Efficacy of Antibody and T cell Therapies for Highly Mutable Viruses like Human Immunodeficiency Virus

by

Pranav M. Murugan

S.B., Electrical Engineering and Computer Science and Physics,
Massachusetts Institute of Technology (2022)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Authored by:	Pranav M. Murugan Department of Electrical Engineering and Computer Science May 19, 2023
Certified by:	Arup K. Chakraborty Department of Chemical Engineering Thesis Supervisor
Accepted by:	Katrina LaCurts Chair, Master of Engineering Thesis Committee

Efficacy of Antibody and T cell Therapies for Highly Mutable Viruses like Human Immunodeficiency Virus

by

Pranav M. Murugan

Submitted to the Department of Electrical Engineering and Computer Science
on May 19, 2023, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

The isolation of broadly neutralizing antibodies (bnAbs) that can neutralize diverse strains of highly mutable viruses like human immunodeficiency virus (HIV) as well as identification of mutationally-constrained regions of the proteome that could be targeted by T cells has led to interest in passive immunotherapies and therapeutic vaccines as promising methods for treating chronic infection. However, the feasibility of creating a sufficiently powerful therapy remains uncertain. In this work, we develop a stochastic computational model of viral dynamics to help characterize the regimes where viral control or cure may be possible. We study the efficacy of either bnAb therapy or therapeutic vaccination that elicits T cell responses that target mutationally-constrained regions, as well as treatments that combine these two therapeutic modalities. Our results show that combination therapy has the best chance of maintaining viral control or achieving a cure. this is because administering combinations of bnAbs with broad coverage of viral strains for a sufficiently long time can potentially clear rare strains from the latent reservoir which are likely to escape T cell responses resulting in viral rebound. We also describe a strong relation between the outcome of treatment and the diversity of the reservoir of latently infected cells, which suggest that the best candidates for immunotherapy are those who started antiretroviral therapy shortly after infection. Importantly, we find that cure is likely to be a rare outcome, and that the average time to cure is long and independent of therapeutic modality as it depends on the rate of activation of the latent reservoir. Our results will help guide the design of new therapeutics, and provide a platform for future computational screening of of the efficacy of new treatment regimes.

Thesis Supervisor: Arup K. Chakraborty
Title: Institute Professor

Acknowledgements

I would like to thank Arup Chakraborty for his tireless guidance and mentorship since the beginning of my undergraduate studies. I am grateful for many scientific and non-scientific conversations with him and other members of the AKC group. I would also like to acknowledge the tremendous help from all of our collaborators, who helped shape the questions explored in this thesis. Finally, I want to thank all my family and friends without whom any of this would not be possible.

This work was supported by the Ragon Institute of Massachusetts General Hospital, Massachusetts Institute of Technology and Harvard.

Contents

List of Figures

List of Tables

Introduction

Higher organisms have an adaptive immune system which is remarkable in that it enables them to mount pathogen-specific responses to a diverse and evolving world of microbes [1]. B and T Lymphocytes (B cells and T cells) play a key role in mediating adaptive immune responses. Each human has billions of B cells and T cells, and most have a receptor on their surface that is distinct from that on another lymphocyte of the same type. If a B cell's receptor can bind sufficiently strongly to a pathogen's surface proteins, like viral spikes, it can get activated and ultimately produce antibodies specific for this pathogen. The antibodies bind to the viral spike and prevent the pathogen from infecting other cells. Antibodies can also cause some clearance of infected cells [2, 3]. T cells can bind to parts derived from a pathogen's internal proteins (peptides) that are displayed on the surface of infected cells, and if a T cell binds sufficiently strongly to a particular viral peptide it can be activated. Activated T cells can coordinate an immune response in many ways. For example, certain kinds of T cells, called CD8 T cells, can kill infected cells that display the peptide that activated them. Upon infection or vaccination, T cells and B cells specific for the pathogen or vaccine antigen are elicited. The resulting antibodies and memory B cells and T cells remain after infection is cleared or after vaccination, and can protect against future infections. Pathogen-specific immunological memory is the basis for vaccination. To date, despite enormous efforts and expense, a universal vaccine that can protect against a highly mutable virus like HIV or even the seasonal influenza is unavailable.

Antibodies have been isolated from some infected persons that can neutralize

diverse strains of HIV and influenza [4, 5]. Such antibodies are called broadly neutralizing antibodies (bnAbs). For HIV, peptides that are constrained from mutating because of severe penalties to the virus' fitness have also been identified, and these are targets for potent T cell responses that could potentially be elicited by vaccination [6–13]. These advances suggest that antibody- and T-cell-based therapies are also promising avenues for the treatment of highly mutable viruses, such as HIV or influenza [14–17]. The effectiveness of these immunotherapies is linked to the breadth of coverage of mutant viral strains by bnAbs, and whether or not some mutations in different sequence backgrounds may still allow the virus to escape T cell responses that target mutationally constrained peptides. The greater the breadth of coverage of bnAbs or mutational constraints on a peptide, the lesser the likelihood that the virus population will be able to produce an escape mutant that will make the treatment ineffective. Thus, current approaches are focused on increasing the breadth of individual bnAbs and vaccinating with multiple mutationally constrained peptides to minimize the chance of viral escape [16, 18, 19].

Immunotherapies aim to treat patients with existent viral infection, for whom there can be a large inter-patient variability in virus sequence diversity at the time of treatment. Much of the virus diversity, however, may not be contained in the actively infectious virus population. Retroviruses, like HIV, which will be the principal example that we will consider, have large latent reservoirs of infected cells which do not produce new virus but instead accumulate over time while avoiding detection by the immune system. Because viruses can remain latent for long periods of time, these latent reservoirs can contain much more diverse virus mutants than the actively infectious viruses, making escape from treatment more likely when these latent populations reactivate [20–23]. Treating viral infections with immunotherapies requires either supplying exogenous bnAbs that can bind to virus proteins with high breadth (passive therapy), or inducing a T cell response that targets mutationally constrained peptides by immunizing with these antigens (therapeutic vaccination).

For HIV, passive immunization with individual bnAbs, or combinations of bnAbs that target different regions of the virus, have shown promise in their ability to reduce viral load in patients [24–27]. These antibody therapies can clear both free virus in plasma as well as some actively infected cells; however, many patients still experience viral escape through the emergence of mutations that evade the coverage of the antibodies [21, 27, 28]. Therapeutic vaccination with identified mutationally constrained peptides to elicit potent T cell responses is still being developed, but such vaccines containing larger parts of the proteome have been tried [29–31]. In contrast to bnAb therapies, however, T cell therapies can only target and clear actively infected cells and thus result in different viral population dynamics during treatment.

Ex-vivo expansion and engineering of T cell populations to treat viral infections are another possible solution [14]. Recent advancements in chimeric antigen receptor (CAR) T cell therapies have shown promise in their ability to control these diverse virus populations [32, 33]. The engineered T cell receptors can bind to a large fraction of viral strains, and as in the antibody case, T cells with multiple receptors can be administered simultaneously to increase the breadth of the treatment [33, 34].

The probability of success of such immunotherapy approaches is difficult to predict for multiple reasons. First, it is challenging to comprehensively measure or compute the effect of specific mutations on the ability of the virus to escape a particular treatment [35, 36]. Second, the effect of pre-treatment diversity of viral sequences in both actively infectious and latently-infected cells on the ability of passive therapies to control virus populations has not been fully characterized [20, 37]. Finally, in the absence of drugs that promote viral gene expression, it is difficult to clear a significant portion of the latent viral reservoir during the relatively brief period during which therapy is administered [24, 38].

Prior attempts to model the dynamics of viral response to treatment used systems of ordinary differential equations to characterize the time evolution of viruses,

infected cells, and immune responses during and after treatment. These models can give us insights into the biological mechanisms underpinning the response to passive immunotherapies or therapeutic vaccines and help explain the variable responses observed in different patients [2, 39]. However, these models do not include stochastic effects that are inherent to real biological systems, and have a limited ability to model viral diversity and heterogeneity [40]. Stochastic versions of a birth-death-mutation process have been studied previously in the context of other immunological processes, showing their ability to provide quantitative and qualitative insight into these complex biological systems [41–43]. There are stochastic models that have been used to study viral dynamics, but they are limited in their ability to describe necessary features of our problem, such as viral diversity or response to immunotherapies [8, 44].

In this work, we develop a stochastic framework to model the growth, death, and mutational dynamics of a virus population under the immune pressure imposed by T cell or antibody therapy. Our model is parameterized with experimental data and recapitulates viral rebound characteristics observed in clinical trials [24]. In addition, our stochastic model highlights the effectiveness of combination therapy, wherein an extended bnAb treatment can help clear rare strains in the latent reservoir that would otherwise escape the T cell response. We find that this cooperative effect between bnAbs and T cells is maximized when viral diversity is low. Finally, we show that the expected time to cure is unaffected by immunotherapies, but can be reduced by methods that increase the rate of activation of the latent reservoir. With these findings, we hope to guide therapy design to maximize the chances of sustained viral control.

Coarse-grained model for viral dynamics with bnAb and T cell therapeutics

A complete description of viral and immune dynamics requires the incorporation of biological details which can be incredibly complex or poorly understood. Thus, our goal is to design a minimal model that contains sufficient features of the underlying biological process that we can recapitulate experimentally-observed trends and gain an understanding of the characteristics relevant to improving bnAb and T cell therapies.

Viral dynamics alone can be a source of great complexity. Free virus in the plasma can circulate and infect new cells, and infected cells can rapidly produce new virus leading to rapid proliferation if left unchecked. Because of error-prone reverse transcriptase machinery and an extensive set of post-translational modifications, the HIV virus mutates often and follows a very complex fitness landscape that makes describing the possible mutational transitions challenging. Thus, we make some simplifying assumptions about the characteristics of the dynamics. First, we work explicitly only with the populations of infected cells. We assume we are in the regime where the number of uninfected, susceptible cells is larger than our infected population. The dynamics of free virus in plasma is then incorporated implicitly through the appropriate choice of constants. While there is evidence that cells can be multiply infected by different virions, we note that the resulting dynamics have been found to be equivalent to singly infected cells [45, 46]. Thus, we do not include any effects of viral recombination and assume each infected cell can be associated with a single viral

strain. In this view of infected cells, we adopt a coarse-grained view of viral diversity, where the virus population can be from one of N discrete strains, ordered from 1 through N. The infected cells can infect new cells with rate constant r and mutate from strain i to j with a rate $\mu_{ij} \equiv \mu^{|i-j|}$ that decreases the further apart i and j are. For ease of computation and interpretability, we decouple replication and mutation in our model; both are simply included as independent first-order reactions.

The infected cells can potentially be cleared upon clearance of these infected cells can be done with the administration of T cells or bnAbs, each of which has a certain probability p of covering a particular viral strain. If the strain is covered, it is cleared with a rate $c > r$ which may be different between T cells and bnAbs. In the context of bnAbs, the physical notion of the coverage fraction is relatively straightforward, analogous to the coverage estimated from traditional neutralization assays [47]. In the case of T cells, the coverage fraction requires a bit more care to interpret. If a T cell targets an epitope from a mutationally vulnerable region of the virus sequence,

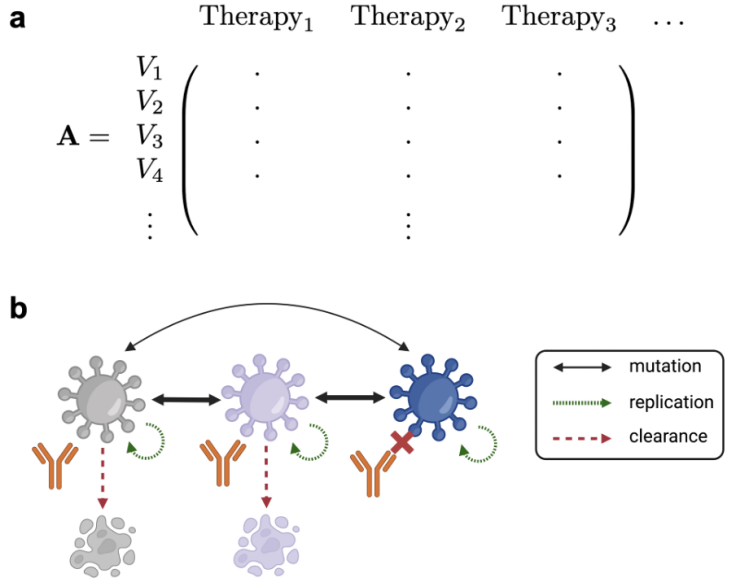

Figure 1: **a.** The coverage matrix structure, representing the clearance rate of each virus strain V_i by each therapy. **b.** Graphical depiction of the dynamics modeled by the stochastic master equation for three strains, including replication, clearance, and mutation with rates that are a function of the distance between strains.

we consider it to have a higher coverage fraction; escape mutations in such a region would be associated with a large fitness cost, rendering a large fraction of the nearby sequence space "covered" by the T cell. On the other hand, targeting an epitope from a highly mutable portion of the virus would result in a T cell with low coverage. In both cases of bnAbs and T cells, we allow for the possibility of combined treatment; in such cases, we assume each therapy targets a different epitope allowing for independent random selection of their covered strains. The clearance of each strain by each treatment is encoded in the random matrix $\mathbf{A}$, outlined schematically in Figure 1a. The difference between bnAb and T cell treatments biologically comes from the ability of antibodies to neutralize free virus. Thus, the presence of bnAbs not only clears infected cells, but also reduces the rate of replication and mutation. We therefore introduce a parameter α_i which is a constant factor α for a covered strain i and 1 otherwise; the replication rate and mutation rates for strain i are then r_i/α_i and μ_{ij}/α_i. For bnAbs we set $\alpha > 1$, and for T cells $\alpha = 1$. A summary of the dynamics experienced by the infected cells is presented in Figure 1b.

Each actively infected cell strain also has a latent counterpart which cannot proliferate but also is not cleared by the administration of treatments. To model the increased diversity in the latent reservoir, we also allow at time 0 there to be strains with cells in the latent reservoir that are not present in the active population. All cells in the latent reservoir can activate with a certain rate $\mu_{L\to v}$ and decay slowly over time. Although we can identify a baseline biological value for $\mu_{L\to v}$ (Appendix A), we can also treat it as a controllable parameter to model the effect of HIV latency reversal drugs which increase the rate of HIV activation [48–50].

Each of the rates above describes a stochastic process, which can be described using master equations, a set of coupled differential equations which describe the transition rates between states of the system. In our case, the states are described by the vectors $V_a \in \mathbb{N}_0^N$, $V_L \in \mathbb{N}_0^N$ which describe the population in each strain of the active and latent population, respectively. Then, for a particular strain i, the master

equations for the probability of being in a particular state read

$$\begin{aligned}
\frac{dp_{a,i}\left(V_a^i,t\right)}{dt} = & \\
& - p_{a,i}\left(V_a^i,t\right) V_a^i \left(r_i + \sum_m A_{im} + \sum_{i\neq j} m u_{ij} + \mu_{v\to L} \right) \\
& + p_{a,i}\left(V_a^i+1,t\right)\left(V_a^i+1\right) \left(\sum_m A_{im} + \sum_{i\neq j} m u_{ij} + \mu_{v\to L} \right) \\
& + p_{a,i}\left(V_a^i-1,t\right)\left(V_a^i-1\right) \Bigg(r_i + \sum_{j\neq i} \mu_{ji} \sum_{V_a^j} V_a^j p_{a,j}\left(V_a^j,t\right) \\
& \qquad\qquad + \mu_{L\to v} \sum_{V_L^j} V_L^j p_{L,j}\left(V_L^j,t\right) \Bigg) \qquad (1)
\end{aligned}$$

and

$$\begin{aligned}
\frac{dp_{L,i}\left(V_L^i,t\right)}{dt} = & - p_{L,i}\left(V_L^i,t\right) V_L^i \left(\delta + \mu_{L\to v} + \mu_{v\to L}\right) \\
& + p_{L,i}\left(V_L^i+1,t\right)\left(V_L^i+1\right)\left(\delta+\mu_{L\to v}\right) \\
& + p_{L,i}\left(V_L^i-1,t\right)\left(V_L^i-1\right)\mu_{v\to L} \qquad (2)
\end{aligned}$$

where $V^i_{(*)}$ is the number of cells in strain i in either the latent (L) or active (a) reservoir, m indexes over the different applied treatments, and δ is the decay rate of the latent reservoir.

The choice of parameters is determined primarily by known biological values. Parameters we are able to estimate this way include the replication rate, clearance rate for bnAb and for T cells, and the latent activation rate. We can also obtain the rate of mutations μ_{ij} in this model, but to do so requires some care. Because each coarse-grained strain in this model can replicate and be targeted by T cell and bnAb therapies, a coarse-grained strain must only reflect replication-competent viruses. Thus, the effective mutation rate in the model will be much lower than the observed *in-vivo* HIV mutation rate of approximately 4×10^{-3}/base pair/cell [51]. We derive

the correction factor to the baseline mutation rate by estimating the fraction of non-synonymous mutations that preserve high fitness [9, 10]. The full calculation of the mutation rate μ, as well as the values of all parameters, are provided in Appendix A.

We can also biologically motivate our choices for the initial conditions of the active and latent reservoirs. We set the size of the initial active pool to $V_{a0} = 10^3$ cells and the initial latent pool to $V_{L0} = 10^4$ cells; such values are in line with the experimental observation that the latent reservoir can often be much larger than the actively infectious virus pool [52]. As a matter of convention, we define rebound to occur when the number of total cells in the active reservoir exceeds $2V_{a0}$, and terminate the simulation at that point.

We model the distribution of these cells among the strains with a geometric distribution – that is, as

$$V_a^i(t=0) = \left\lfloor V_{a0} \cdot \frac{\left(1-\frac{1}{\lambda}\right)^{i-1}}{\lambda} \right\rfloor \tag{3}$$

and similarly for $V_L^i(t=0)$, where we defined a parameter λ that controls the mean of the distribution. By varying λ, we can model reservoirs with different diversities; a larger λ corresponds to a more diverse reservoir. This geometric distribution along the strain axis is analogous to the distribution of sequences observed some time after infection by a founder strain. In such a case, we would expect that many of the strains would have high sequence similarity to the founder (i.e. populate coarse-grained strain 1) and cells with a large number of mutations would be rare. Because we expect HIV sequence diversity to grow over time (in the absence of treatment), the low diversity distribution can also be interpreted as the reservoir of an individual who started antiretroviral therapy soon after infection. A high diversity would then correspond to a patient with longer-term uncontrolled HIV infection.

Finally, we must determine the total size of the coarse-grained sequence space

N. Although we know the full length of the HIV sequence, the total number of real strains is challenging to simulate because of the combinatorial size of the sequence space. For the parameters we identify above, we find empirically that an N=30 is sufficiently large to not qualitatively affect the dynamics while remaining computational tractable. A full table of parameters and sources is provided in Table 1 in Appendix A.

To simulate the master equations, we take two approaches. First is to treat the underlying processes as a chemical reaction network and simulate it directly with a custom variant of Gillespie's stochastic simulation algorithm (SSA) designed to run on GPUs for high throughput and parallelization [53]. This method is faster than available versions running on conventional CPUs in regimes where the population size is bounded [54]. The other method is derived from recent results that provide a framework for analytically solving monomolecular reaction networks with Doi-Peliti field theory [55]. These results were derived primarily by collaborator Henrik Pinholt, and thus are beyond the scope of this thesis. The results presented hereafter are computed by the SSA method described above.

Results

I. Single-therapy results highlight limitations in time to viral extinction

We will begin by investigating the administration of bnAb or T cell therapies for a duration much longer than the time for all the cells in the latent reservoir to activate. This is a realistic scenario for T cell vaccination, as CD8 T cell responses to vaccination have been shown to persist for extended periods of time [56, 57]. However, passive immunization with bnAbs decay on the order of weeks, which means long-duration bnAb coverage requires continual administrations of the treatment which is not practical. However, it is useful to study this limit to better understand the behavior of the model. During extended administration of a single therapy, since the clearance rate of the therapy is larger than the replication rate of the infected cells,

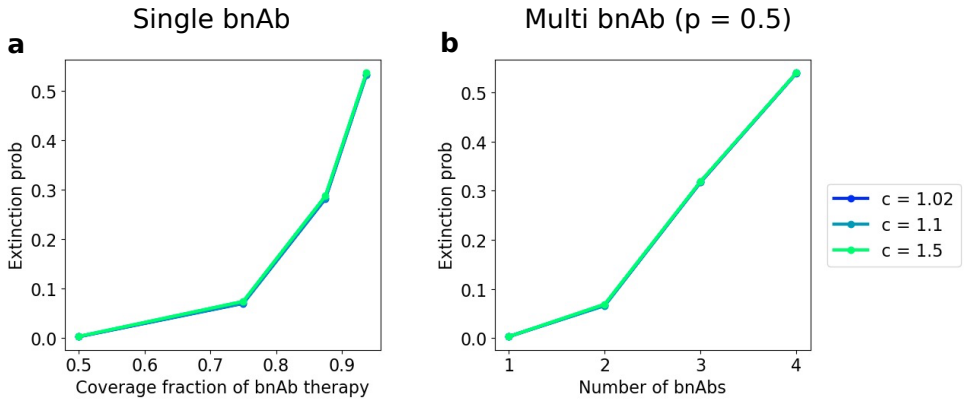

Figure 2: **a.** Extinction probability of virus population for bnAbs administered with given clearance rate and coverage fraction. **b.** Extinction probability for multiple bnAbs administered at the same time, where each bnAb has a fixed coverage fraction. We note that the observed extinction probability in the two cases is nearly identical.

we expect all outcomes to be either extinction or rebound events. To characterize these outcomes, we look at three quantities for each simulation: the time to extinction, the time to rebound, and the probability of extinction. Since both the coverage matrix and the resulting viral dynamics are both stochastic processes, we compute averages of the three observables over 10000 realizations of the coverage matrix and its resulting dynamics.

The main characteristics that define the effectiveness of a treatment are its clearance rate of infected cells and the fraction of viral strains that the treatment can effectively cover. We can vary these two parameters in the case of bnAb therapy, and observe the corresponding extinction probabilities in Figure 2. These outcomes are independent of the clearance rate of the bnAb, as long as it is greater than the replication rate of the infected cells. We also see that combining multiple bnAbs, whose coverage of individual strains is drawn independently at random, is equivalent to administering a single, higher-coverage bnAb such that the probability of covering any one strain in both situations is the same. Thus, we now focus primarily on the single-bnAb case, with the understanding that very high coverage fractions (i.e., > 0.95) could be achieved by either engineering an exceptional bnAb or combining multiple bnAbs with more modest coverage fractions. We observe this pattern for T cell therapy as well (Figure S1).

In Figure 3, we find that the mean time to extinction is independent of both the clearance rate and coverage fraction of the bnAb. This is because any treatment that targets only active infected cells must wait for the latent reservoir to activate, which happens on a timescale much longer than the dynamics of infected cell clearance. As the situation considered here is more realistic for T cell-based therapeutic vaccination (see above), unfortunately, this implies that simply increasing the potency of the magnitude of T cell treatments cannot decrease the average time to cure for HIV-infected patients. However, Figure 3 also shows that increasing the coverage fraction of the bnAb does increase the average time to rebound, primarily because the time

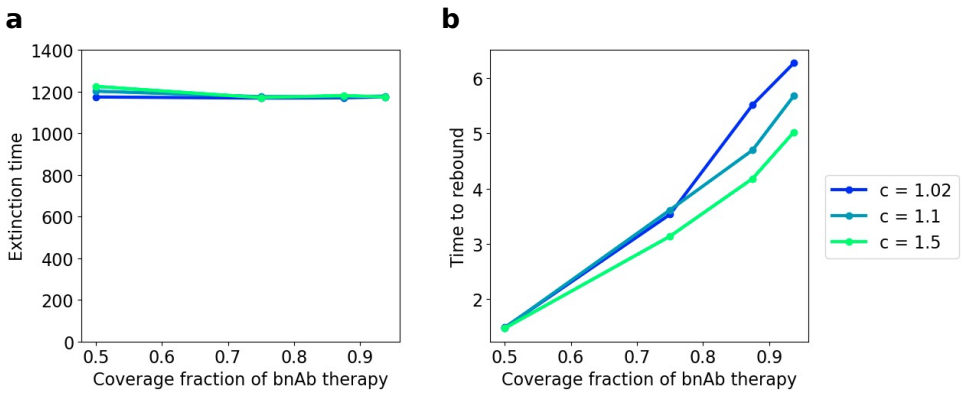

Figure 3: **a.** Mean time to extinction for the virus population when administered bnAb therapy. **b.** Mean time to rebound for the virus population when administered bnAb therapy. Each color corresponds to a different viral clearance rate of the bnAb therapy. The extinction time is independent of both the clearance rate and coverage fraction, while time to rebound increases with increasing coverage.

to find an uncovered strain by accumulating mutations or activating a rare latent cell increases. We find a similar effect for T cells, indicating the increase in rebound time is a result of increased coverage and not the mutational inhibition effects particular to bnAbs.

II. Influence of latent reservoir characteristics on combination treatment outcomes

Recent findings suggest that administration of bnAbs leads to the formation of immune complexes that result in a T cell response that lasts for much longer than the bnAb therapy [18]. In addition, new trials combining simultaneous injection of bnAbs with T cell vaccination demonstrate a promising avenue for combination therapy [30]. Thus, our next goal is to use our model to investigate these scenarios with combined bnAb and T cell therapy based on vaccination.

To model the transition between bnAb-mediated and T cell-mediated clearance, we simulate viral dynamics under solely bnAb therapy for some time t, and then switch instantly to T cell dynamics until the end of the simulation. Physically, this

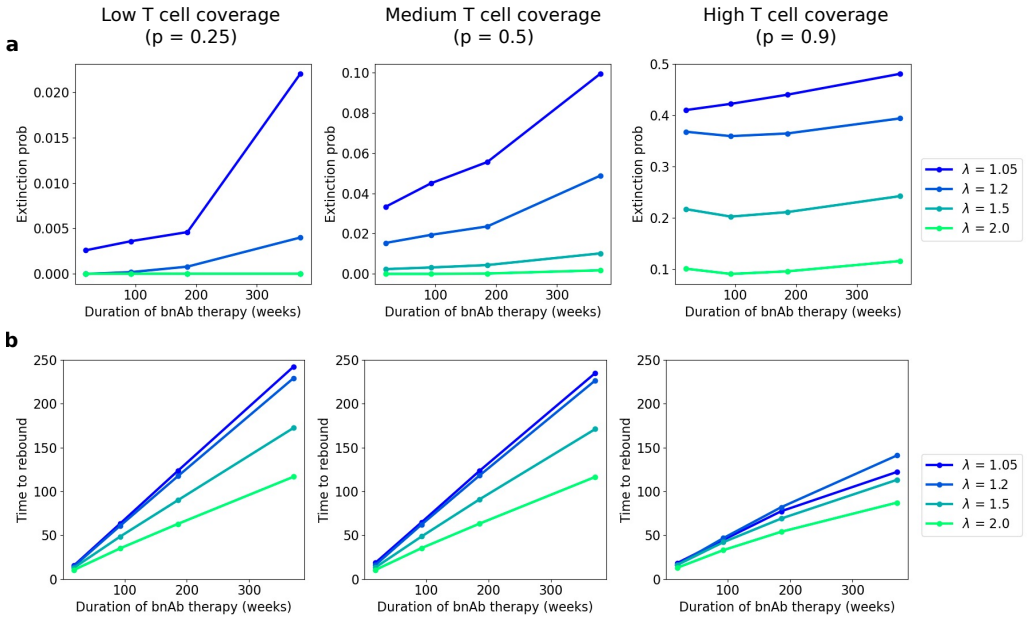

Figure 4: **a.** Extinction probability and **b.** time to rebound for bnAb therapy followed by a T cell response as a function of the duration of bnAb therapy with low, medium, or high T cell coverage. Different colors indicate the diversity of the latent reservoir, parameterized by λ (larger λ implies more diverse). Both extinction probability and time to rebound increase with longer duration bnAb therapy, and decrease with increasing latent diversity.

corresponds to administering bnAbs periodically for a fixed duration, and then giving a therapeutic T cell vaccine just before bnAb titers decay. As we can see in Figure 4, the probability of extinction depends strongly on the strength of the T cell response, for which an increase in coverage from 25% to 50% yields an order of magnitude increase in extinction probability. In this regime of low T cell coverage, even relatively modest improvements in coverage fraction result in significantly greater chances of covering escape mutations. These results highlight the how small improvements in T cells' ability to target mutationally constrained peptides can have significant therapeutic benefit over an unspecific response as elicited by whole-proteome vaccination. Further comparison with the 90% T cell coverage case indicates that the duration of bnAb therapy has the greatest percentage impact on extinction probability when the T cell coverage is low. Thus, when administering highly potent T cells, we do not need to rely as heavily on bnAb therapy to help control the virus population until the rare strains in the latent reservoir are eliminated.

The average time to rebound, however, consistently maintains a linear relationship with the duration of bnAb therapy. This suggests that the virus populations that are controlled by the bnAb therapy but can escape the T cell response rebound at a fixed time after the end of bnAb treatment if they are not cleared before then. In all cases, the extinction probability and time to rebound decrease as the diversity of the latent reservoir increases, which is in agreement with our intuition that a more diverse latent reservoir is more challenging to fully cover with a given therapy.

We investigate further how the properties of the latent reservoir affect viral rebound and escape; Figure 5a suggests that decreasing the size of the latent reservoir can significantly improve outcomes, and increases the effect of bnAb therapy duration because there is a greater chance of clearing entire strains from the latent

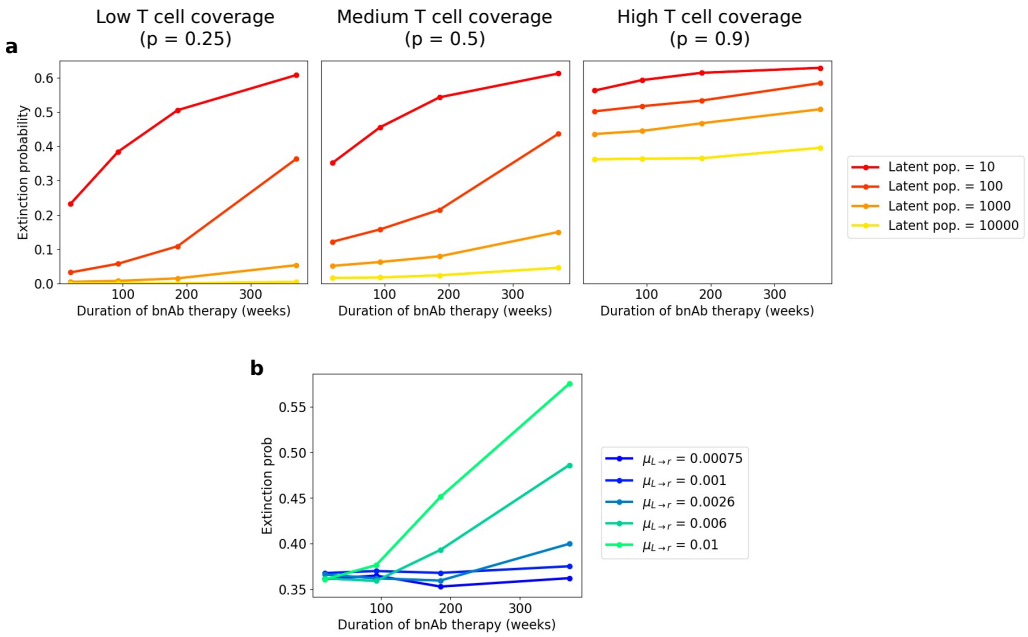

Figure 5: Extinction probability of the virus population for bnAb followed by T cell therapy, as a function of the duration of bnAb therapy. **a.** Different colors denote the size of the latent reservoir, with a fixed distribution. High T cell coverage after the end of bnAb therapy helps mitigate the chances of escape for large latent reservoir sizes. **b.** Different colors denote the activation rate of cells in the latent reservoir. The extinction probability is independent of the duration of bnAb therapy until a critical value of the latent activation rate, which is close to the biological parameter estimate of $\mu_{L\to v} = 0.0026$.

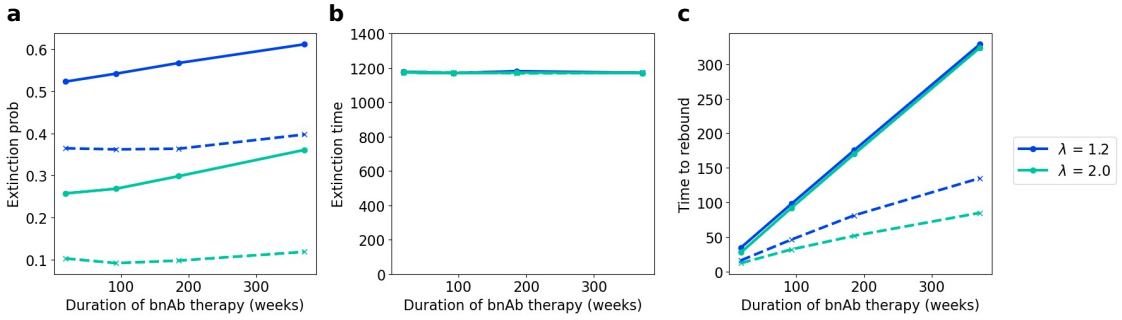

Figure 6: **a.** Extinction probability of the virus population, **b.** mean time to extinction and **c.** mean time to rebound for bnAb + T cell combination therapy (solid) and bnAb followed by T cell therapy (dashed), as a function of the duration of bnAb therapy. Different colors indicate the diversity of the latent reservoir, parameterized by λ. The coverage of the bnAb-only treatment is the same as the bnAb coverage in the combination treatment.

reservoir. Since early initiation of antiretroviral therapy is associated with smaller HIV reservoirs, we again highlight the therapeutic benefit of early detection and intervention before expansion of the latent reservoir [58]. We can increase the chance of clearing strains from the latent reservoir during the course of bnAb treatment by increasing the latent activation rate, analogously to transcription-promoting drugs that have recently been tested in combination with bnAb regimens [30, 38]. As the latent activation rate increases, we note an interesting transition in Figure 5b at the biologically-estimated activation parameter of $\mu_{L\to v} = 0.0026$ (Table 1). For slower activation rates, the extinction probability is independent of the duration of bnAb therapy, because the timescale of clearing the more populated strains is longer than the course of bnAb therapy. However, as we increase the latent activation rate beyond the biological baseline, we uncover a strong dependence of extinction probability on the duration of bnAb therapy, since there is now sufficient time to clear strains from the latent reservoir before the onset of T cell-mediated clearance.

We can also look at the case where bnAb therapy and a T cell vaccine are administered simultaneously, with bnAb therapy terminating at a fixed time and the T cell response continuing until the end of the simulation. Figure 6 shows how for any degree of diversity in the latent reservoir, the combination therapy outperforms the sequential treatment substantially, regardless of the bnAb treatment duration or di-

versity of the latent reservoir. In particular, increasing the duration of bnAb therapy has a much greater effect on the time to rebound when it is administered in combination with T cells. This is because the viral reservoir must escape both bnAb and T cell clearance to rebound early, which is much less likely than escaping bnAb therapy alone. We can confirm this hypothesis by comparing the bnAb-T cell combination with a bnAb-T cell sequential therapy where the bnAb has a 99% coverage fraction to be equivalent to two independent 90% coverage treatments. In Figure S2, we see that the sequential case yields comparable or better extinction probabilities and times to rebound because of its additional suppression of replication and mutation. However, a bnAb therapy with such high coverage is only achievable using combinations of bnAbs; thus, it is promising to see that a combination of a single bnAb with T cell vaccination can yield comparable efficacy to a multi-bnAb treatment.

III. Exploring the space of achievable therapy outcomes

As we consider the clinical implications of these results, it becomes apparent that certain therapeutic outcomes (i.e., duration and probability of viral control) are only achievable under certain treatment conditions. Thus, we now leverage our model to explore the space of achievable therapy outcomes as a function of changeable parameters, such as treatment coverage, duration, and latent reservoir characteristics.

To better contextualize our results above, we look at the fraction of simulations that maintain viral control as a function of time when undergoing a course of bnAb followed by T cell therapy, which we can characterize using a diverse range of parameters (Figure 7). Across all parameter regimes, we observe the same characteristic curve for the survival probability in which rebound primarily occurs at the onset of new treatments. In other words, if the virus population is able to escape the bnAb, it does so soon after treatment onset at $t = 0$, and the viral populations that are controlled by the bnAb but survive long enough to escape the T cell similarly tend to escape soon after bnAb therapy ends. This bimodal distribution in rebound time

explains why the mean time to rebound is earlier than the end of bnAb therapy, and validates the decrease in time to rebound for sufficiently fast-activating latent cells which either escape early or get completely cleared.

We then construct “phase diagrams", to understand how changes in treatment parameters and latent reservoir characteristics move us in the space of probability of control vs time to rebound. There are several qualitative heuristics we can obtain from Figure 8. First, increasing the coverage fraction of the bnAb has a bigger effect on the probability of indefinite viral control than on the average time to rebound for shorter bnAb courses. The bnAb therapy can only control the virus population for as long as it is administered, so for short-duration treatments the virus rebounds quickly

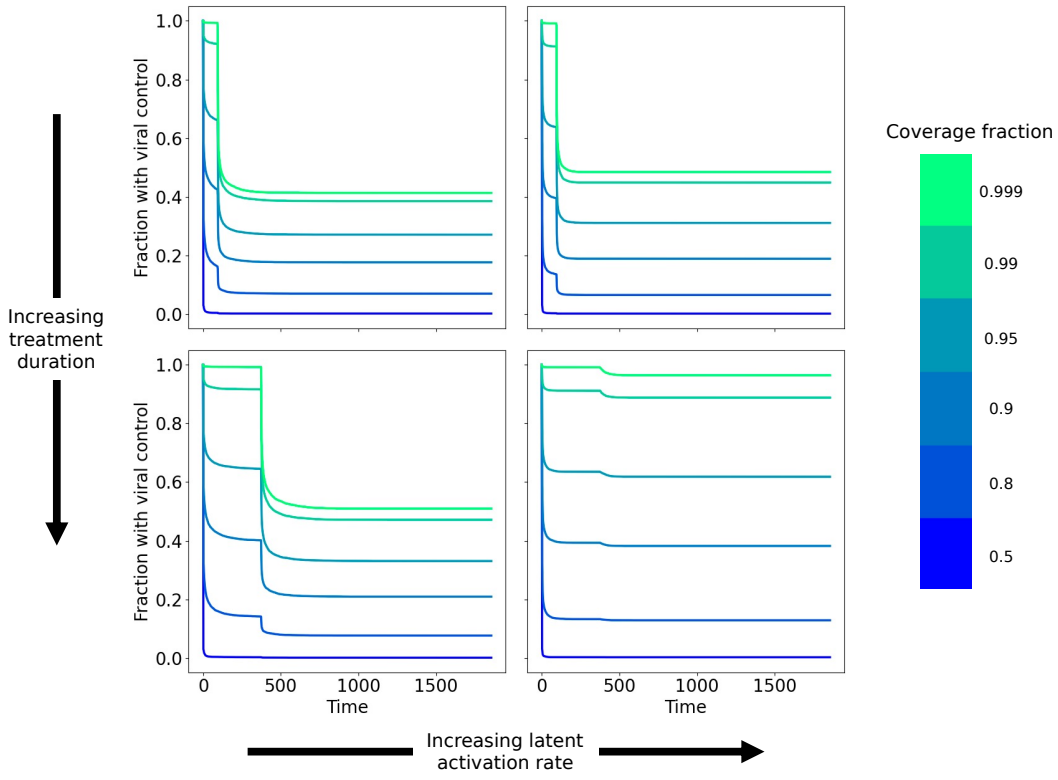

Figure 7: Survival probability plots showing the fraction of simulations which maintain viral control as a function of time. Each color specifies the coverage fraction of the bnAb therapy. Each row in the grid of plots has a fixed bnAb treatment duration (top to bottom: 93, 370), and each column has a specified activation rate for cells in the latent reservoir (left to right: 2.6e-3, 1e-2).

unless control is achieved. As a result, the most effective way to delay the average time to rebound is to administer the bnAb for a longer period of time.

Second, varying the diversity of the latent reservoir does not affect the time to rebound, and simply changes the probability of control. Increasing diversity does not significantly change the bimodal shape of the rebound time distribution, because the majority of rebound events are due to highly populated latent strains that activate soon after coverage is lost. These plots also show the effect of the ballistic trajectories to viral rebound discussed above, where for fast-activating latent cells exposed to long durations of bnAb therapy the average time to rebound is close to zero. In a clinical context, this would manifest as a few patients rebounding early while the remainder maintain low virus populations for an extended period of time.

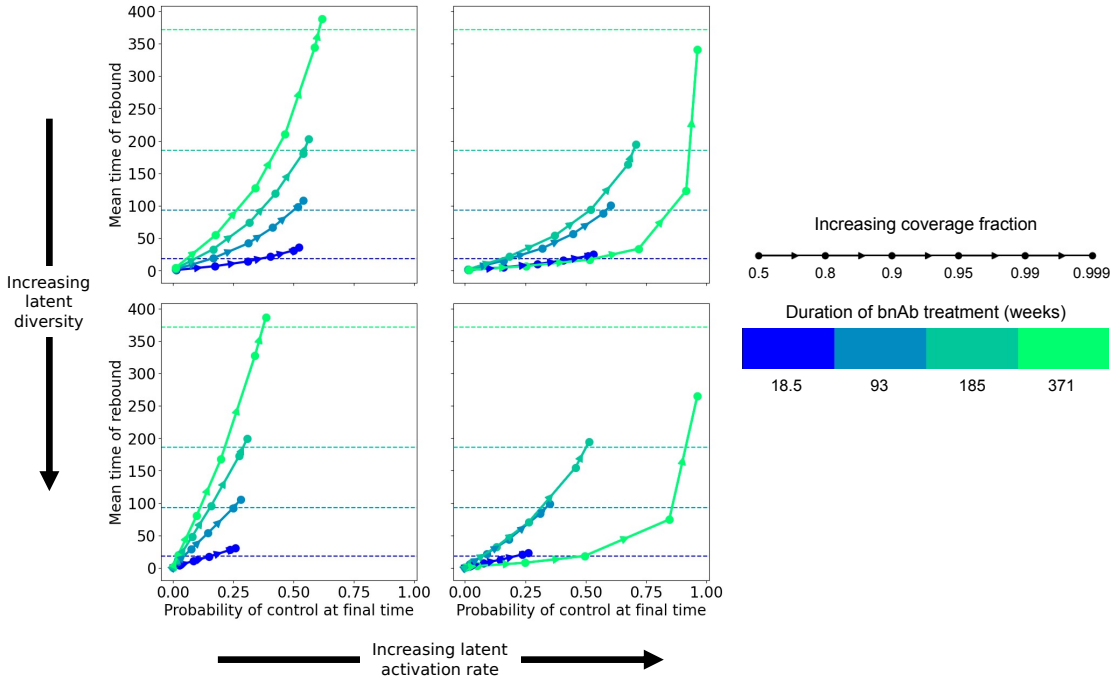

Figure 8: bnAb followed by T cell treatment outcome plotted in the space of mean time of rebound vs probability of control at final time. Each color specifies the duration of the bnAb treatment, and the arrows show the trajectory of the outcome as the bnAb coverage fraction increases. Each row in the grid of plots has a given latent reservoir diversity (top to bottom: mean = 1.2, 2.0), and each column has a specified activation rate for cells in the latent reservoir (left to right: 2.6e-3, 1e-2). The dashed horizontal lines indicate the duration of bnAb treatment corresponding to the color of the line.

These phase diagrams also give us a way to understand the dynamics of rebound after the end of bnAb therapy. Specifically, if we look at the time to rebound given that the simulation maintained control until the end of bnAb therapy, we find that increasing the coverage fraction of the bnAb therapy has no effect on the time to rebound (Figure S3). This suggests that after bnAb treatment end, patients that then rebound will tend to do around the same time regardless of the coverage fraction of the bnAb therapy. Because the probability of rebound changes, we conclude that the bnAb therapy must only be eliminating the rare strains in the latent reservoir, while the well-populated strains are not significantly depleted. If we recall the equivalence we demonstrated above between increasing coverage fraction and administering more bnAbs, we note that this finding is in agreement with clinical trials of bnAb combination therapies. For example, Figure 1b of Gaebler, et al. shows how the time to rebound after the termination of bnAb therapy is not statistically different for different numbers of bnAbs [24].

Discussion

Designing effective therapies for those who have been infected by highly mutable viruses such as HIV is a major medical and scientific challenge. This task is made more difficult by the dynamics of HIV infection, during which many infected cells may remain latent and undetectable to the immune system. Thus, the ultimate goal is to design therapeutic regimens which are robust to the inter-patient variability of the sizes and diversities of their active and latent virus populations. In this paper, we develop a model which incorporates active and latent viral diversity as part of the viral dynamics and which can distinguish between the differing biological mechanisms behind bnAb- and T cell-based therapies. This model was designed to be both analytically and computationally tractable, which enabled us to explore large sections of the parameter space to better understand the theoretical limits of these therapies and uncover heuristics for their design.

Our model suggests that while combining multiple therapies can provide significant improvement to the probability of viral extinction, conventional bnAb and T cell approaches do not have any effect on the time to viral extinction (i.e., cure). Combining bnAb and T cell therapies does perform better than either therapy alone, and we find that the best way of delaying viral rebound is to administer bnAb therapy for as long a duration as possible. Administering bnAb and T cell vaccination in combination performs better than sequential administration, and is comparable in performance to a multiple-bnAb sequential therapy which would be more cost-prohibitive in reality. We also find that the efficacy of a particular bnAb or T cell is determined primarily by its coverage breadth, as there are minimal returns to increases in clearance rate once it exceeds the replication rate of the virus.

Clinically and experimentally, our results indicate a path forward for designing therapeutic regimens. New drugs that increase the transcription rate of latently infected cells we find can have massive benefits, in particular when given with combination therapy. In addition, bnAb treatments that are given for a short duration should be optimized to clear rare strains from the latent reservoir to maximize the chance of subsequent T cell control. Finally, we see a strong dependence of the outcome on the initial diversity of the latent reservoir; even for reservoirs of the same size, a reservoir of low diversity can be controlled for longer with greater probability than a reservoir with high diversity. This highlights the benefits for early intervention for these HIV treatments, before the accumulation of a large, diverse latent reservoir. Even in cases where early therapeutic vaccination regimens are not possible, placing a recently-infected patient on antiretroviral therapy to help limit the expansion of the latent reservoir can significantly increase the chances of success of bnAb or T cell therapy in the future.

Our model can be extended to incorporate more biological features that are known to be important for viral infection. For example, the decay of bnAb concentration over time could potentially lead to increased viral resistance to those bnAbs, a feature that is possible to learn from the mutational landscape of the model [25]. We hope that our work will serve as a platform for further theoretical investigation of immunotherapies, and that future findings can shape the direction of bnAb and T cell therapeutic design.

References

1. Janeway, C. A., Travers, P., Walport, M. & Shlomchik, M. J. *Immunobiology* 5th ed. (CRC Press, Boca Raton, FL, June 2001).
2. Lu, C.-L. *et al.* Enhanced clearance of HIV-1–infected cells by broadly neutralizing antibodies against HIV-1 in vivo. *Science* **352,** 1001–1004 (May 2016).
3. Lu, L. L., Suscovich, T. J., Fortune, S. M. & Alter, G. Beyond binding: antibody effector functions in infectious diseases. *Nature Reviews Immunology* **18,** 46–61 (Oct. 2017).
4. McCoy, L. E. & Burton, D. R. Identification and specificity of broadly neutralizing antibodies against scpHIV/scp. *Immunological Reviews* **275,** 11–20 (Jan. 2017).
5. Zhou, T. *et al.* Multidonor Analysis Reveals Structural Elements, Genetic Determinants, and Maturation Pathway for HIV-1 Neutralization by VRC01-Class Antibodies. *Immunity* **39,** 245–258 (Aug. 2013).
6. Dahirel, V. *et al.* Coordinate linkage of HIV evolution reveals regions of immunological vulnerability. *Proceedings of the National Academy of Sciences* **108,** 11530–11535 (June 2011).
7. Ferguson, A. L. *et al.* Translating HIV Sequences into Quantitative Fitness Landscapes Predicts Viral Vulnerabilities for Rational Immunogen Design. *Immunity* **38,** 606–617 (Mar. 2013).
8. Shekhar, K. *et al.* Spin models inferred from patient-derived viral sequence data faithfully describe HIV fitness landscapes. *Physical Review E* **88,** 062705 (Dec. 2013).
9. Louie, R. H. Y., Kaczorowski, K. J., Barton, J. P., Chakraborty, A. K. & McKay, M. R. Fitness landscape of the human immunodeficiency virus envelope protein that is targeted by antibodies. *Proceedings of the National Academy of Sciences* **115** (Jan. 2018).
10. Murakowski, D. K. *et al.* Adenovirus-vectored vaccine containing multidimensionally conserved parts of the HIV proteome is immunogenic in rhesus macaques. *Proceedings of the National Academy of Sciences* **118** (Jan. 2021).
11. Gaiha, G. D. *et al.* Structural topology defines protective CD8 sup/sup T cell epitopes in the HIV proteome. *Science* **364,** 480–484 (May 2019).

12. Moyo, N. *et al.* Tetravalent Immunogen Assembled from Conserved Regions of HIV-1 and Delivered as mRNA Demonstrates Potent Preclinical T-Cell Immunogenicity and Breadth. *Vaccines* **8,** 360 (July 2020).

13. Mothe, B. *et al.* Therapeutic Vaccination Refocuses T-cell Responses Towards Conserved Regions of HIV-1 in Early Treated Individuals (BCN 01 study). *EClinicalMedicine* **11,** 65–80 (May 2019).

14. Lam, S. & Bollard, C. T-cell therapies for HIV. *Immunotherapy* **5,** 407–414 (Apr. 2013).

15. Margolis, D. M., Koup, R. A. & Ferrari, G. HIV antibodies for treatment of HIV infection. *Immunological Reviews* **275,** 313–323 (Jan. 2017).

16. Julg, B. & Barouch, D. Broadly neutralizing antibodies for HIV-1 prevention and therapy. *Seminars in Immunology* **51,** 101475 (Jan. 2021).

17. Ekiert, D. C. & Wilson, I. A. Broadly neutralizing antibodies against influenza virus and prospects for universal therapies. *Current Opinion in Virology* **2,** 134–141 (Apr. 2012).

18. Caskey, M., Klein, F. & Nussenzweig, M. C. Broadly neutralizing anti-HIV-1 monoclonal antibodies in the clinic. *Nature Medicine* **25,** 547–553 (Apr. 2019).

19. Schommers, P. *et al.* Restriction of HIV-1 Escape by a Highly Broad and Potent Neutralizing Antibody. *Cell* **180,** 471–489.e22 (Feb. 2020).

20. Deng, K. *et al.* Broad CTL response is required to clear latent HIV-1 due to dominance of escape mutations. *Nature* **517,** 381–385 (Jan. 2015).

21. Saha, A. & Dixit, N. M. Pre-existing resistance in the latent reservoir can compromise VRC01 therapy during chronic HIV-1 infection. *PLOS Computational Biology* **16** (ed Kouyos, R. D.) e1008434 (Nov. 2020).

22. Murray, A. J., Kwon, K. J., Farber, D. L. & Siliciano, R. F. The Latent Reservoir for HIV-1: How Immunologic Memory and Clonal Expansion Contribute to HIV-1 Persistence. *The Journal of Immunology* **197,** 407–417 (July 2016).

23. Colby, D. J. *et al.* Rapid HIV RNA rebound after antiretroviral treatment interruption in persons durably suppressed in Fiebig I acute HIV infection. *Nature Medicine* **24,** 923–926 (June 2018).

24. Gaebler, C. *et al.* Prolonged viral suppression with anti-HIV-1 antibody therapy. *Nature* **606,** 368–374 (Apr. 2022).

25. Sneller, M. C. *et al.* Combination anti-HIV antibodies provide sustained virological suppression. *Nature* **606,** 375–381 (June 2022).

26. Scheid, J. F. *et al.* HIV-1 antibody 3BNC117 suppresses viral rebound in humans during treatment interruption. *Nature* **535,** 556–560 (June 2016).

27. Bar, K. J. *et al.* Effect of HIV Antibody VRC01 on Viral Rebound after Treatment Interruption. *New England Journal of Medicine* **375,** 2037–2050 (Nov. 2016).

28. Schoofs, T. *et al.* HIV-1 therapy with monoclonal antibody 3BNC117 elicits host immune responses against HIV-1. *Science* **352,** 997–1001 (May 2016).

29. Pereyra, F. *et al.* HIV Control Is Mediated in Part by CD8 sup/sup T-Cell Targeting of Specific Epitopes. *Journal of Virology* **88** (ed Silvestri, G.) 12937–12948 (Nov. 2014).

30. Peluso, M. *et al. Rebound dynamics following immunotherapy with an HIV vaccine, TLR-9 agonist, and broadly neutralizing antibodies* in *Conference on Retroviruses and Opportunistic Infections* (Feb. 2023).

31. Korber, B. & Fischer, W. T cell-based strategies for HIV-1 vaccines. *Human Vaccines &amp Immunotherapeutics* **16,** 713–722 (Oct. 2019).

32. Kuhlmann, A.-S., Peterson, C. W. & Kiem, H.-P. Chimeric antigen receptor T-cell approaches to HIV cure. *Current Opinion in HIV and AIDS* **13,** 446–453 (Sept. 2018).

33. Herzig, E. *et al.* Attacking Latent HIV with convertibleCAR-T Cells, a Highly Adaptable Killing Platform. *Cell* **179,** 880–894.e10 (Oct. 2019).

34. Maldini, C. R. *et al.* Dual CD4-based CAR T cells with distinct costimulatory domains mitigate HIV pathogenesis in vivo. *Nature Medicine* **26,** 1776–1787 (Aug. 2020).

35. Dingens, A. S., Haddox, H. K., Overbaugh, J. & Bloom, J. D. Comprehensive Mapping of HIV-1 Escape from a Broadly Neutralizing Antibody. *Cell Host &amp Microbe* **21,** 777–787.e4 (June 2017).

36. Barton, J. P. *et al.* Relative rate and location of intra-host HIV evolution to evade cellular immunity are predictable. *Nature Communications* **7** (May 2016).

37. Cohen, Y. Z. *et al.* Relationship between latent and rebound viruses in a clinical trial of anti–HIV-1 antibody 3BNC117. *Journal of Experimental Medicine* **215,** 2311–2324 (Aug. 2018).

38. Gunst, J. D. *et al.* Early intervention with 3BNC117 and romidepsin at antiretroviral treatment initiation in people with HIV-1: a phase 1b/2a, randomized trial. *Nature Medicine* **28,** 2424–2435 (Oct. 2022).

39. Conway, J. M. & Perelson, A. S. Post-treatment control of HIV infection. *Proceedings of the National Academy of Sciences* **112,** 5467–5472 (Apr. 2015).

40. Conway, J. M. & Perelson, A. S. Residual Viremia in Treated HIV Individuals. *PLOS Computational Biology* **12** (ed Antia, R.) e1004677 (Jan. 2016).

41. Baake, E. & Gabriel, W. Biological evolution through mutation, selection, and drift: An introductory review. *Annual Reviews of Computational Physics* **7,** 203–264 (2000).

42. Leibler, S. & Kussell, E. Individual histories and selection in heterogeneous populations. *Proceedings of the National Academy of Sciences* **107,** 13183–13188 (July 2010).

43. Ganti, R. S. & Chakraborty, A. K. Mechanisms underlying vaccination protocols that may optimally elicit broadly neutralizing antibodies against highly mutable pathogens. *Physical Review E* **103** (May 2021).

44. Rouzine, I. M., Weinberger, A. D. & Weinberger, L. S. An Evolutionary Role for HIV Latency in Enhancing Viral Transmission. *Cell* **160,** 1002–1012 (Feb. 2015).

45. Dixit, N. M. & Perelson, A. S. HIV dynamics with multiple infections of target cells. *Proceedings of the National Academy of Sciences* **102,** 8198–8203 (May 2005).

46. Jung, A. *et al.* Multiply infected spleen cells in HIV patients. *Nature* **418,** 144–144 (July 2002).

47. Karuna, S. T. & Corey, L. Broadly Neutralizing Antibodies for HIV Prevention. *Annual Review of Medicine* **71,** 329–346 (Jan. 2020).

48. Kim, J. T. *et al.* Latency reversal plus natural killer cells diminish HIV reservoir in vivo. *Nature Communications* **13** (Jan. 2022).

49. Wightman, F., Ellenberg, P., Churchill, M. & Lewin, S. R. HDAC inhibitors in HIV. *Immunology &amp Cell Biology* **90,** 47–54 (Nov. 2011).

50. Spivak, A. M. & Planelles, V. Novel Latency Reversal Agents for HIV-1 Cure. *Annual Review of Medicine* **69,** 421–436 (Jan. 2018).

51. Cuevas, J. M., Geller, R., Garijo, R., López-Aldeguer, J. & Sanjuán, R. Extremely High Mutation Rate of HIV-1 In Vivo. *PLOS Biology* **13** (ed Rowland-Jones, S. L.) e1002251 (Sept. 2015).

52. Sengupta, S. & Siliciano, R. F. Targeting the Latent Reservoir for HIV-1. *Immunity* **48,** 872–895 (May 2018).

53. Gillespie, D. T. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics* **22,** 403–434 (Dec. 1976).

54. Harris, L. A. *et al.* BioNetGen 2.2: advances in rule-based modeling. *Bioinformatics* **32,** 3366–3368 (July 2016).

55. Vastola, J. J. Solving the chemical master equation for monomolecular reaction systems and beyond: a Doi-Peliti path integral view. *Journal of Mathematical Biology* **83** (Oct. 2021).

56. Davenport, M. P., Ribeiro, R. M. & Perelson, A. S. Kinetics of Virus-Specific CD8 sup/sup T Cells and the Control of Human Immunodeficiency Virus Infection. *Journal of Virology* **78,** 10096–10103 (Sept. 2004).

57. Collins, D. R., Gaiha, G. D. & Walker, B. D. CD8 T cells in HIV control, cure and prevention. *Nature Reviews Immunology* **20,** 471–482 (Feb. 2020).

58. Jain, V. *et al.* Antiretroviral Therapy Initiated Within 6 Months of HIV Infection Is Associated With Lower T-Cell Activation and Smaller HIV Reservoir Size. *The Journal of Infectious Diseases* **208,** 1202–1211 (July 2013).

59. 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* **271,** 1582–1586 (Mar. 1996).

60. Sohail, M. S., Louie, R. H. Y., McKay, M. R. & Barton, J. P. MPL resolves genetic linkage in fitness inference from complex evolutionary histories. *Nature Biotechnology* **39,** 472–479 (Nov. 2020).

61. Bournazos, S., Gazumyan, A., Seaman, M. S., Nussenzweig, M. C. & Ravetch, J. V. Bispecific Anti-HIV-1 Antibodies with Enhanced Breadth and Potency. *Cell* **165,** 1609–1620 (June 2016).

62. Georgiev, I. S. *et al.* Antibodies VRC01 and 10E8 Neutralize HIV-1 with High Breadth and Potency Even with Ig-Framework Regions Substantially Reverted to Germline. *The Journal of Immunology* **192,** 1100–1106 (Jan. 2014).

Appendix A: Model parameters

Parameter	Value	Source
r	0.38 days^{-1}	[59]
c	$1.1r$ — $5r$	[2]
μ	$0.014r$	[9, 10, 51, 60]
$\mu_{L\to v}$	$0.0026r$	[39]
$\mu_{v\to L}$	$10^{-6}r$	[39]
p	0.9	[61, 62]
δ	$0.0005r$	[39]

Table 1: Model parameters and sources.

With the exception of the mutation rate μ, we find the specified parameters above directly from the associated reference, converting units such that the unit of time is r^{-1}. Below, we describe how μ was estimated.

We want the mutational distance between coarse-grained strains to be representative of the observed distances between HIV strains in human reservoirs. Thus, with data from Sohail, et al., we compute the average pairwise hamming distance between plasma HIV gp160 sequences as 1.8 edits [60]. To estimate the rate at which this 1.8 fitness-preserving nonsynonymous mutation occurred, we assume mutations occur uniformly on a random sequence background which has fitness cost distribution given by Louie, et al., 2018 and Murakowski, et al., 2021 [9, 10]. Thus, from the baseline HIV mutation rate of 4×10^{-3}/base pair/replication, we weight by the probability of a nonsynonymous mutation that is fitness preserving and convert to units of r to get the estimated $\mu = 0.014r$. Thus, each coarse-grained strain in our model is defined such that they are sequentially distinct and have sufficiently high fitness to survive innate immune responses.

Supplementary Figures

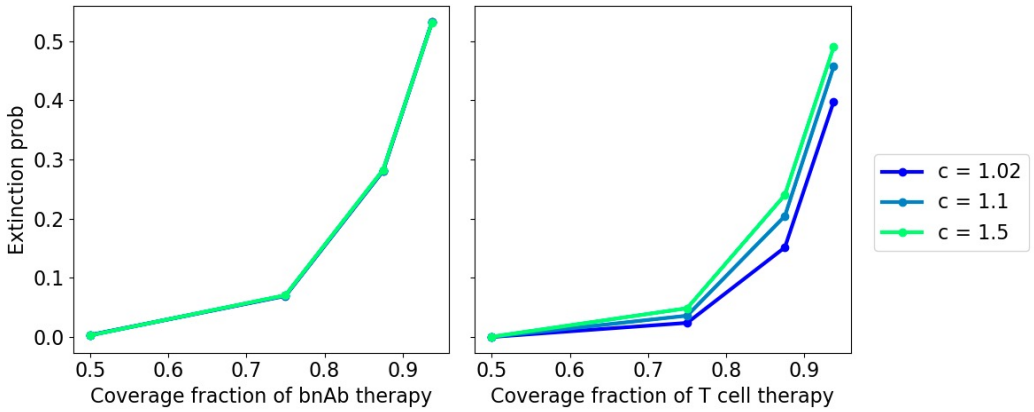

Figure S1: Extinction probability as a function of coverage fraction for left: single bnAb and right: single T cell therapy.

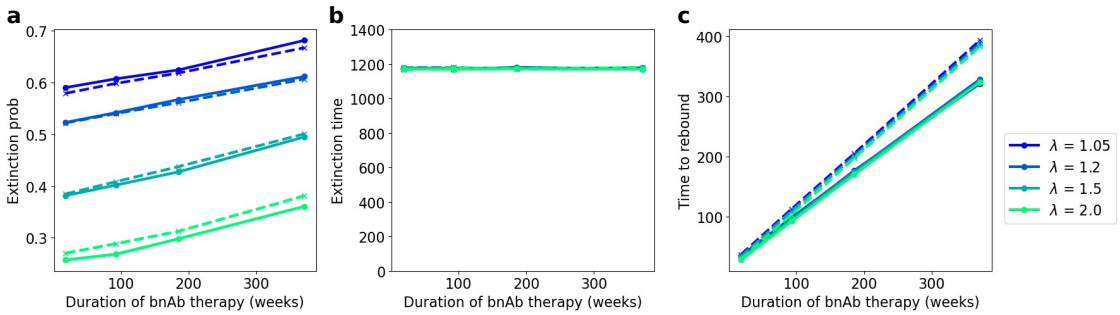

Figure S2: **a.** Extinction probability of the virus population, **b.** mean time to extinction and **c.** mean time to rebound for bnAb + T cell combination therapy (solid) and bnAb followed by T cell therapy (dashed), as a function of the duration of bnAb therapy. The coverage of the bnAb-only treatment was selected to have the same effective coverage fraction as the bnAb + T cell combination treatment. The extinction probabilities are similar, although the bnAb-only treatment leads to a longer average duration of viral control.

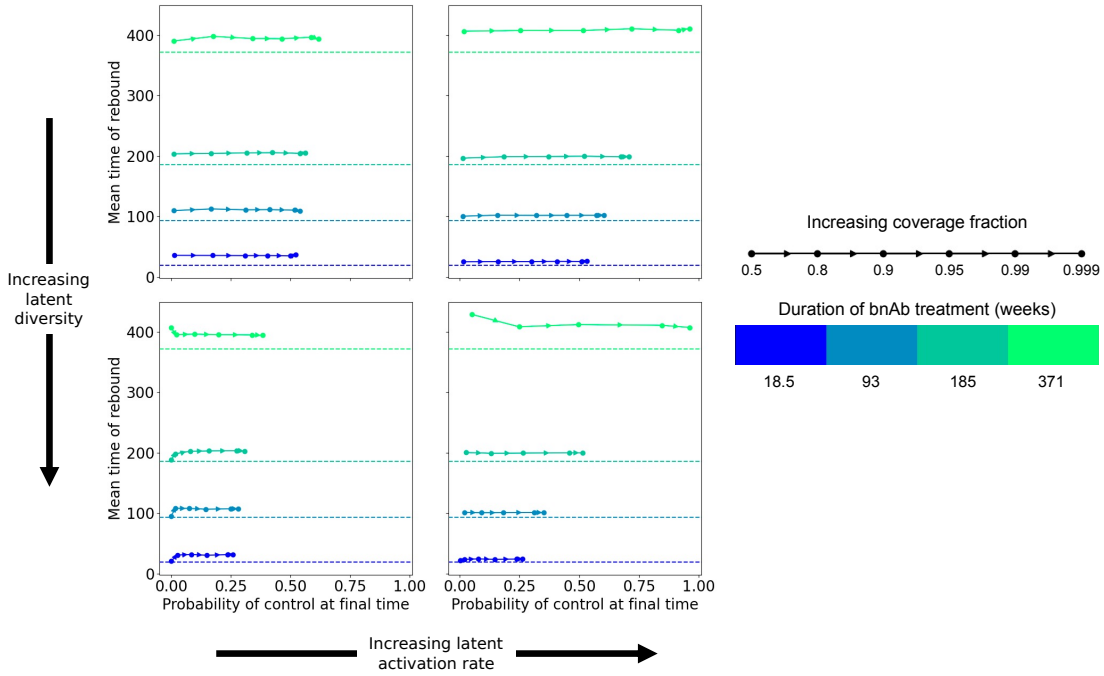

Figure S3: bnAb followed by T cell treatment outcome plotted in the space of mean time of rebound (given control is maintained until the end of bnAb therapy) vs probability of control at final time. Each color specifies the duration of the bnAb treatment, and the arrows show the trajectory of the outcome as the bnAb coverage fraction increases. Each row has a given latent reservoir diversity (top to bottom: mean = 1.2, 2.0), and each column has a specified activation rate for cells in the latent reservoir (left to right: 2.6e-3, 1e-2). The dashed horizontal lines indicate the duration of bnAb treatment corresponding to the color of the line.