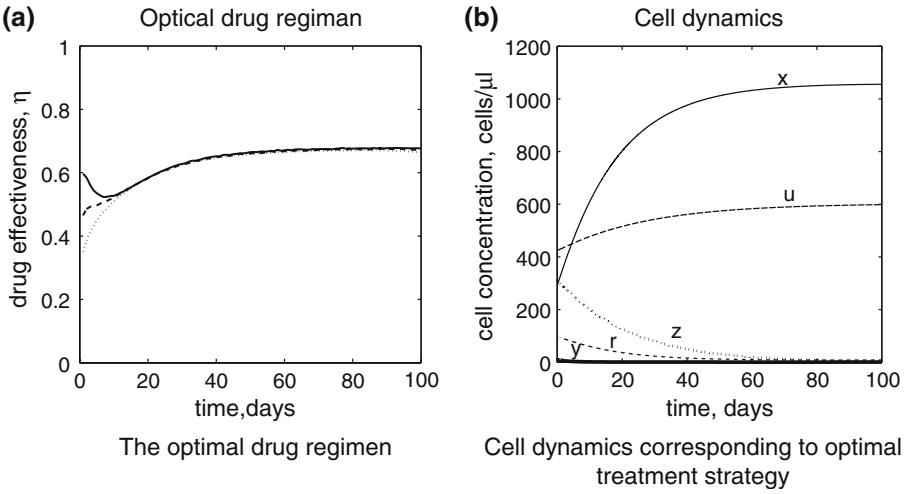

Fig. 2 The best possible drug regimen when the tolerance for drug-resistant mutations is low. Tolerance for side effects is either high, $q_1 = 796$, $q_2 = 0.18$, (*solid line*), intermediate, $q_1 = 1,194$, $q_2 = 0.10$, (*dashed line*), or low $q_1 = 1,592$, $q_2 = 0.02$, (*dotted line*). Note that these three cases are virtually indistinguishable. Numerical solution of the maximization problem (7) was determined as for Fig. 1. In *panel* 2(b) cell dynamics are shown for the case of high drug tolerance

drug time course for one such case. In this situation the populations of activated $CD4^+$ and $CD8^+$ are very small at all times, and the infected cell population is maintained at a low level throughout therapy.

4.3 Sensitivity analysis

As stated previously, analytical sensitivity analysis suggested that the parameter ρ, the rate at which $CD8^+$ T cells kill infected cells, was uniquely sensitive to the estimates for other parameters and equilibrium values taken from the literature. We therefore recomputed both the best possible and best practical treatment strategies for a 20% increase or decrease in ρ, keeping all other parameters constant.

We found that our results were almost completely insensitive to changes of this magnitude in ρ. The best possible strategies were almost indistinguishable from those illustrated in Figs. 1 and 2. For the best practical treatment regimens, of the 96 entries in Table 5 (including the horizontal lines when treatment is not beneficial), 2 entries differed. Thus, although there is a large degree of uncertainty surrounding our estimate of ρ, our main conclusions are not sensitive to the value assumed.

5 Discussion

We have carefully weighed the main costs of antiviral treatment against its benefits, predicting optimal therapeutic strategies and systematically evaluating

Table 5 Best practical treatment strategies of a days on drug therapy, followed by b days off therapy

Tolerance for drug-resistant mutations:		High		High		Low		Low	
Tolerance for side effects		High		Low		High		Low	
Drug	IC_{50}	a	b	a	b	a	b	a	b
Amprenavir	High	1	2	–	–	1	2	–	–
	Low	2	4	–	–	2	4	–	–
Indinavir	High	3	1	–	–	3	1	–	–
	Low	2	1	–	–	2	1	–	–
Ritonavir	High	25	6	22	9	–	–	–	–
	Low	3	2	–	–	1	1	–	–
Delavirdine	High	26	6	22	9	–	–	–	–
	Low	1	3	–	–	1	3	–	–
Efavirenz	High	16	9	13	12	–	–	–	–
	Low	1	19	–	–	1	18	–	–
Nevirapine	High	2	4	–	–	1	3	–	–
	Low	1	7	–	–	1	7	–	–

A horizontal line indicates that it is best not to treat in this case

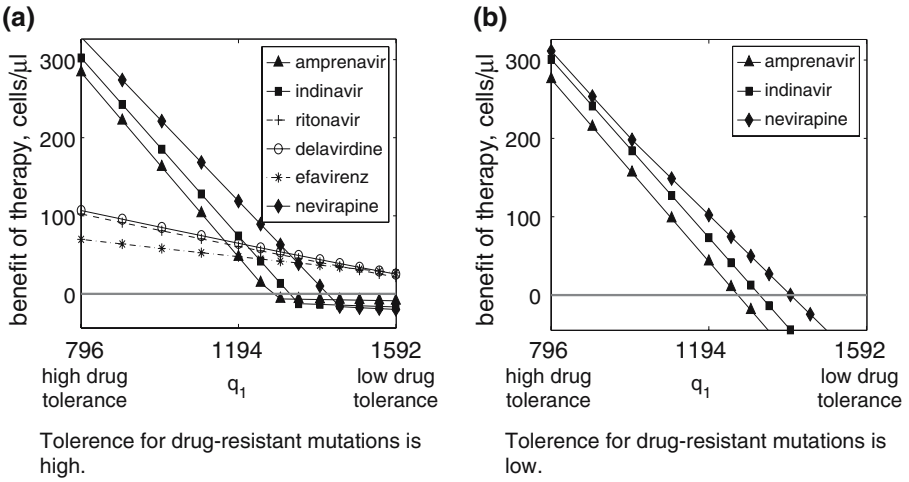

Fig. 3 Net Benefit of therapy versus tolerance for side effects. The maximum possible benefit of therapy was determined (see text for details) for high ($q_2 = 0$, *panel* 3(a)), and low (q_2 from Eq. (9), *panel* 3(b)) tolerances for drug-resistant mutations. Note that values below zero indicate that no treatment is preferred. *Panel* 3(b) shows only those drugs for which treatment is recommended at some value of q_1. The treatment pattern is a fixed number of "on" days followed by a fixed number of "off" days, as described in Table 5. Parameter values as provided in Tables 1, 2, 3, and 4

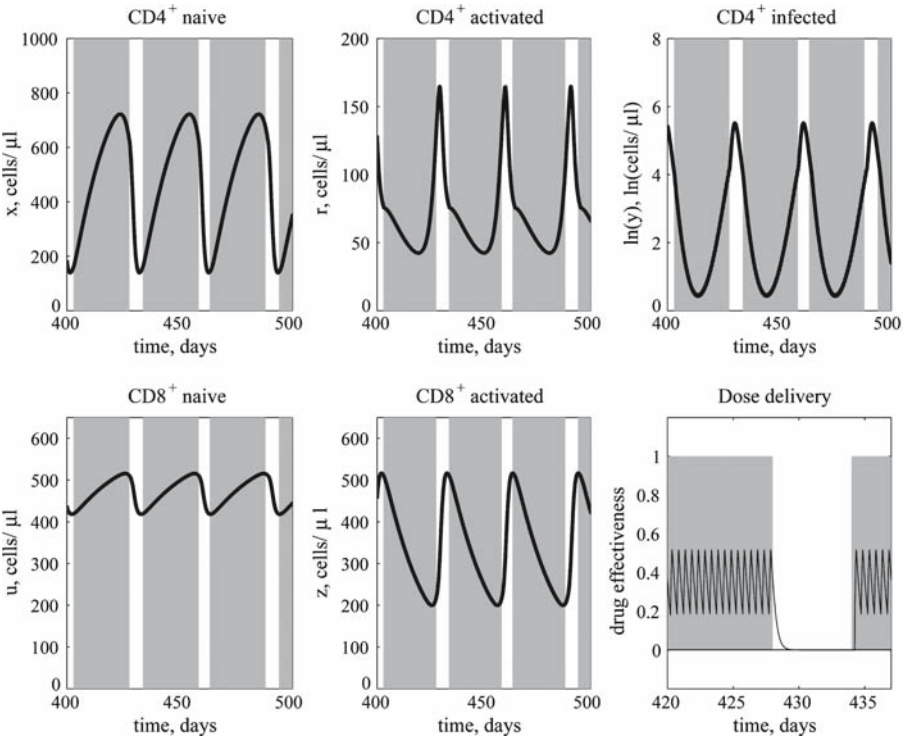

Fig. 4 Cell dynamics during the best practical treatment regimen: ritonavir, high IC_{50}. Note that breaks in treatment boost the immune system. Equation (7) was maximized by comparing all possible repeating combinations of a days of therapy followed by a break of b days, as described in the text. Shaded regions correspond to blocks of therapy. Parameter values as provided in Tables 1, 2, 3, and 4. In this example tolerances for both side effects and drug-resistant mutations are high

STI regimens. In the latter evaluation, we use parameters specific to currently prescribed antiviral drugs. Nonetheless, the main purpose of this study is not to make specific predictions for therapeutic use, but rather to draw more generally applicable conclusions. We outline the main predictions of our model below.

We repeatedly find that non-constant treatment strategies are superior to constant treatment regimens. In particular, *drug regimens which boost the immune system are often best*. These strategies reduce drug exposure, subsequently increasing the population of infected cells, and therefore stimulate the immune response. In such cases, the drug works in concert with the immune response (Fig. 4). The best drugs to use in this case are those with low dose to tissues and longer half-lives (see Fig. 6).

We also predict that *the lower the patient's tolerance for drugs, the higher the benefit of boosting the immune system*. This is not surprising and was common to both the best possible and best practical treatment cases. Conversely, *the lower the tolerance for the risk of resistance mutations, the lower the benefit of boosting the immune system*. In practice, low tolerance might correspond to patients

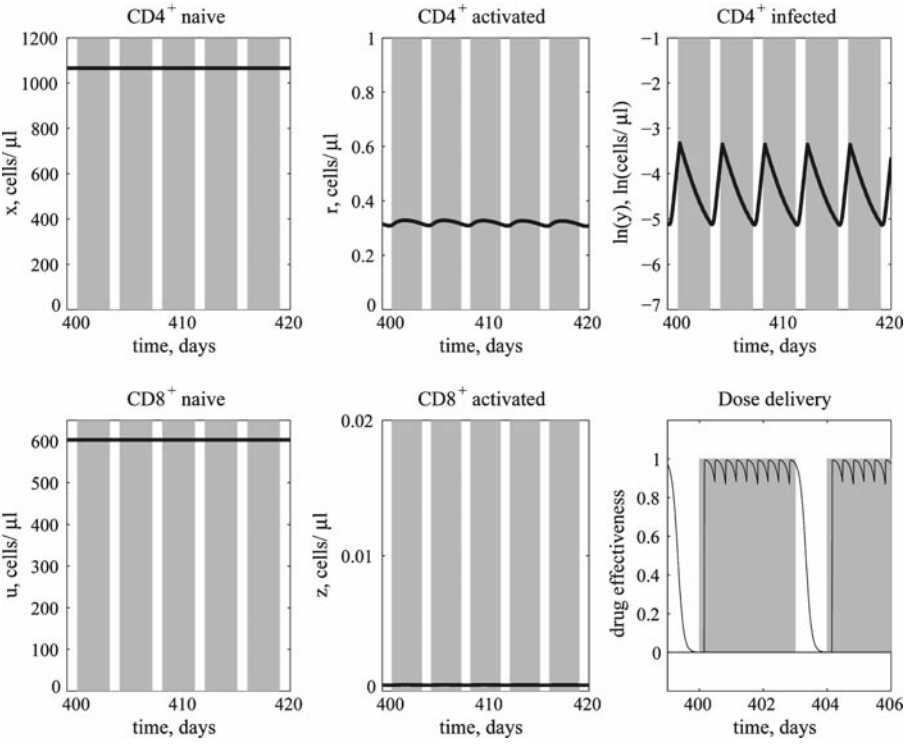

Fig. 5 Cell dynamics during the best practical treatment regimen: indinavir, high IC_{50}. Note that the pool of infected cells is maintained at low levels at all times, thus activated $CD4^+$ and $CD8^+$ T cell populations are also low. See legend to Fig. 4 for details

who are unable to frequently switch to new drugs, for instance, patients with limited therapeutic options because of previous exposure to drug-resistance. Drug regimens which boost the immune system are not recommended in such cases.

When boosting the immune system is not recommended, the best possible strategy is to maintain the infected cell population at the lowest possible level (Fig. 5). This has the adverse effect of eliminating immunological pressure. In such cases the best possible drugs are those with sufficient effectiveness to reduce R_0 below one, and the longest possible half-life (see Fig. 6). Greater effectiveness only increases the cost of side effects; shorter half-lives make it difficult to maintain a continually high dose.

When drug effectiveness can be continually adjusted to any desired level, our model predicts that starting from the infected equilibrium, *the best treatment regimen always begins with a gradual increase in drug levels over the first few weeks of therapy*. Thus as antigen exposure decreases during drug therapy, the declining immune response is gradually replaced by drugs. Interestingly, this corresponds to current clinical practice, in that lower doses are typically prescribed when therapy is initiated, in order to limit side effects. We also find that

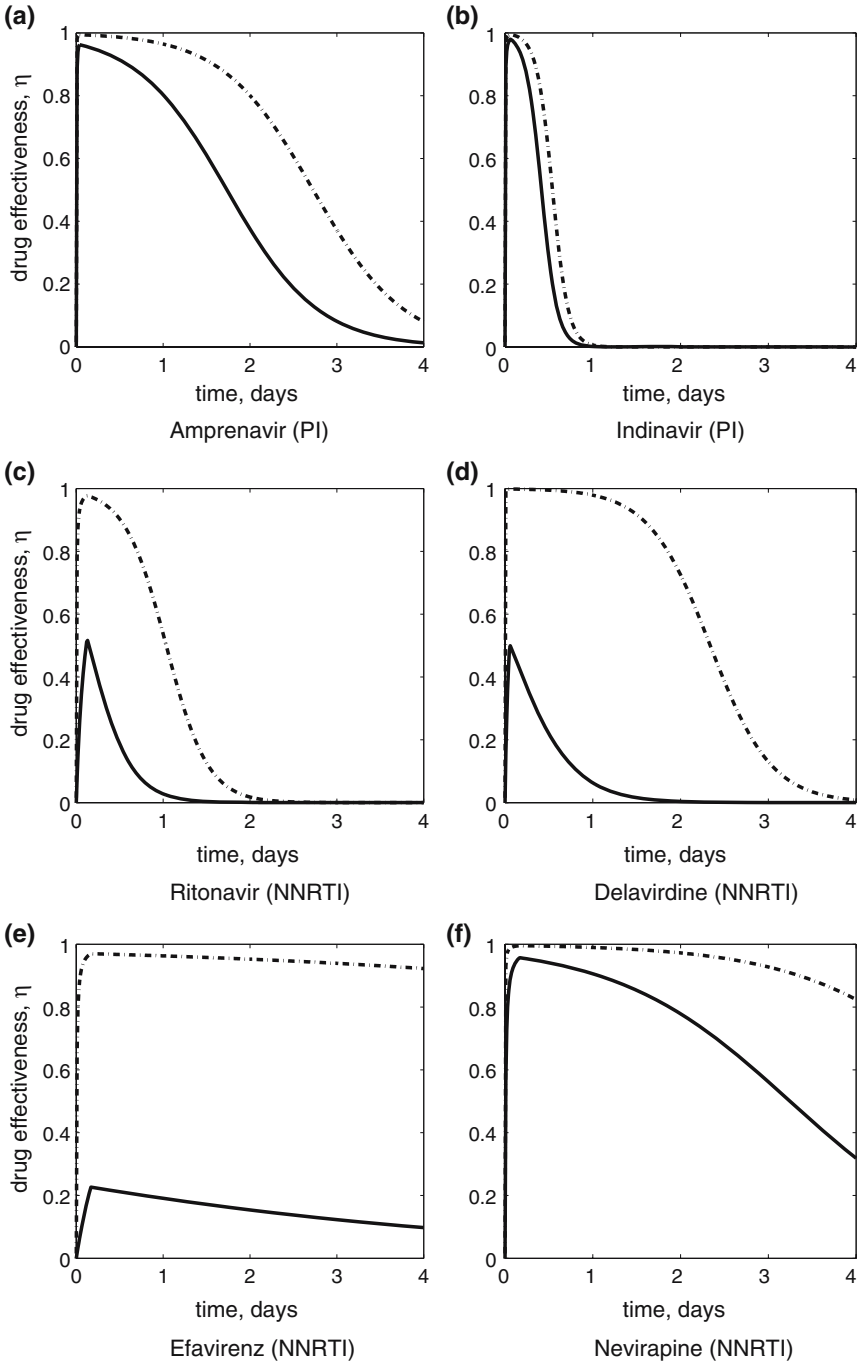

Fig. 6 Time course of drug effectiveness after a single dose. The *dotted line* corresponds to the lowest estimated value of the IC_{50} and the *solid line* to the highest value for each drug, as provided in Table 4. Dose effectiveness is determined by Eq. (11), with the IC_{50} value adjusted for protein binding

the optimal therapy in this case involves lower drug concentrations than would be required to set $R_0 < 1$, since these regimens rely on both the drug and the immune system to contain the infection.

We find that *it is not possible with every drug to devise a beneficial treatment plan when limited to full days "on" and "off" therapy*. This is because the drug concentration decays during the drug holiday, exposing the patient to intermediate drug levels which greatly increase the probability of developing drug-resistance.

The approach reported here is limited in a number of important ways. As a first avenue for future work, we plan to extend this model to include double and triple drug therapy. Also, it is important to note that we have assumed perfect adherence to therapy; in computing the target function, doses are never missed. Extending our model for imperfect adherence is an obvious and necessary extension.

Acknowledgments This work is supported by the Natural Sciences and Engineering Research Council of Canada, the Ontario Ministry of Science, Technology and Industry and by the SHARCNET parallel computing facility. We are indebted to Jane Heffernan for many helpful discussions, to Tyson Whitehead for his invaluable technical expertise, and to an anonymous referee whose comments substantially strengthened the paper.

References

1. Adams, B.M., Banks, H.T., Davidian, M., Kwon, H.D., Tran, H.T., Wynne, S.N.: HIV dynamics: modeling, data analysis, and optimal treatment protocols. J. Comput. Appl. Math. **184**(1), 10–49 (2004)
2. Adams, B.M., Banks, H.T., Tran, H.T., Kwon, H.: Dynamic multidrug therapies for HIV: optimal and STI control approaches. Math. Biosci. Eng. **1**, 223–241 (2004)
3. Bajaria, S.H., Webb, G., Kirschner, D.E.: Predicting differential responses to structured treatment interruptions during HAART. Bull. Math. Biol. **66**(5), 1093–1118 (2004)
4. Boffito, M., Back, D.J., Blaschke, T.F., Rowland, M., Bertz, R.J., Gerber, J.G., Miller, V.: Protein binding in antiretroviral therapies. AIDS Res. Hum. Retroviruses **19**(9), 825–835 (2003)
5. Bonhoeffer, S., May, R.M., Shaw, G.M., Nowak, M.A.: Virus dynamics and drug therapy. Proc. Natl. Acad. Sci. USA **94**, 6971–6976 (1997)
6. Bonhoeffer, S., Nowak, M.A.: Pre-existence and emergence of drug resistance in HIV-1 infection. Proc. R. Soc. Lond. B Bio. **264**(1382), 631–637 (1997)
7. Bonhoeffer, S., Rembiszewski, M., Ortiz, G.M., Nixon, D.F.: Risks and benefits of structured antiretroviral drug therapy interruptions in HIV-1 infection. AIDS **14**(15), 2313–2322 (2000)
8. Butler, S., Kirschner, D., Lenhart, S.: Optimal control of chemotherapy affecting the infectivity of HIV. In: Arino, O., Axelrod, D., Kimmel, M., Langlais, M. (eds.) Advances in Mathematical Population Dynamics: Molecules, Cells, Man, pp. 104–120. World Scientific Publishing, (1997)
9. Cavert, W., Notermans, D.W., Staskus, K., Wietgrefe, S.W., Zupancic, M., Gebhard, K., Henry, K., Zhang, Z.Q., Mills, R., McDade, H., Schuwirth, C.M., Goudsmit, J., Danner, S.A. Haase, A.T.: Kinetics of response in lymphoid tissues to antiretroviral therapy of HIV-1 infection. Science **276**(5314), 960–964 (1997)
10. Chun, T.W., Fauci, A.S.: Latent reservoirs of HIV: obstacles to the eradication of virus. Proc. Natl. Acad. Sci. USA **96**(20), 10958–10961 (1999)
11. Culshaw, R.V., Ruan, S., Spiteri, R.J.: Optimal HIV treatment by maximising immune response. J. Math. Biol. **48**(5), 545–562 (2004)
12. De Boer, R.J., Homann, D., Perelson, A.S.: Different dynamics of $CD4^+$ and $CD8^+$ T cell responses during and after acute lymphocytic choriomeningitis virus infection. J. Immunol. **171**(8), 3928–3935 (2003)

13. Dorman, K.S., Kaplan, A.H., Lange, K., Sinsheimer, J.S.: Mutation takes no vacation: can structured treatment interruptions increase the risk of drug-resistant HIV-1?. J. Acq. Imm. Defic. Syndr. **25**(5), 398–402 (2000)
14. Finzi, D., Blankson, J., Siliciano, J.D., Margolick, J.B., Chadwick, K., Pierson, T., Smith, K., Lisziewicz, J., Lori, F., Flexner, C., Quinn, T.C., Chaisson, R.E., Rosenberg, E., Walker, B., Gange, S., Gallant, J., Siliciano, R.F.: Latent infection of $CD4^+$ T cells provides a mechanism for lifelong persistence of HIV-1, even in patients on effective combination therapy. Nat. Med. **5**, 512–517 (1999)
15. Fister, K.R., Lenhart, S., McNally, J.S.: Otimizing chemotherapy in an HIV model. Electr. J. Diff. Eq. **32**, 1–12 (1998)
16. Frost, S.D.: Dynamics and evolution of HIV-1 during structured treatment interruptions. AIDS Rev. **4**(3), 119–127 (2002)
17. Frost, S.D., Martinez-Picado, J., Ruiz, L., Clotet, B., Brown, A.J.: Viral dynamics during structured treatment interruptions of chronic Human Immunodeficiency Virus type 1 infection. J. Virol. **76**(3), 968–987 (2002)
18. Funk, G.A., Fischer, M., Joos, B., Opravil, M., Gunthard, H.F., Ledergerber, B., Bonhoeffer, S.: 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)
19. HIV ePharmacotherapy Network Home page. http://hiv.buffalo.edu. Cited 22, 2005
20. Jeffrey, A.M., Xia, X., Craig, J.K.: When to initiate HIV therapy: a control theoretic approach. IEEE Trans. Biomed. Eng. **50**(11), 1213–1219 (2003)
21. Joshi, H.R.: Optimal control of an HIV immunology model. Optim. Contr. Appl. Math. **4**, 199–213 (2002)
22. Julg, B., Goebel, F.D.: Treatment interruption in HIV therapy: a SMART strategy? Infection **34**(3):186–188 (2006)
23. Kaech, S.M., Wherry, E.J., Ahmed, R.: Effector and memory T cell differentiation: implications for vaccine development. Nat. Rev. Immunol. **2**(4), 251–262 (2002)
24. Kirschner, D., Lenhart, S., Serbin, S.: Optimizing chemotherapy of HIV infection: scheduling, amounts and initiation of treatment. J. Math. Biol. **35**, 775–792 (1997)
25. Kirschner, D., Webb, G.F.: A model for treatment strategy in the chemotherapy of AIDS. Bull. Math. Biol. **58**(2), 367–390 (1996)
26. Komarova, N.L., Barnes, E., Klenerman, P., Wodarz, D.: Boosting immunity by antiviral drug therapy: a simple relationship among timing, efficacy, and success. Proc. Natl. Acad. Sci. USA **100**(4), 1855–1860 (2003)
27. Kutch, J.J., Gurfil, P.: Optimal control of HIV infection with a continuously-mutating viral population. In: Proceedings of American Control Conference, pp. 4033–4038 (2002)
28. Levy, J.A., Ramachandran, B., Barker, E., Guthrie, J., Elbeik, T.: Plasma viral load, $CD4^+$ cell counts, and HIV-1 production by cells. Science **271**(5249), 670–671 (1996)
29. Mansky, L.M., Temin, H.M.: 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)
30. Markowitz, M., Louie, M., Hurley, A., Sun, E., Di Mascio, M., Perelson, A.S., Ho, D.D.: 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**, 5037–5038 (2003)
31. Mohri, H., Perelson, A.S., Tung, K., Ribeiro, R.M., Ramratnam, B., Markowitz, M., Kost, R., Hurley, A., Weinberger, L., Cesar, D., Hellerstein, M.K., Ho, D.D.: Increased turnover of T lymphocytes in HIV-1 infection and its reduction by antiretroviral therapy. J. Exp. Med. **194**(9), 1277–1287 (2001)
32. Murphy, R.L., Sommadossi, J.P., Lamson, M., Hall, D.B., Myers, M., Dusek, A.: Antiviral effect and pharmacokinetic interaction between nevirapine and indinavir in persons infected with Human Immunodeficiency Virus type 1. J. Infect. Dis. **179**(5), 1116–1123 (1999)
33. NIAD Division of AIDS (DAIDS) Home page. http://www.niaid.nih.gov/daids/. Cited 22, 2005
34. Nowak, M.A., May, R.M.: Virus Dynamics: Mathematical Principles of Immunology and Virology, Chap. 10. Oxford University Press, Oxford, New York (2000)
35. Pfizer Inc. Home page. http://www.pfizer.ca/. Cited 22, 2005

36. Ramratnam, B., Mittler, J.E., Zhang, L., Boden, D., Hurley, A., Fang, F., Macken, C.A., Perelson, A.S., Markowitz, M., Ho, D.D.: The decay of the latent reservoir of replication-competent HIV-1 is inversely correlated with the extent of residual viral replication during prolonged anti-retroviral therapy. Nat. Med. **6**(1), 82–85 (2000)
37. Ribeiro, R.M., Mohri, H., Ho, D.D., Perelson, A.S.: 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)
38. Richman, D.D.: The implications of drug resistance for strategies of combination antiviral chemotherapy. Antiviral Res. **29**(1), 31–33 (1996)
39. RxList, The Internet Drug Index Home page. http://www.rxlist.com/. Cited 15 October 2005
40. Stengel, R.F., Ghigliazza, R.: Stochastic optimal therapy for enhanced immune response. Math. Biosci. **191**, 123–142 (2004)
41. The Immunodeficiency Clinic—University Health Network Home Page. http://www.tthhiv-clinic.com/. Cited 15 October 2005
42. van Heeswijk, R.P., Veldkamp, A.I., Mulder, J.W., Meenhorst, P.L., Wit, F.W., Lange, J.M., Danner, S.A., Foudraine, N.A., Kwakkelstein, M.O., Reiss, P., Beijnen, J.H., Hoetelmans, R.M.: The steady-state pharmacokinetics of nevirapine during once daily and twice daily dosing in HIV-1 infected individuals. AIDS **14**(8), F77–F82 (2000)
43. Wahl, L.M., Nowak, M.A.: Adherence and resistance: Predictions for therapy outcome. Proc. Biol. Sci. **267**(1445), 835–843 (2000)
44. Walensky, R.P., Goldie, S.J., Sax, P.E., Weinstein, M.C., Paltiel, A.D., Kimmel, A.D., Seage, G.R., Losina, E., Zhang, H., Islam, R., Freedberg, K.A.: Treatment for primary HIV infection: projecting outcomes of immediate, interrupted, or delayed therapy. J. Acquir. Immune. Defic. Syndr. **31**, 27–37 (2002)
45. Wein, L.M., Zenios, S.A., Nowak, M.A.: Dynamic multidrug therapies for HIV: a control-theoretic approach. J. Theor. Biol. **185**, 15–29 (1997)
46. Wodarz, D., Page, K.M., Arnaout, R.A., Thomsen, A.R., Lifson, J.D., Nowak, M.A.: A new theory of cytotoxic T-lymphocyte memory: implications for HIV treatment. Philos. Trans. R. Soc. Lond. B Biol. Sci. **355**(1395), 329–343 (2000)
47. Zand, M.S., Briggs, B.J., Bose, A., Vo, T.: Discrete event modeling of $CD4^+$ memory T cell generation. J. Immunol. **173**(6), 3763–3772 (2004)