Mechanisms underlying vaccination protocols that may optimally elicit broadly neutralizing antibodies against highly mutable pathogens

Raman S. Ganti[1] and Arup K. Chakraborty[1,2,3]

[1] *Institute of Medical Engineering and Sciences, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA*

[2] *Ragon Institute of MGH, MIT and Harvard, Cambridge, Massachusetts 02139, USA*

[3] *Department of Chemical Engineering, Physics, Chemistry, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA*

(Dated: October 7, 2020)

Effective prophylactic vaccines usually induce the immune system to generate potent antibodies that can bind to an antigen and thus prevent it from infecting host cells. B cells produce antibodies by a Darwinian evolutionary process called affinity maturation (AM). During AM, the B cell population evolves in response to the antigen. Antibodies that bind specifically and strongly to the antigen are thus produced. Highly mutable pathogens pose a major challenge to the development of effective vaccines because antibodies that are effective against one strain of the virus may not protect against a mutant strain. Antibodies that can protect against diverse strains of a mutable pathogen are called broadly neutralizing antibodies (bnAbs). In spite of extensive experimental and computational studies that have led to important advances, an effective vaccination strategy that can generate bnAbs does not exist for any highly mutable pathogen. Here we study a minimal model of AM in different time-varying antigenic environments to explore the mechanisms underlying optimal vaccination protocols that maximize the production of bnAbs. We find that the characteristics of the time-varying Kullback-Leibler distance (KLD) between the B cell population distribution and the fitness landscape imposed by antigens is a key determinant of bnAb evolution. The optimal vaccination protocol requires a relatively low KLD in the beginning in order to increase the entropy (diversity) of the B cell population so that the surviving B cells have a high chance of evolving into bnAbs upon subsequently increasing the KLD. For a discretized two-step variation in antigenic environment, there are optimal values of the KLDs for the first and second steps. Phylogenetic tree analysis further reveals the evolutionary pathways that lead to bnAbs. The connections between our results and recent simulation studies of bnAb evolution and the general problem of evolution of generalists versus specialists are discussed.

I. INTRODUCTION

Vaccines are a major contributor to the dramatic decline in childhood mortality over the past century, and have mitigated the threat of infectious diseases around the world. Indeed, vaccination has saved more lives than any other medical procedure. The importance of vaccines for the economy and for human health has been made vivid during the COVID-19 pandemic because we do not have one yet. Vaccines elicit pathogen-specific immune responses that can be rapidly recalled upon natural infection with the same pathogen, thus preventing it from establishing infection. Most such prophylactic vaccines induce our immune systems to produce antibodies that can thwart infection by specific pathogens.

B lymphocytes (B cells), an important part of the immune systems of vertebrates, express a receptor on their surface that is called the B cell receptor (BCR). Humans have about 100 billion B cells each and most possess a BCR that is distinct from that of a different B cell. Antibodies are produced by a Darwinian evolutionary process called affinity maturation (AM) [1]. Upon stimulation by a pathogen or surrogate of a pathogen (antigen) that is used as a vaccine component, BCR on B cells can bind to proteins on the surface of the pathogen or antigen. B cells that bind with a free energy above a threshold can seed structures called Germinal Centers (GCs) where they undergo AM [2]. The BCR of these cells mutate at a high rate [3]. The mutated B cells compete with each other to bind to the pathogenic surface protein, and those with BCRs that can bind more strongly have a better chance to be positively selected. B cells that are not positively selected die. A few positively selected B cells exit the GC, and some of them secrete their BCR in soluble form. The secreted product is an antibody. Other positively selected B cells become so called memory B cells that can subsequently be stimulated rapidly upon exposure to the same pathogen. Most positively selected B cells undergo further rounds of mutation and selection. Therefore, as time ensues, the antibodies that are produced bind more strongly to the pathogenic surface protein [4].

To see how antibodies can prevent pathogens from infecting new cells, consider viruses. Viruses have spikes on their surface comprised of proteins. To infect a cell, a virus' spike binds to a protein on the surface of the host cell. Antibodies can bind to the spike proteins of a virus, and thus interfere with its ability to bind to a host cell protein, which prevents infection [5]. Thus, if potent antibodies that bind specifically to the spike proteins of a virus are elicited by a vaccine, it can protect humans from disease. Highly mutable pathogens can, however, present a major challenge to vaccination using this strategy. This is because mutations can emerge in their spike

proteins, and the antibodies generated by vaccination, which bind to a particular spike protein, are no longer effective [6].

The spike proteins of even highly mutable pathogens, like HIV and influenza, have some regions that cannot mutate because these regions bind to host cell proteins to enable infection [7–9]. For example, a part of the spike protein of HIV is relatively conserved. Antibodies that can bind to these conserved regions would be effective against diverse strains of the pathogens [10]. However, generating such antibodies by vaccination is a non-trivial challenge [11]. Again, consider HIV as an example. The relatively conserved region on the HIV spike is smaller in size than the typical size of an antibody's antigen binding region, and the conserved part is surrounded by highly variable residues of the spike protein [12]. So, the challenge is to design vaccination protocols that can guide the evolutionary process of AM to produce antibodies that bind only to the conserved bits of the spike and avoid the highly variable parts [13]. Antibodies that can achieve such specificity for the conserved parts of the spike, and can thus effectively neutralize diverse mutant virus strains, are called broadly neutralizing antibodies (bnAbs).

Antibodies that can bind to particular strains of a virus' spike are specialists, while bnAbs may be viewed to be generalists [14, 15]. AM is a stochastic non-equilibrium mutation and selection process. The challenge of influencing this process in a way that results in the evolution of generalists, rather than specialists, is a problem that lies at the intersection of statistical physics, evolutionary biology, and immunology. Fundamental insights into this problem will be of significant pragmatic relevance to society.

Recently, a number of simulation studies have been carried out with the aim of shedding light on these issues [2, 16–20]. For those pertinent to the HIV example noted above, a number of insights emerged from such studies. It is evident that in order to generate bnAbs, one would have to stimulate AM with multiple variant antigens that share conserved residues, but have different variable parts. Wang et al [2] showed that sequential stimulation of AM with variant antigens, is likely a more effective way to generate bnAbs than a cocktail of the same antigens. Subsequently, experiments in mice have shown that antibodies resembling HIV bnAbs emerge upon sequential immunization with variant antigens [21]. Upon vaccinating with one variant antigen, AM produces memory B cells that can be stimulated and can undergo further evolution upon vaccinating with a second variant antigen, and so on. Every time a new variant antigen is introduced, the environment in which the existing memory B cells evolved changes, and the memory B cell population is driven from one steady-state to another [22]. So, sequential immunization results in the evolution of B cells by AM in a time-varying environment.

The non-equilibrium response of heterogeneous populations to time varying environments has been explored by numerous authors [23–25]. Clonal populations stochastically switch phenotypes in response to changes in the environment. Recently, stochastic phenotype switching was directed towards searching for optimal conditions that increase the fraction of generalists within a background population of both specialists and generalists [14, 15]. Two recent studies have provided further insights into how vaccination protocols could optimize the evolution of bnAbs, or generalists. Sprenger et al [17] have carried out simulations of the AM process upon sequential immunization and reported that there is an optimal extent to which the existing B cell population should be driven out of equilibrium (from a previous steady-state) in order to maximize the chance of producing bnAbs. Furthermore, they found that this optimum extent is sequentially higher as each new variant antigen is used to stimulate AM. Recently, Sachdeva et. al [15] studied the effect of cycling between antigenic environments. Their simulations showed that if a population of specialist B cells is subjected to time-varying environments cycled at an optimal resonance frequency, specialists will evolve to become generalists while the reverse process is prevented. Interestingly, it was observed that gradually increasing the frequency from slow to fast resulted in an optimal point for generating a large fraction of generalists. This result is related to the shift of the optimum to higher values at each subsequent step of sequential immunization that Sprenger et al reported [17].

In this study, we define a clear information-theoretic metric for the extent to which the existing B cell population is driven out of equilibrium upon sequential stimulation of AM. Using a very simple model of AM, we then explore how B cells gain information by evolution to become bnAbs upon sequential stimulation with variant antigens. Our study provides insights into the evolutionary forces that define optimal vaccination protocols. The key features of trajectories that lead to the evolution of generalist bnAbs are also described.

II. MINIMAL MODEL FOR THE EVOLUTION OF BROADLY NEUTRALIZING ANTIBODIES

As we described in our introductory remarks, upon natural infection or vaccination, antibodies that bind more strongly to the stimulating antigen evolve by a Darwinian evolutionary process. Our focus here is the evolution of bnAbs, which bind to regions of the virus' spike proteins that are relatively conserved across strains. Let us focus here on HIV, a virus against which no effective vaccine exists after over 30 years of research.

The virus' spike protein, or a part of it, has to be the immunogen, or the active component of the vaccine. This is because the B cells that evolve during AM must learn to focus their binding footprint on the relatively conserved residues in the context of the actual geometry of the virus' spike. But, the spike has many regions that do not contain any conserved residues. Human B

cells have a huge diversity of BCRs, and so vaccinating with the spike protein would most likely result in activating B cells that do not bind to the region containing the conserved residues [11]. This is because there are many more regions without conserved residues. AM would then proceed, and generate antibodies that bind well to the variable residues of the particular strain of spike protein that was employed in the vaccine. To address this challenge, immunogens have been developed that can activate the so called germline B cells that bind to parts of the spike containing the conserved region [10]. These immunogens are not the spike protein, but something much simpler. Now, upon vaccination with mimics of the full spike, can bnAbs evolve? If so, what is the optimal strategy? These are the questions we consider here. That is, how to choose characteristics of the variable regions of the immunogens in the vaccine that share the conserved residues, and the concentrations at which they are administered, that would optimize the production of bnAbs.

To obtain essential physical insights into this problem, we construct a minimal model. The affinity between a BCR and an antigen to which it binds is, in principle, defined by many variables, such as the amino acids of the BCR's antigen binding region and the antigen. A simplified representation, called shape space [16], has proven useful for reducing the dimensionality down to a few (5 or 6) abstract variables [26–28]. The breadth of evolving BCRs will similarly be defined by many variables. Inspired by the shape-space model, and in the spirit of Occam's razor, we consider the breadth of coverage of a BCR to be defined by a single dimension. This dimension may be considered to be an appropriate projection of a higher-dimensional manifold, and we will refer to it as the "breadth space".

The origin (0.0 on the abscissa of (Fig. 1(a)) denotes the state of highest possible breadth. If a B cell with a particular BCR sequence is at the origin, it binds with the highest possible affinity to conserved epitopes on the antigens and avoids binding to the surrounding variable region as best as possible. B cells traverse the breadth dimension by mutations. The breadth decreases upon moving left or right from the origin. There is substantially less diversity in the BCR sequences that bind strongly to the conserved residues and avoid the variable residues than those that bind to both the conserved and variable residues [29, 30]. Therefore, for entropic reasons, most mutations are likely to reduce breadth. Some mutations can result in a BCR that does not fold into the proper shape or confer some other grossly deleterious feature. So, any BCR sequence, with some probability, can acquire a lethal mutation [31]. Also, GC B cells are intrinsically apoptotic; i.e., if not positively selected, they die [32, 33]. For computational convenience, we discretize breadth space into a set of $K-1$ bins where K is the total number of states and the additional state corresponds to death (Fig. 1(b)).

The immunogen that leads to the activation and maturation of B cells that can bind to the conserved residues on the spike protein produces a pool of memory B cells that have the potential to evolve into bnAbs. However, these B cells are not bnAbs. The immunogens that mimic the entire virus' spike and are administered subsequently are meant to evolve this B cell population to produce bnAbs. These immunogens, the variant antigens that share conserved residues, impose selection forces on this population of B cells; i.e., the population evolves by mutation to respond to the fitness landscape imposed by these immunogens. At the end of AM induced by a variant antigen, a new pool of evolved memory B cells results. If another immunogen is now administered, this B cell population evolves in response to the fitness landscape imposed by the second immunogen.

In our simple model, the fitness landscape is one-dimensional, and the administered variant antigen imposes different fitnesses to B cells that have acquired different extents of breadth. Since the goal is to evolve bnAbs, as shown in Fig. 1(a), the immunogens are chosen such that the fitness landscape has a peak at the bin corresponding to bnAbs. The choice of immunogen and immunization protocol determines the variance of the fitness landscape that is imposed. For example, if a cocktail of variant antigens that share conserved residues but are very different in their variable regions, is administered, only the B cells that have evolved high breadth will be strongly selected. The corresponding fitness landscape will be sharply peaked (orange and green curves in Fig. 1(a)). If a single immunogen is first administered, then a greater diversity of B cells can be positively selected, and the fitness landscape is characterized by a higher variance (blue curve in Fig. 1(a)). For sequential immunization, as noted above, the fitness landscape changes in discrete time steps. This change corresponds to a time-varying antigenic environment in which a heterogeneous B cell population evolves through replication, mutation, and selection.

The birth-death master equations describing time evolution of the probability distribution of the B cell population vector subjected to mutation and selection is given by [24]:

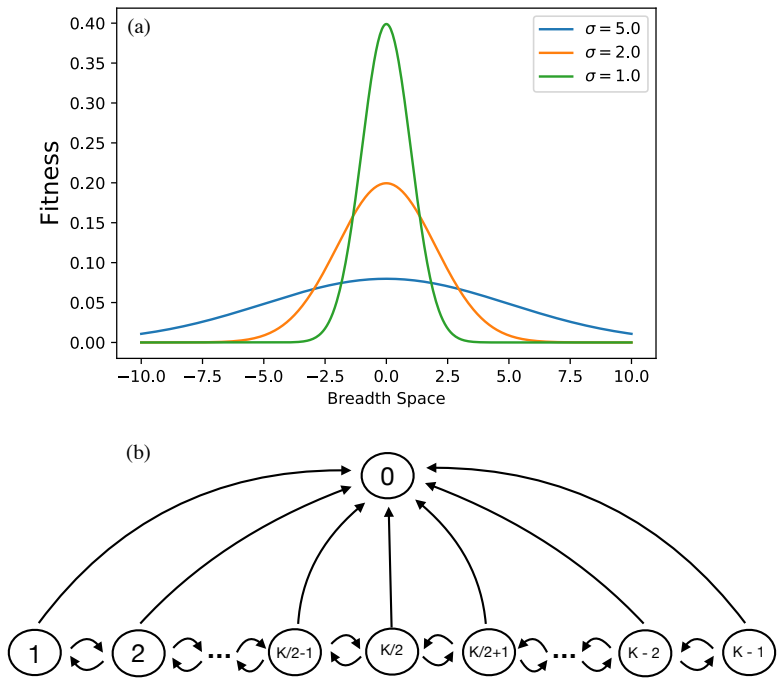

FIG. 1: Definition of breadth space. (a) Depiction of the 1D fitness landscape in a continuous breadth space. The origin denotes the point of highest breadth and the fitness distribution imposed on the B cell population by the immunogen is Gaussian. Reducing the variance of the fitness focuses selection pressure on B cells that bind with high affinity to conserved epitopes. (b) Breadth space is binned into $K-1$ states. $K/2$ is the highest breadth state corresponding to the origin in (a), 1 and $K-1$ are states of lowest breadth, and 0 is the death state.

$$\frac{dp_i(n,t)}{dt} = -p_i(n,t)\left[n(f_i + \sum_{i\neq j}\mu_{ij}) + \sum_{j\neq i}\mu_{ji}\sum_m m p_j(m,t)\right] + \sum_{i\neq j}\mu_{ij}(n+1)p_i(n+1,t) + p_i(n-1,t)\left[(n-1)f_i + \sum_{j\neq i}\mu_{ji}\sum_m m p_j(m,t)\right] \tag{1}$$

where i denotes the index of the breadth bin, n is the number of sequences in bin i, m is the number of sequences in bin j, μ_{ij} is the mutation rate per cell from bin i to j, f_i is the fitness or probability that a B cell in bin i replicates. The combined effects of lethal mutation and basal death rate of GC B cells corresponds to a rate of B cell death, denoted by μ_{i0}. The state space describing the master equations is shown in Fig. 1(b). B cells that occupy bin $K/2$ have attained maximum possible breadth whereas those in states 1 and $K-1$ have lowest breadth. State 0 represents an absorbing boundary condition as it corresponds to B cell death.

Directly solving the master equations is numerically cumbersome. Instead, taking the expectation value of Eq (1) and following the procedure described in the SI of [24], we derive the following mean-field equations:

$$\frac{dN_i(t)}{dt} = (f_i - \sum_{i\neq j}\mu_{ij})N_i(t) + \sum_{j\neq i}\mu_{ji}N_j(t) \tag{2}$$

where N_i is the number of cells occupying bin i and N_j is the number in bin j. Eq (2) clarifies the processes which regulate the population dynamics in each breadth state. Cells in a particular bin in breadth space replicate with a probability that depends on the fitness of cells in the

bin and the number of cells that occupy that bin. They leave bin, i, with a propensity equivalent to the product of the transition rate from i to all other states j and the population of cells in bin, i; entrance into bin, i, is determined by rates j to i and occupancy of states j. Of course, if $i = 0$, the first term in the right-hand side of Eq (2) vanishes.

The master equations (Eq (1)) can be solved by transforming the mean field equations (Eq (2)) into a set of "chemical reactions" and using the Gillespie method [34, 35] to generate stochastic trajectories of the B cell birth-death-mutation process; each trajectory contains the set of all reactions that occur within a single GC. In addition to the reactions provided by Eq (2), there are two important stop conditions that end a GC trajectory: the B cell population dies ($\sum_{i\neq 0} N_i(t) = 0$) or the total number of B cells in the GC approaches a sufficiently large size. The former condition represents extinction, and the latter is a proxy for the B cells having consumed all antigens present in the GC.

The objective of our stochastic simulations is to understand how a time-evolving fitness profile or changing antigenic environment affects the production of bnAbs, and thus determine the mechanistic underpinnings of how certain vaccination protocols may optimize bnAb production. In experimental studies [36], antigen concentration and mutational distance between the variable regions of sequentially administered immunogens are the principal parameters that can be controlled. Changing these parameters is tantamount to changing the characteristics of the fitness landscapes.

How do we relate changing mutational distance between sequentially administered immunogens to changes in the fitness landscape with time? In response to an immunogen, the existing memory B cell population evolves with respect to the imposed fitness landscape (Fig. 2(b), green). The amount of information that the B cell population needs to gain in order to adapt perfectly to the antigenic environment is quantified by the Kullback-Leibler (KL) distance or the relative entropy of the B cell distribution p_i and the fitness f_i [37]:

$$D(p||f) = \sum_{i=1}^{K-1} p_i \log(p_i/f_i) \tag{3}$$

where p_i is the probability that a B cell occupies breadth state i. In the context of Bayesian inference, $D(p||f)$ quantifies the information gained by changing conclusions derived from the prior probability distribution f to the posterior p. The KL divergence measures the maximum information that the B cell population can gain during the evolutionary process. The greater this quantity, the more challenging the non-equilibrium adaptation process. An alternative interpretation of Eq (3) stems from relating $D(p||f)$ to the difference in "free energies" corresponding to the distributions, p and f [38]. In this analogy, the KL distance quantifies how far from equilibrium the initial memory B cell population is being driven by f [39]. In either interpretation, a larger KL divergence corresponds to a greater mutational distance between sequentially administered immunogens. Changing the concentration of the administered immunogen also changes the fitness landscape. Higher concentrations correspond to fitness landscapes with a larger variance.

The remaining parameters, μ_{ij} and μ_{i0}, must be fixed so that the resulting dynamics are consistent with the expected behavior of GCs and B cell sequence evolution [17]. Given that breadth space is discretized into $K-1$ bins and the fitness landscape is normalized, $\mu_{i0} < 1/(K-1)$ so that the population does not become extinct when a uniform fitness profile is imposed. The fitness landscape, not the basal death rate, should primarily determine the probability of GC survival. It is known from clinical data on HIV that the initially activated B cells need to acquire many mutations in order to become bnAbs [40, 41]. Thus, even when bnAbs evolve naturally, it usually takes a long time. In our coarse-grained minimal model, this observation translates to a low mutation rate between bins. This is because each mutation in our model corresponds to many mutations at the residue level. Also, breadth-enhancing mutations are far less likely to occur than those that reduce breadth [29]. In biological terms, this is because B cells are more likely to mutate away from the small number of sequences that have high affinity for the conserved residues than to mutate towards them. In other words, the sequence entropy of high-breadth B cells is lower than that of low-breadth cells; as a consequence, an entropic force pushes B cells to mutate outward from low (high breadth) to high (low breadth) sequence entropy states.

An initial population of 50 cells is sampled from the occupancy distribution (black, Fig. 2(b)). These are the B cells that were activated by the simple immunogen that selects for B cells that bind to the region of the virus' spike that includes the conserved residues [36, 42]. As noted earlier, they do not lead to bnAbs, which is why the initial occupancy distribution of B cells is chosen as shown in Fig. 2(b). A fitness landscape is then imposed on this B cell population (blue, Fig. 2(a)). As a consequence, B cell occupancy for all GCs that did not go extinct shifts from low to higher breadth bins (blue, Fig. 2(b)). We run 100 GC simulations, and obtain stochastic trajectories of the reactions derived from Eq (2) until either stop condition is met: $N_{total} = 0$ or 200. From each surviving GC, 50 new memory B cells are sampled from the 200 that exit after prime. Corresponding to the second immunogen, a new fitness profile is imposed and this B cell population then evolves stochastically until either stop condition is met. The procedure can be repeated for subsequent immunizations, and the total number of B cells occupying the highest breadth state ($i = K/2$) summed over all surviving GCs is equivalent to the total bnAb titers produced by the immunization protocol. For the results that are shown, $K = 16$, $\mu_{i0} = 0.02$, and $\mu_{ij} = 0.05$. The mutation rates $\mu_{i,i+1} = 0.125\mu_{ij}$ and $\mu_{i+1,i} = 0.875\mu_{ij}$ if $i < K/2$. Conditions flip for $i > K/2$

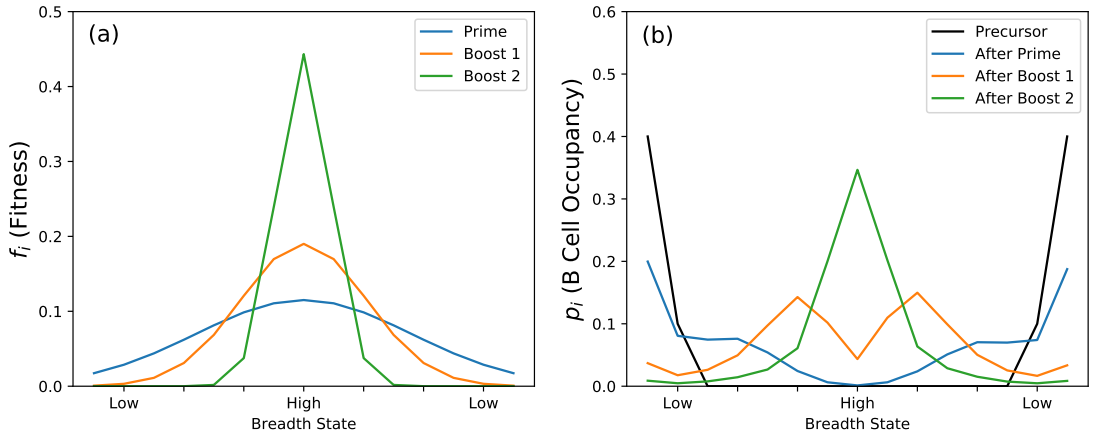

FIG. 2: Fitness landscape and response of the B cell population. (a) Imposed fitness landscape (f_i) during prime (blue), boost 1 (orange), and boost 2 (green). (b) Response of B cell population distribution p_i to fitness landscapes shown in (a).

and for $i = K/2$, $\mu_{i,i+1} = \mu_{i,i-1} = 0.5\mu_{ij}$.

III. RESULTS

A. Optimal protocols for the first and second immunizations

We will focus here on studying vaccination protocols with two sequentially administered immunogens. Following standard terminology from vaccination, we will refer to the first as "prime", and to the second as "boost". There is a large continuous space of choices for immunogens, or possible fitness landscapes, for prime and boost. We first asked if there is an optimal combination of the imposed fitness landscapes, f_1 and f_2, respectively, that maximizes bnAb production. Since the fitness landscape corresponding to the boost immunogen acts on the B cell distribution that evolved in response to the fitness profile corresponding to the priming immunogen, it is clear that optimal f_2 will depend on the choice of f_1. Using the previously described parameter values, we first searched over the space of variances in f_1 profiles that are imposed on the B cell population (p_0) that developed after activation of the right germline B cells. We then searched over the space of variances of f_2 profiles that can be imposed on p_1, the distribution that evolves after prime immunization. Our results (Fig. 3) show that an optimal setting of f_1 exists (yellow line), occurring at a KL distance $D(p_0||f_1) = 2.76$.

For each choice of f_1, the number of bnAbs produced after boost is also graphed as a function of $D(p_1||f_2)$, the distance between the boost fitness f_2 and the B cell distribution that results after prime p_1 (Fig. 3(a)). In all cases, there is also an optimal setting of $D(p_1||f_2)$, which exceeds $D(p_0||f_1)$. From the standpoint of vaccination, the last result implies that boost immunization needs to more aggressively focus selection on high breadth B cells than prime in order to maximize production of bnAbs. This result is consistent with what has been found with more elaborate computer simulations [17]. Fig. 3(b) shows the number of bnAbs produced at the maximal points on the curves calculated in Fig. 3(a). Clearly, an optimal pair of $D(p_0||f_1)$ and $D(p_1||f_2)$ exists.

At these optimal points, we computed the probability of GC survival; i.e., the probability that the B cell population is not extinguished during AM. Fig. 4(a) shows that P(GC Survival) = 1.0 for low values of $D(p_0||f_1)$ during prime immunization. At the optimal point ($D(p_0||f_1) = 2.76$), P(GC Survival) drops to ~ 0.9. Beyond the optimal value of f_1, there is a sharp drop in GC survival probability (Fig. 4(a)) and bnAb titers (Fig. 3(b)). Fig. 4(b) provides an explanation for the abrupt onset of GC death past the optimal point. For low KL distances ($D(p_0||f_1) \leq 2.56$), the fitness of all B cell breadth states exceeds the intrinsic death rate $\mu_{i0} = 0.02$ (blue dotted line, Fig. 4(b)). As a result, B cells sampled from the precursor population (black, Fig. 2(b)) that predominantly occupy low breadth states proliferate quickly, leading to GC survival. As these B cells rapidly internalize antigen, N quickly reaches a value of 200, and AM ends. Thus, there is limited time for mutations that allow B cells to transition from low to high breadth states. For $D(p_0||f_1) > 2.76$, we observe that fitness in the lowest breadth states drops below the death rate. Furthermore, the occupancy of B cells in states adjacent to the bnAb state is low. So, very few of these B cells multiply, and most mutations that arise during proliferation are transitions to lower breadth states. This mutational flux to low breadth leads to accumulation of B cells in states at the edge where they die. This is the reason that there is a sharp increase in the likelihood of extinction events. At $D(p_0||f_1) = 2.76$, the fitness of the lowest breadth states (yellow, Fig. 4(b)) becomes commensurate with the death rate. But, the B cells in states further from the edge states have reasonable fitness

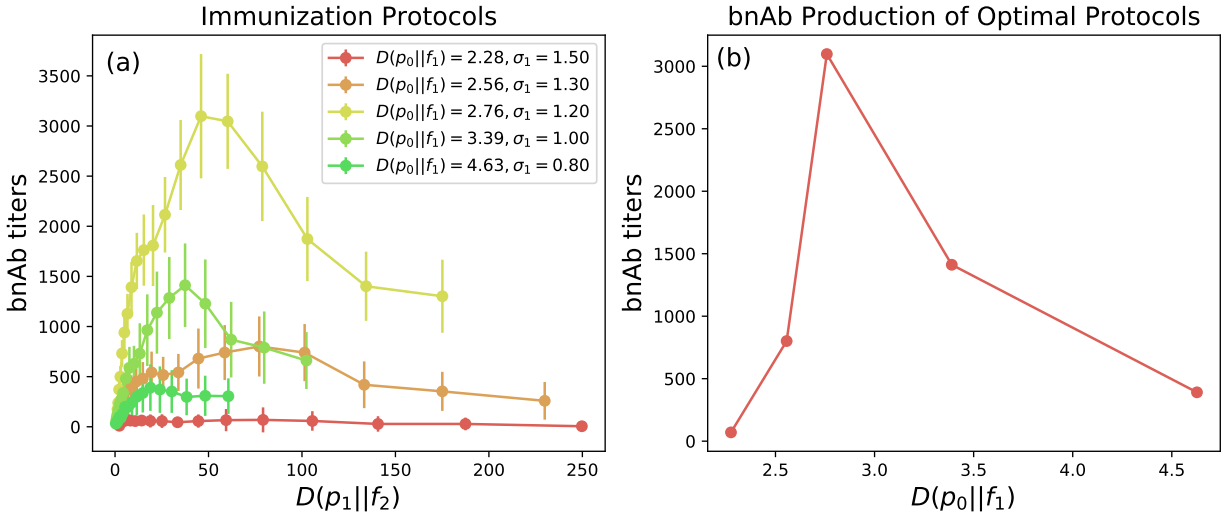

FIG. 3: Response to different vaccination protocols. (a) For each prime immunization choice of KL distance, $D(p_0||f_1)$ (colored curves), the number of bnAbs produced is shown over a range of boost immunization choices, $D(p_1||f_2)$. (b) Graphing maximal bnAb titers for each of the curves shown in (a) versus $D(p_0||f_1)$ shows that an optimal prime-boost protocol exists at $D(p_0||f_1) \sim 2.76$.

and can replicate. Some of these B cells mutate toward higher breadth states and multiply more, while others transition to low breadth states and die. This leads to a balance of these fluxes that causes neither rapid extinction nor rapid proliferation that could end the AM process. Under these conditions, GC reactions continue for some time, enabling the B cell population to acquire the many mutations required to become bnAbs.

B. Key features of evolutionary trajectories for optimal prime-boost combinations

After prime immunization, as the B cell population evolves in response to the imposed fitness distribution, it acquires information about the landscape. The initial population is mostly comprised of B cells with low breadth. As the population acquires mutations that confer higher breadth, it diversifies to occupy bins corresponding to higher breadth states. Thus, the entropy of the B cell distribution increases. Fig. 5(a) shows a histogram of the entropy change ($\Delta H(p)$) of the B cell distribution within all GCs. At low $D(p_0||f_1)$, all GCs survive but produce B cell populations dominated by low breadth cells (orange, Fig. 5(b)). As a result, there is a minimal increase in B cell diversity and entropy.

Optimal $D(p_0||f_1)$ (yellow, Fig. 5(a)) leads to larger changes in entropy or B cell diversity within most of the GCs. This indicates greater occupancy of higher breadth states in GCs (yellow, Fig. 5(b)) with respect to the initial B cell population distribution (black, Fig. 5(b)). A small number of GCs ($\sim 10\%$) die and experience entropy loss (yellow, Fig. 5(a)), indicating that those B cell populations have lost all information about the antigenic environment. If $D(p_0||f_1)$ (green, Fig. 5(a)) is too large, high rates of GC extinction result in entropy loss within most of the GCs ($\sim 80\%$). The small number of GCs that survive do so because they manage to produce high breadth B cells by stochastic chance.

In order to understand the optimal choice of f_2 given f_1, for each B cell that exits as a bnAb after optimal boost, we computed the breadth state of its ancestral B cell that initially seeded the GC at the beginning of boost. Fig. 6 shows the number of bnAb trajectories generated given that they originated from a starting sequence of breadth i after the prime. The optimal $D(p_1||f_2)$ following low $D(p_0||f_1)$ must be large in order to generate bnAb trajectories from high breadth states (orange, Fig. 6) since occupancy within those states is very low after the prime (orange, Fig. 5(b)). For high $D(p_0||f_1)$, the optimal $D(p_1||f_2)$ is small since the few GCs that survive prime mostly contain high breadth B cells (green, Fig. 5(b)). The best choice of $D(p_1||f_2)$ following the optimal setting of $D(p_0||f_1)$ falls between the optimal values of $D(p_1||f_2)$ for high and low values of $D(p_0||f_1)$; i.e., $D(p_1||f_2)_{high\ prime} < D(p_1||f_2)_{optimal\ prime} < D(p_1||f_2)_{low\ prime}$. The optimal setting for the boost successfully proliferates bnAb trajectories from several states near the highest breadth state (e.g. $i = 5 - 7$ or $i = 9 - 11$) that are populated by selection forces imposed by the optimal setting for the prime. The optimal priming immunogen generates the right kind of B cell diversity, which makes possible many high probability evolutionary trajectories that mature into bnAbs. These trajectories ensue upon imposing the fitness landscape corresponding to the optimal boost.

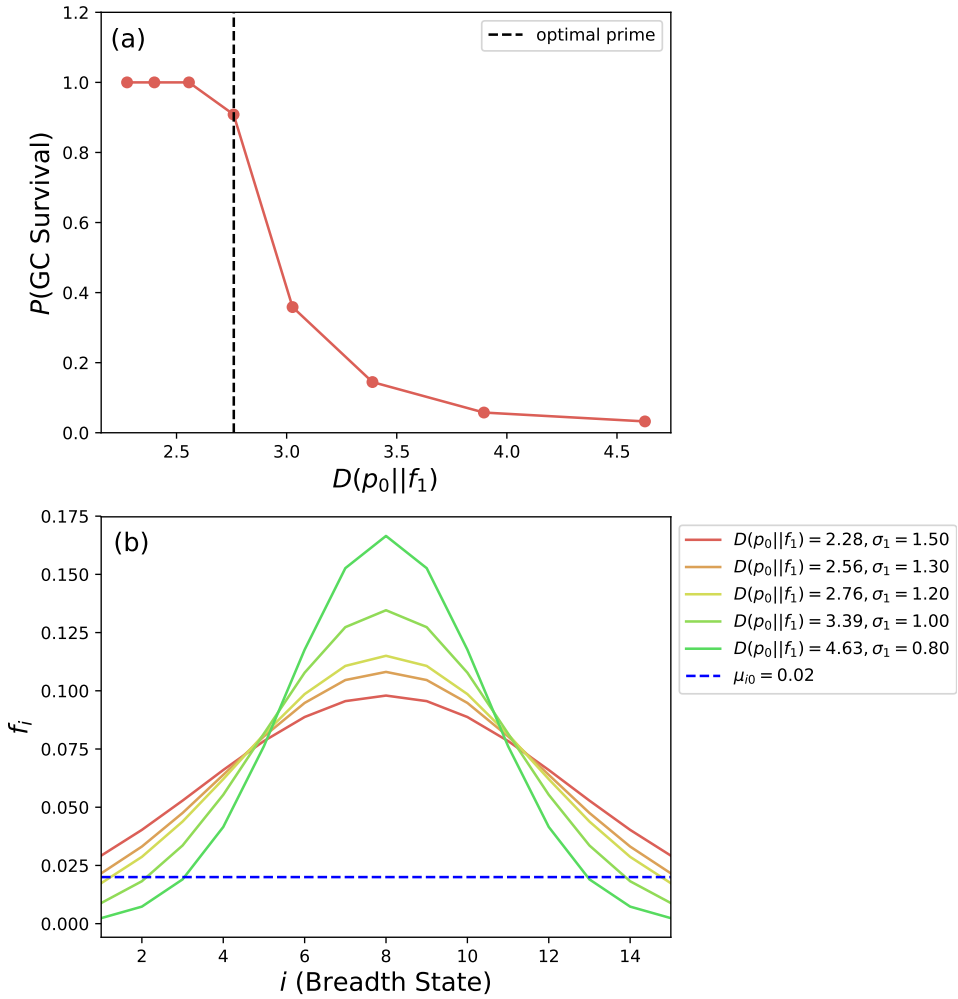

FIG. 4: Germinal center (GC) survival (B cell population is not extinguished). (a) Fraction of GCs that survive prime immunization for a range of values of $D(p_0||f_1)$. The fraction falls to ~ 0.9 at the optimal setting and drops significantly beyond this point. (b) Sample fitness landscapes at the points located on the curve shown in (a).

C. Phylogenetic analyses of evolving B cell populations

The phylogenetic trees shown in Figs. 7-9 tell a deeper story about the three possible immunization protocols analyzed in Figs. 5-6. For each of the 50 B cells that initially seed the GC, we compute the birth-death-mutation trajectories that emerge.

At low $D(p_0||f_1)$, most of the initial crop of B cells manage to seed trajectories that survive until the end of prime (Fig. 7(a), yellow region). Yet, since selection is favorable for low breadth, these trajectories predominantly generate low breadth B cells; when a high $D(p_1||f_2)$ is set during boost (see Fig. 7(a), pink region), nearly all of the trajectories sampled from these cells die. In rare cases, one of the precursor trajectories will stochastically manage to acquire enough breadth-enhancing mutations to generate high breadth cells before the end of prime. The top panel in Fig. 7(b) is an expanded view of the green box shown in Fig. 7(a). One of the trajectories leads to production of 6 high breadth B cells (green circles), which get selected at the beginning of boost and eventually lead to the proliferation of 124 bnAbs (orange circles, Fig. 7(b), bottom panel, an expanded view of the purple box in Fig. 7(a)). These rare events allow the GC to survive when selection for binding of conserved epitopes is low during prime and high through boost.

The opposite effect is observed at high $D(p_0||f_1)$. Under strong selection for high breadth, most trajectories generated by the precursor B cells die quickly during prime (see Fig. 8(a), yellow region). Yet, due to rare stochastic fluctuations, one of the precursor trajectories may generate a B cell of sufficiently high breadth that under strong selection generates a burst of replication and branching events (Fig. 8(a)); in the top panel of Fig. 8(b), these bursting events culminate in the production of 15 high breadth B cells (yellow circles) at the end of prime. During boost, these cells rapidly replicate and proliferate 139 bnAbs (orange circles, Fig. 8(b), bottom panel). Thus, low and high $D(p_0||f_1)$ protocols can in-

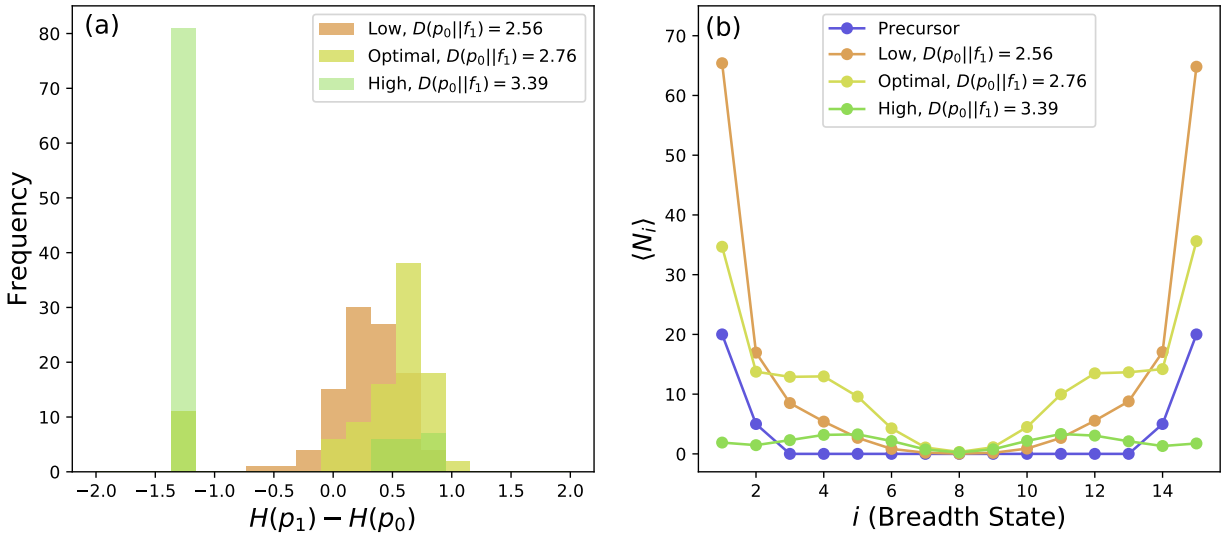

FIG. 5: Optimal diversity of the B cell population. (a) For each GC, the change in entropy or diversity of the B cell distribution after prime immunization is recorded and graphed as a histogram for low (orange), optimal (yellow), and high (green) $D(p_0||f_1)$. The GCs wherein the B cell population is extinguished experience a large entropy loss ($\Delta H \sim -1.2$) (b) The entropy increase is a consequence of B cells on average populating higher breadth states after prime immunization.

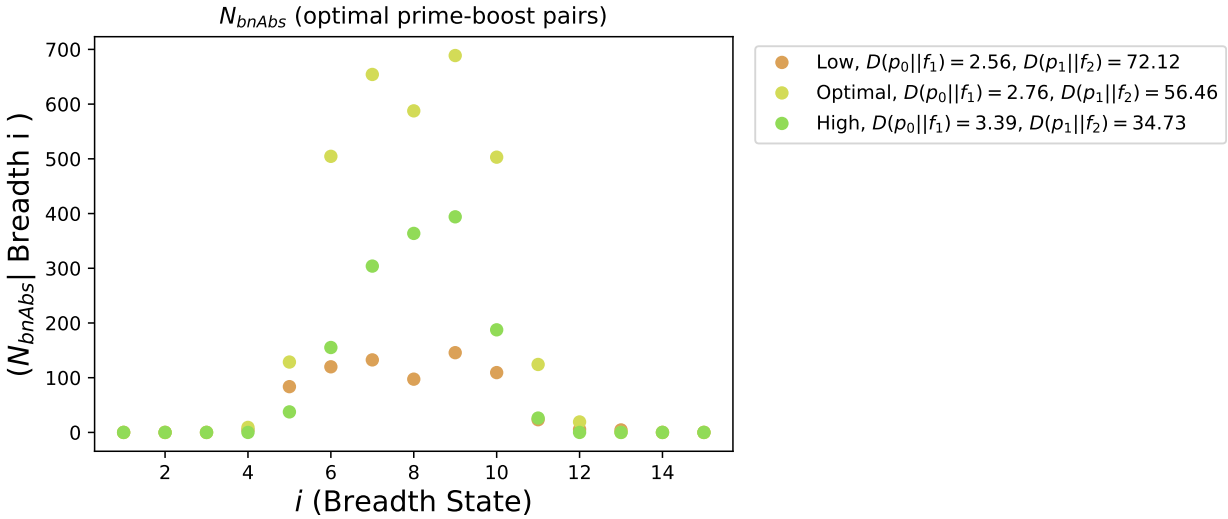

FIG. 6: Characterizing optimal protocols. For the three prime protocols shown in Fig. 5, the numbers of evolutionary trajectories that exit the GCs as bnAbs after optimal boost are graphed as a function of the breadth state of the B cells that initially seed the trajectories at the beginning of boost.

duce bnAb production through rare events; the dynamics of these events differ depending on the strength of selection pressure during prime.

Phylogenetic trees for the optimal prime-boost protocol (Fig. 9(a)) present features which differ dramatically from the sub-optimal trees. Most noticeably, the tree exhibits considerably greater complexity with the presence of highly dense replication and mutation events. Due to the flux balance that prevents extinction and rapid proliferation of the population, GC lifetime during prime lasts significantly longer than for sub-optimal immunization. Thus, the precursor B cells follow longer trajectories that search breadth space more effectively. As a consequence, the probability of generating large numbers of higher breadth B cells is greatly enhanced. The top panel of Fig. 9(b) shows that, at the end of prime, a large number of high breadth sequences ($n \sim 35$, light purple circles) is produced. It is highly likely that at the beginning of boost (top panel, Fig. 9(b)), there is at least one B cell sufficiently close to the bnAb state (yellow circle) which eventually seeds the rapid proliferation of bnAbs by the end of boost (bottom panel, Fig. 9(b)).

An additional notable feature is evidence of clonal interference. For the GC trajectory shown in Fig. 9, the 50 initial B cells have the following distribution: $N_1 = 20$, $N_2 = 4$, $N_{14} = 9$, and $N_{15} = 17$. Interestingly, a single breadth 15 B cell (red box, Fig. 9(a)) manages to generate progeny that acquire breadth-enhancing mutations sooner than any of the other initially higher breadth cells. As a result, significant clonal expansion of this B cell lineage outcompetes trajectories produced by all other precursor cells. Eventually, all of the B cells that exit prime and seed the GC during boost originate from this single precursor cell (blue box, Fig. 9(a)). The exhaustive search process during optimal prime immunization leads to occupancy of high breadth states that maximize the likelihood of proliferating trajectories which eventually find the bnAb state during boost (bottom panel, Fig. 9(b)).

IV. CONCLUSIONS

Developing effective vaccines and immunization protocols that could induce the production of antibodies that neutralize diverse strains of highly mutable viruses is a major global health challenge. To date, a vaccination strategy that can generate such bnAbs against viruses like HIV or influenza have not been developed. Antibodies are produced by the Darwinian evolutionary process of AM. The immunogens used in a vaccination strategy impose selection forces on the B cell population. Generating bnAbs requires that the selection forces are such that generalists that can neutralize diverse mutant strains of the virus, rather than specialists that neutralize specific strains, evolve. In this paper, we combined stochastic simulations of a minimal model for AM with an information theoretic metric to study how B cells acquire information during AM induced by different selection forces (immunogens). In particular, we strove to understand the mechanistic reasons underlying why certain vaccination protocols are optimal for generating bnAbs.

We studied a process where two immunizations were allowed, a prime and a boost. Our principal findings can be summarized as follows. The KL distance between the existing B cell population and the fitness landscape imposed by an immunogen is a measure of the maximum information that the B cells can acquire by evolution during AM. The KL distance can also be thought of as the extent to which the existing B cell population is placed out of equilibrium by a new immunogen. We find that there is an optimal combination of KL distances during prime and boost that maximizes the generation of bnAbs. For a given KL distance that defines the prime, the optimal KL distance for the boost is always higher. Our results are consistent with previous analysis in regards to optimal prime-boost immunization strategies. The degree of conflicting selection forces [17] or the cycling frequency between antigenic environments [15] must be low early in AM and then increased over time to ensure survival of the B cell population and maximize production of bnAbs. Here, the KL distance also quantifies the degree of non-equilibrium perturbation and information gain with respect to the current state of the B cell population.

If the KL distance is too small during the prime, most B cell lineages successfully survive the prime, but they do not gain much information about the bnAb state. These lineages mostly die during the boost, which has to be very aggressive in order for any bnAbs to evolve. The few bnAbs that do evolve are the product of rare stochastic trajectories that acquire the right breadth-enhancing mutations to become nearly bnAbs during the prime. These are then rapidly expanded during the boost. If the KL distance during the prime is too large, most B cell lineages become extinct because their fitness is very low. Again, the B cell population does not gain information about the bnAb state. The few bnAbs that ultimately evolve are again the product of rare lucky trajectories that survive strong selection pressure during prime and access the bnAb state during the boost. In the latter case, the dynamics are ballistic-like as shown by the phylogenetic trees (Fig. 8(a)). The optimal priming condition sets the B cell population off-equilibrium sufficiently that B cells must evolve to higher breadth states to proliferate effectively, but the population does not become extinct with high probability. This balanced selection force during the prime results in complex evolutionary trajectories that generate B cells in states close to the bnAb state. During the boost, these lineages can evolve into bnAbs. In the process, there is evidence of clonal interference. Therefore, the optimal prime results in the B cell population acquiring information about the bnAb state and the right kind of diversity of lineages is generated.

Recently, it has been shown that phylogenetic trees exhibit different topological asymptotic scaling laws which depend on whether they are balanced or unbalanced [43]. Upon visual inspection of Figs. 7-9, it is clear that the trees produced by the optimal protocol exhibit a vastly different topology from those constructed by sub-optimal protocols. A scaling analysis could reveal that optimal and sub-optimal protocols lead to different asymptotic behaviors with respect to phylogeny, a result that could reflect a complex interplay between ecological and evolutionary processes during AM [44].

We hope our results will guide more detailed simulations and experiments designed to generate immunogens that can impose an optimal prime/boost protocol. We also hope that our work will motivate theoretical studies on how evolutionary forces can select for generalists in realistic conditions that are not periodic oscillations in environments.

[1] L. Mesin, J. Ersching, and G. D. Victora, Germinal center b cell dynamics, Immunity **45**, 471 (2016).

[2] S. Wang, J. Mata-Fink, B. Kriegsman, M. Hanson, D. J. Irvine, H. N. Eisen, D. R. Burton, K. D. Wittrup, M. Kardar, and A. K. Chakraborty, Manipulating the selection forces during affinity maturation to generate cross-reactive hiv antibodies, Cell **160**, 785 (2015).

[3] J. Zhang and E. I. Shakhnovich, Optimality of mutation and selection in germinal centers, PLoS computational biology **6** (2010).

[4] H. N. Eisen and G. W. Siskind, Variations in affinities of antibodies during the immune response, Biochemistry **3**, 996 (1964).

[5] A. K. Abbas, A. H. Lichtman, and S. Pillai, *Basic Immunology E-Book: Functions and Disorders of the Immune System* (Elsevier Health Sciences, 2019).

[6] D. Fera, A. G. Schmidt, B. F. Haynes, F. Gao, H.-X. Liao, T. B. Kepler, and S. C. Harrison, Affinity maturation in an hiv broadly neutralizing b-cell lineage through reorientation of variable domains, Proceedings of the National Academy of Sciences **111**, 10275 (2014).

[7] J.-P. Julien, A. Cupo, D. Sok, R. L. Stanfield, D. Lyumkis, M. C. Deller, P.-J. Klasse, D. R. Burton, R. W. Sanders, J. P. Moore, *et al.*, Crystal structure of a soluble cleaved hiv-1 envelope trimer, Science **342**, 1477 (2013).

[8] D. Lyumkis, J.-P. Julien, N. de Val, A. Cupo, C. S. Potter, P.-J. Klasse, D. R. Burton, R. W. Sanders, J. P. Moore, B. Carragher, *et al.*, Cryo-em structure of a fully glycosylated soluble cleaved hiv-1 envelope trimer, Science **342**, 1484 (2013).

[9] D. D. Raymond, G. Bajic, J. Ferdman, P. Suphaphiphat, E. C. Settembre, M. A. Moody, A. G. Schmidt, and S. C. Harrison, Conserved epitope on influenza-virus hemagglutinin head defined by a vaccine-induced antibody, Proceedings of the National Academy of Sciences **115**, 168 (2018).

[10] D. R. Burton and L. Hangartner, Broadly neutralizing antibodies to hiv and their role in vaccine design, Annual review of immunology **34**, 635 (2016).

[11] B. F. Haynes, G. Kelsoe, S. C. Harrison, and T. B. Kepler, B-cell–lineage immunogen design in vaccine development with hiv-1 as a case study, Nature biotechnology **30**, 423 (2012).

[12] B. F. Haynes and D. R. Burton, Developing an hiv vaccine, Science **355**, 1129 (2017).

[13] X. Wu, Z.-Y. Yang, Y. Li, C.-M. Hogerkorp, W. R. Schief, M. S. Seaman, T. Zhou, S. D. Schmidt, L. Wu, L. Xu, *et al.*, Rational design of envelope identifies broadly neutralizing human monoclonal antibodies to hiv-1, Science **329**, 856 (2010).

[14] S. Wang and L. Dai, Evolving generalists in switching rugged landscapes, PLoS computational biology **15**, e1007320 (2019).

[15] V. Sachdeva, K. Husain, J. Sheng, S. Wang, and A. Murugan, Tuning environmental timescales to evolve and maintain generalists, Proceedings of the National Academy of Sciences **117**, 12693 (2020).

[16] J. S. Shaffer, P. L. Moore, M. Kardar, and A. K. Chakraborty, Optimal immunization cocktails can promote induction of broadly neutralizing abs against highly mutable pathogens, Proceedings of the National Academy of Sciences **113**, E7039 (2016).

[17] K. G. Sprenger, J. E. Louveau, P. M. Murugan, and A. K. Chakraborty, Optimizing immunization protocols to elicit broadly neutralizing antibodies, Proceedings of the National Academy of Sciences (2020).

[18] M. Oprea and A. S. Perelson, Somatic mutation leads to efficient affinity maturation when centrocytes recycle back to centroblasts., The Journal of Immunology **158**, 5155 (1997).

[19] S. Chaudhury, J. Reifman, and A. Wallqvist, Simulation of b cell affinity maturation explains enhanced antibody cross-reactivity induced by the polyvalent malaria vaccine ama1, The Journal of Immunology **193**, 2073 (2014).

[20] L. M. Childs, E. B. Baskerville, and S. Cobey, Trade-offs in antibody repertoires to complex antigens, Philosophical Transactions of the Royal Society B: Biological Sciences **370**, 20140245 (2015).

[21] A. Escolano, J. M. Steichen, P. Dosenovic, D. W. Kulp, J. Golijanin, D. Sok, N. T. Freund, A. D. Gitlin, T. Oliveira, T. Araki, *et al.*, Sequential immunization elicits broadly neutralizing anti-hiv-1 antibodies in ig knockin mice, Cell **166**, 1445 (2016).

[22] A. Nourmohammad, J. Otwinowski, and J. B. Plotkin, Host-pathogen coevolution and the emergence of broadly neutralizing antibodies in chronic infections, PLoS genetics **12**, e1006171 (2016).

[23] E. Kussell and S. Leibler, Phenotypic diversity, population growth, and information in fluctuating environments, Science **309**, 2075 (2005).

[24] S. Leibler and E. Kussell, Individual histories and selection in heterogeneous populations, Proceedings of the National Academy of Sciences **107**, 13183 (2010).

[25] E. Kussell, S. Leibler, and A. Grosberg, Polymer-population mapping and localization in the space of phenotypes, Physical review letters **97**, 068101 (2006).

[26] A. S. Perelson and G. F. Oster, Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination, Journal of theoretical biology **81**, 645 (1979).

[27] D. J. Smith, S. Forrest, R. R. Hightower, and A. S. Perelson, Deriving shape space parameters from immunological data, Journal of theoretical biology **189**, 141 (1997).

[28] A. Lapedes and R. Farber, The geometry of shape space: application to influenza, Journal of theoretical biology **212**, 57 (2001).

[29] K. Wiehe, T. Bradley, R. R. Meyerhoff, C. Hart, W. B. Williams, D. Easterhoff, W. J. Faison, T. B. Kepler, K. O. Saunders, S. M. Alam, *et al.*, Functional relevance of improbable antibody mutations for hiv broadly neutralizing antibody development, Cell host & microbe **23**, 759 (2018).

[30] I. H. Moal and J. Fernández-Recio, Skempi: a structural kinetic and energetic database of mutant protein interactions and its use in empirical models, Bioinformatics **28**, 2600 (2012).

[31] C. Berek and C. Milstein, Mutation drift and repertoire shift in the maturation of the immune response, Immunological reviews **96**, 23 (1987).

[32] A. D. Gitlin, Z. Shulman, and M. C. Nussenzweig, Clonal selection in the germinal centre by regulated proliferation

and hypermutation, Nature **509**, 637 (2014).

[33] S. Crotty, A brief history of t cell help to b cells, Nature Reviews Immunology **15**, 185 (2015).

[34] D. T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, Journal of computational physics **22**, 403 (1976).

[35] D. T. Gillespie, Exact stochastic simulation of coupled chemical reactions, The journal of physical chemistry **81**, 2340 (1977).

[36] J. M. Steichen, D. W. Kulp, T. Tokatlian, A. Escolano, P. Dosenovic, R. L. Stanfield, L. E. McCoy, G. Ozorowski, X. Hu, O. Kalyuzhniy, *et al.*, Hiv vaccine design to target germline precursors of glycan-dependent broadly neutralizing antibodies, Immunity **45**, 483 (2016).

[37] T. M. Cover and J. A. Thomas, *Elements of information theory* (John Wiley & Sons, 2012).

[38] J. C. Baez and B. S. Pollard, Relative entropy in biological systems, Entropy **18**, 46 (2016).

[39] A. Nourmohammad and C. Eksin, Optimal evolutionary control for artificial selection on molecular phenotypes, arXiv preprint arXiv:1912.13433 (2019).

[40] J. F. Scheid, H. Mouquet, B. Ueberheide, R. Diskin, F. Klein, T. Y. Oliveira, J. Pietzsch, D. Fenyo, A. Abadir, K. Velinzon, *et al.*, Sequence and structural convergence of broad and potent hiv antibodies that mimic cd4 binding, Science **333**, 1633 (2011).

[41] F. Klein, R. Diskin, J. F. Scheid, C. Gaebler, H. Mouquet, I. S. Georgiev, M. Pancera, T. Zhou, R.-B. Incesu, B. Z. Fu, *et al.*, Somatic mutations of the immunoglobulin framework are generally required for broad and potent hiv-1 neutralization, Cell **153**, 126 (2013).

[42] J. Jardine, J.-P. Julien, S. Menis, T. Ota, O. Kalyuzhniy, A. McGuire, D. Sok, P.-S. Huang, S. MacPherson, M. Jones, *et al.*, Rational hiv immunogen design to target specific germline b cell receptors, Science **340**, 711 (2013).

[43] E. A. Herrada, C. J. Tessone, K. Klemm, V. M. Eguíluz, E. Hernández-García, and C. M. Duarte, Universal scaling in the branching of the tree of life, Plos one **3**, e2757 (2008).

[44] C. Xue, Z. Liu, and N. Goldenfeld, Scale-invariant topology and bursty branching of evolutionary trees emerge from niche construction, Proceedings of the National Academy of Sciences **117**, 7879 (2020).

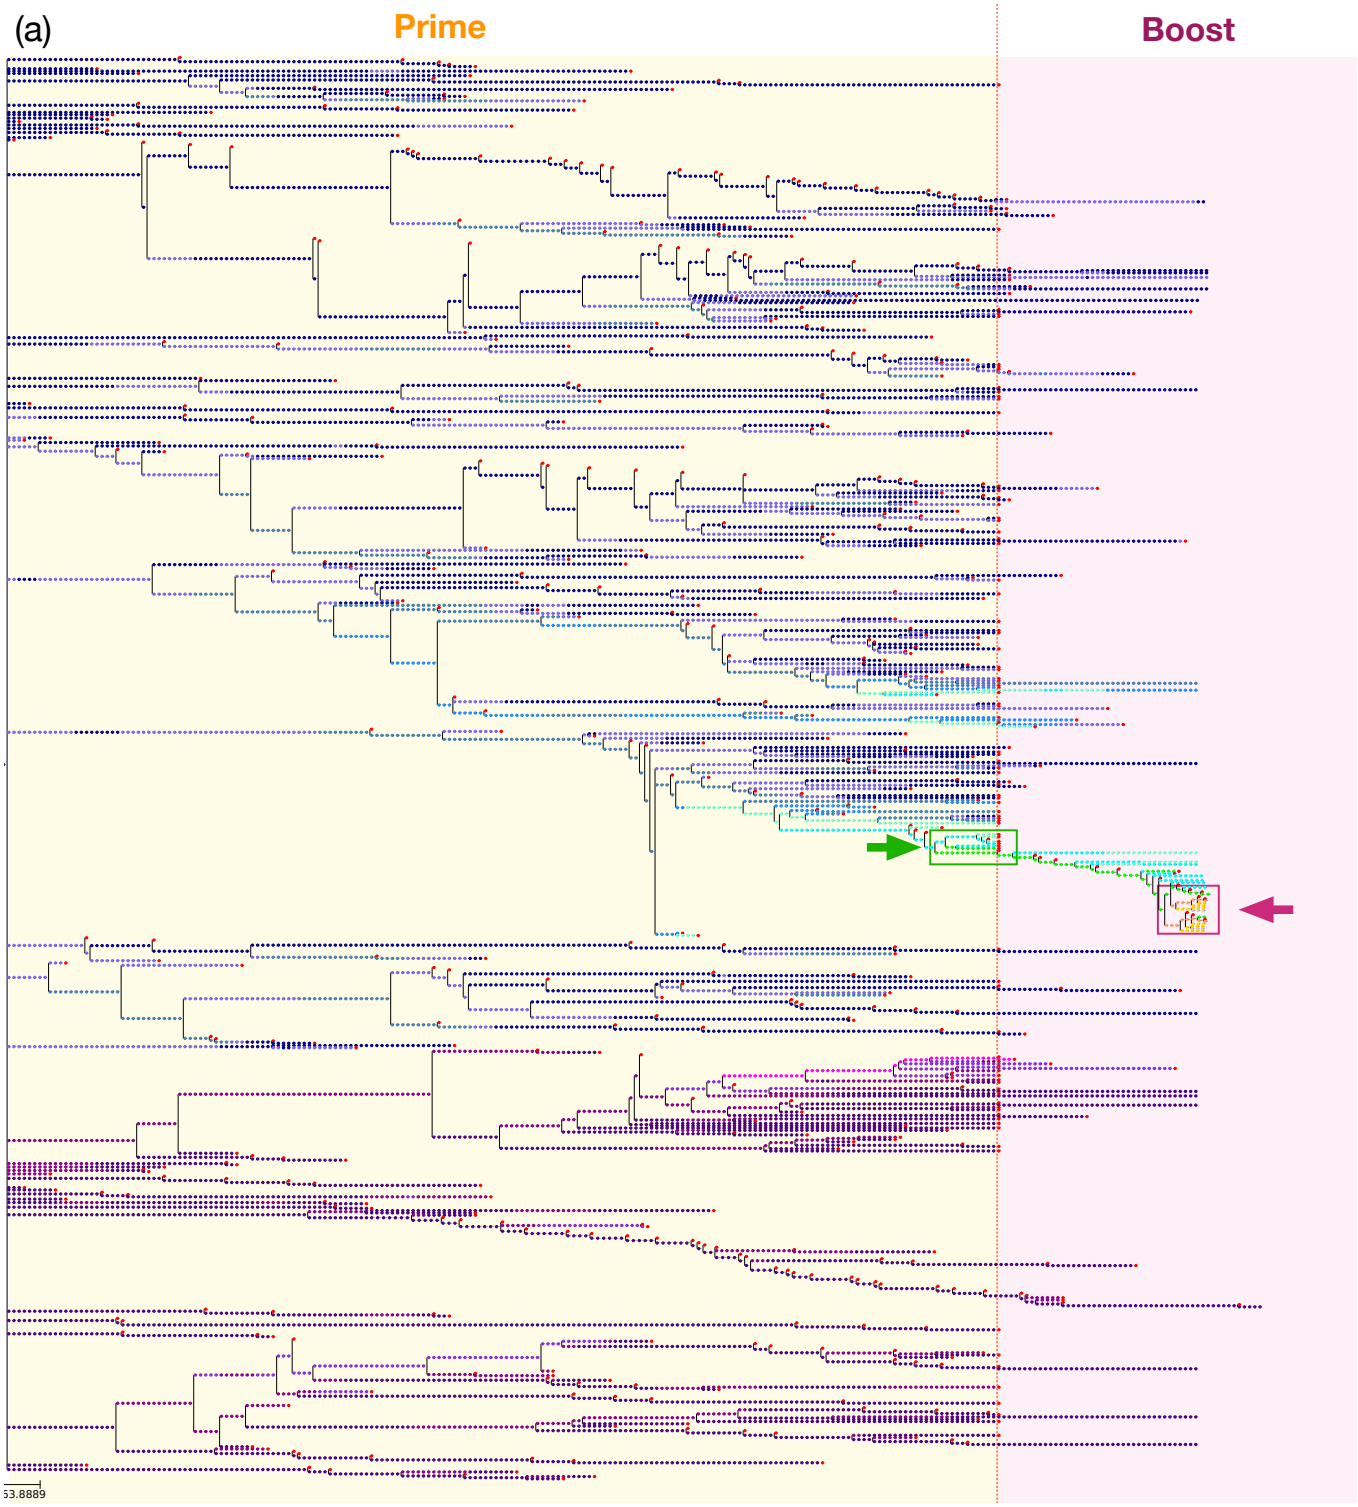
(a)
Prime
Boost
53.8889

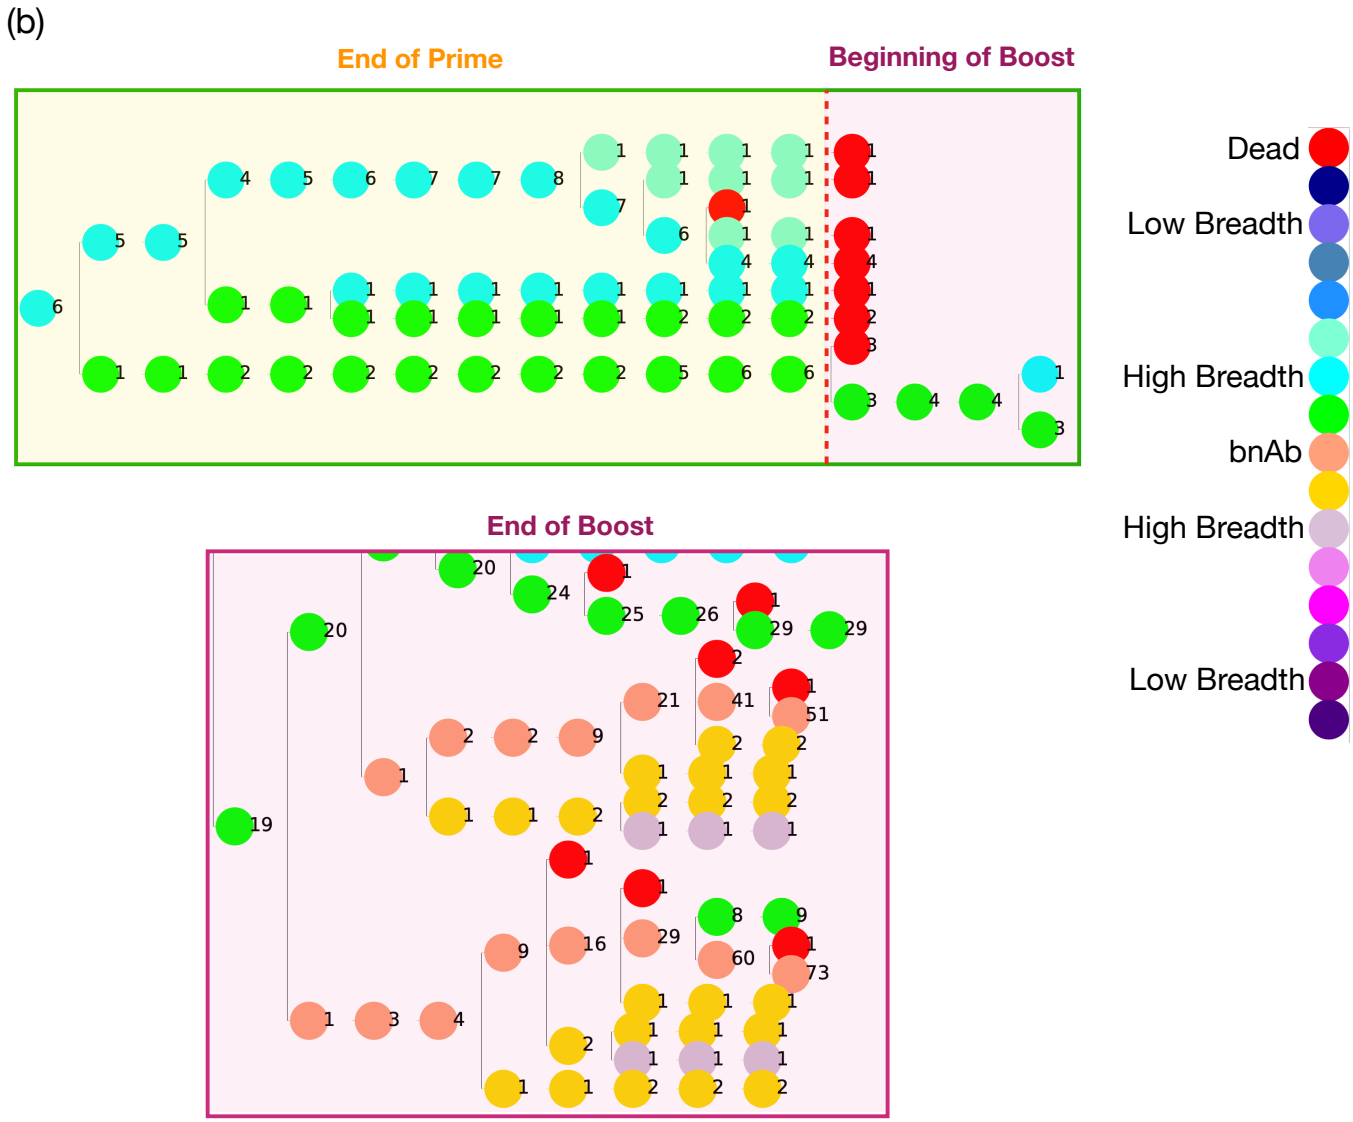

FIG. 7: Phylogeny of evolving B cell lineages for low $D(p_0||f_1)$ during prime followed by the optimal boost. (a) An example of a population of B cells evolving in a germinal center. The initial B cells were sampled from the precursor distribution (black, Fig. 2(b)) subjected to $D(p_0||f_1) = 2.56$ and $D(p_1||f_2) = 72.12$. The yellow region corresponds to the prime and the pink region to the boost. (b) The top panel is an expanded view of the green box (green arrow) in (a), which shows high breadth B cells generated at the end of prime and the beginning of boost. The bottom panel corresponds to the purple box (purple arrow) in (a), which occurs at the end of boost. Colors of circles depicting B cells denote their breadth states, and the numbers to the right of each circle quantify the number of B cells in that state.

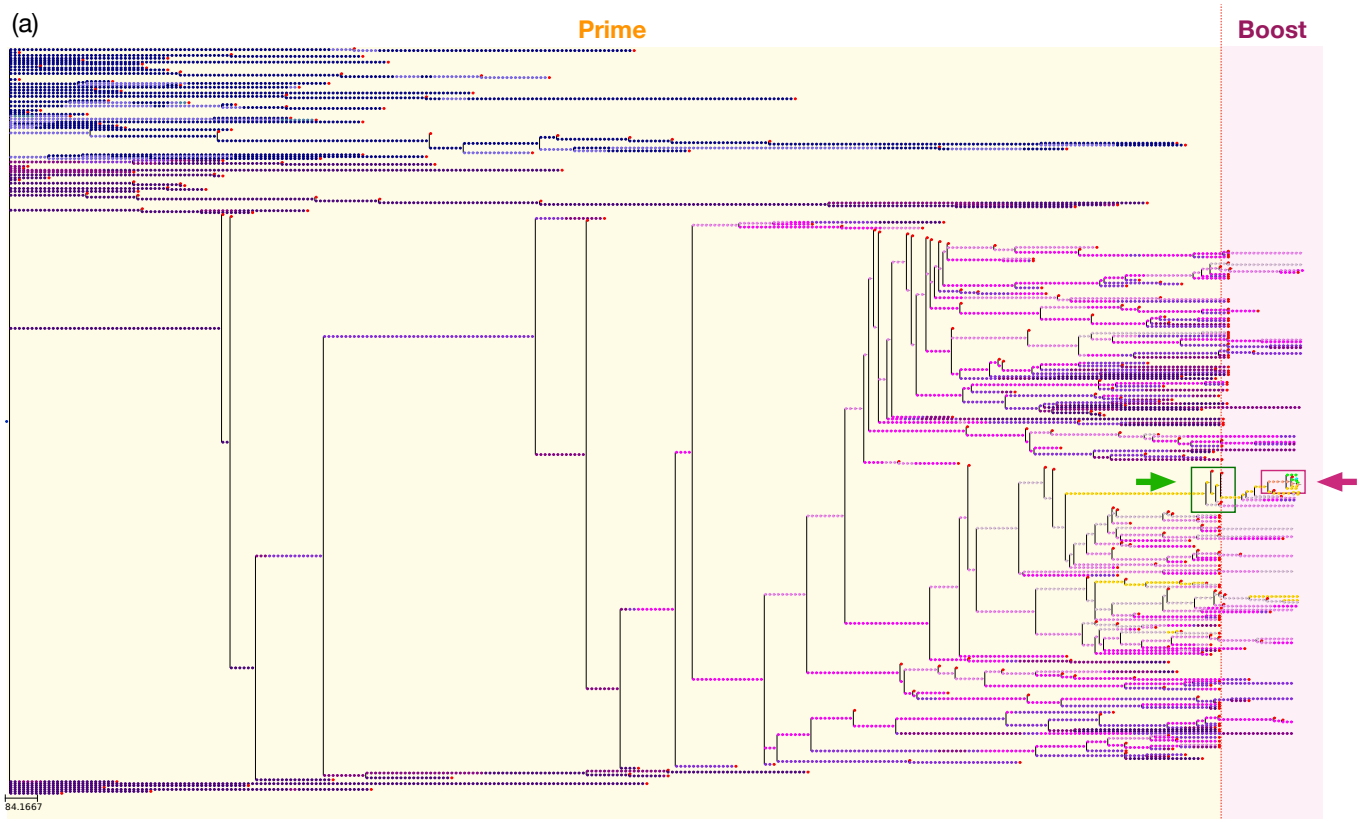
(a)
Prime
Boost
84.1667

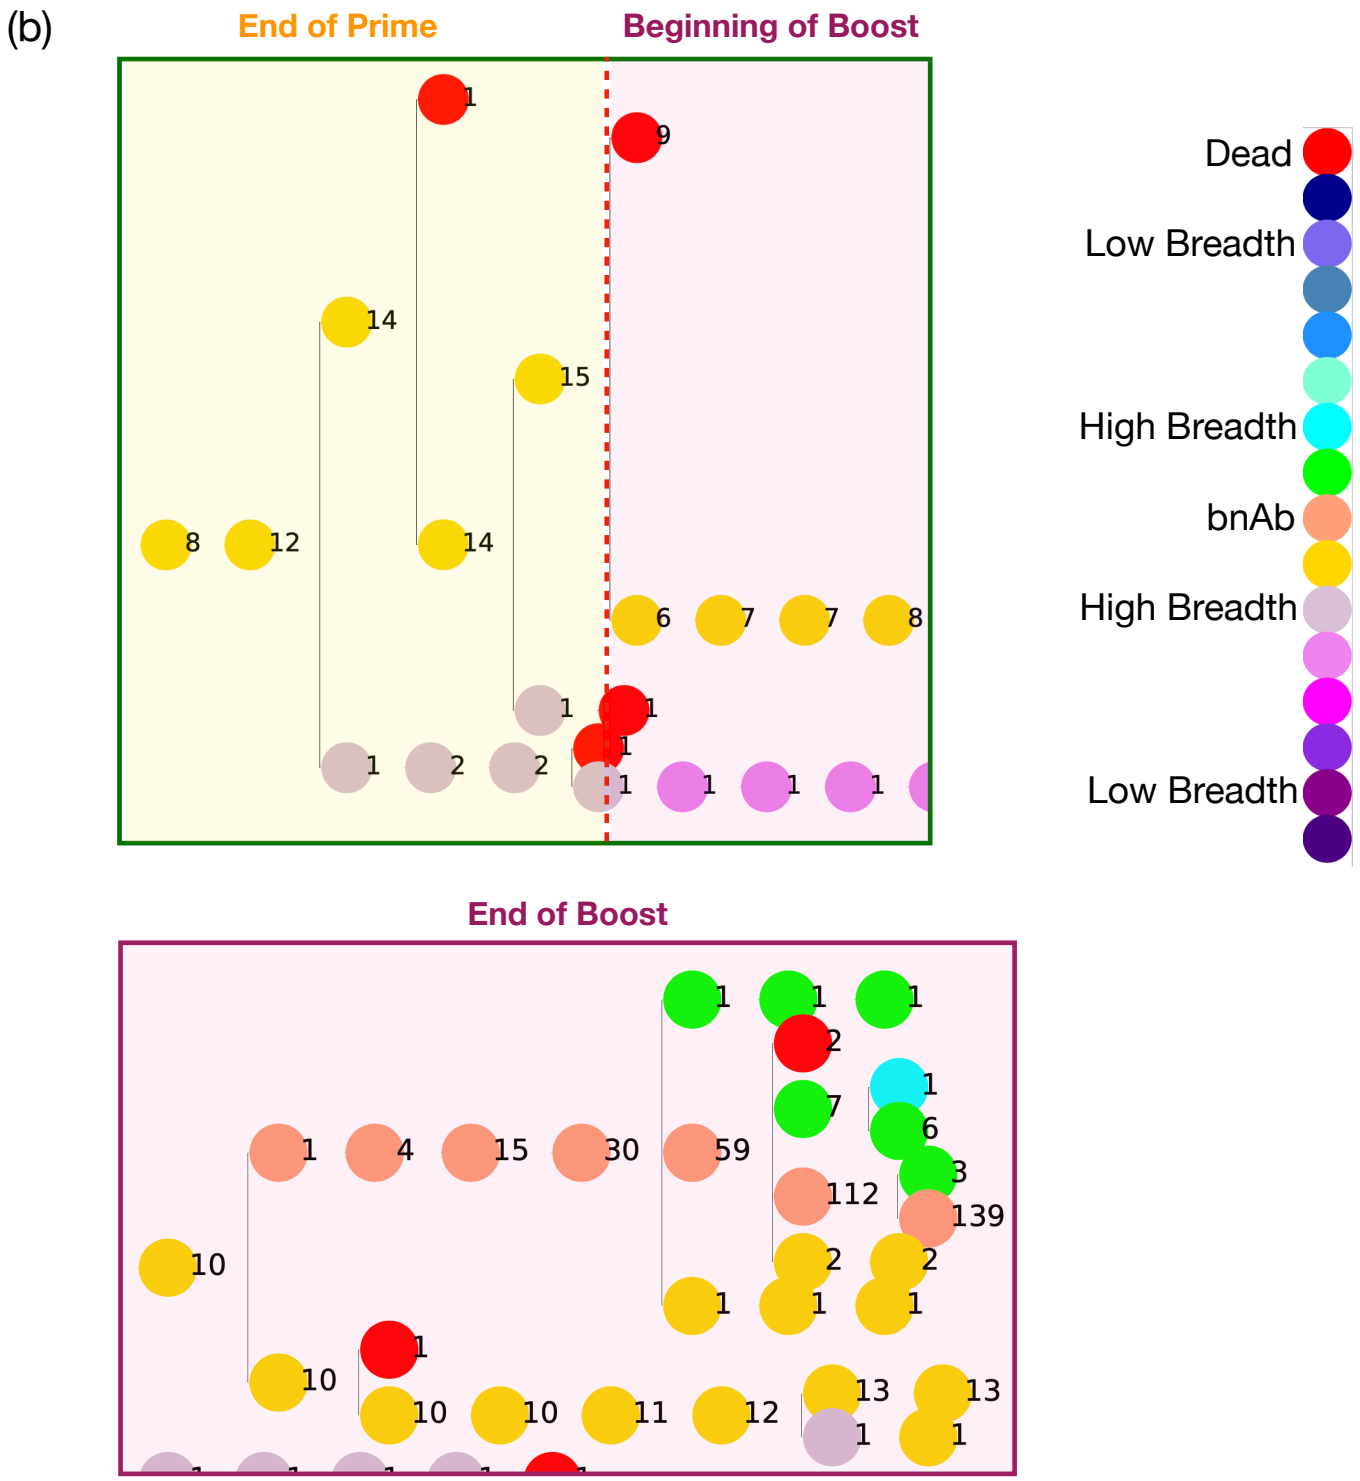

FIG. 8: Phylogeny of evolving B cell lineages for high $D(p_0||f_1)$ during prime followed by the optimal boost. (a) An example of a population of B cells evolving in a germinal center. The initial B cells were sampled from the precursor distribution (black, Fig. 2(b)) subjected to $D(p_0||f_1) = 3.39$ and $D(p_1||f_2) = 34.73$. The yellow region corresponds to the prime and the pink region to the boost. (b) The top panel is an expanded view of the green box (green arrow) in (a), which shows high breadth B cells generated at the end of prime. The bottom panel corresponds to the purple box (purple arrow) in (a), which corresponds to the end of boost. Colors of circles depicting B cells denote their breadth states, and the numbers to the right of each circle quantify the number of B cells in that state.

(a)

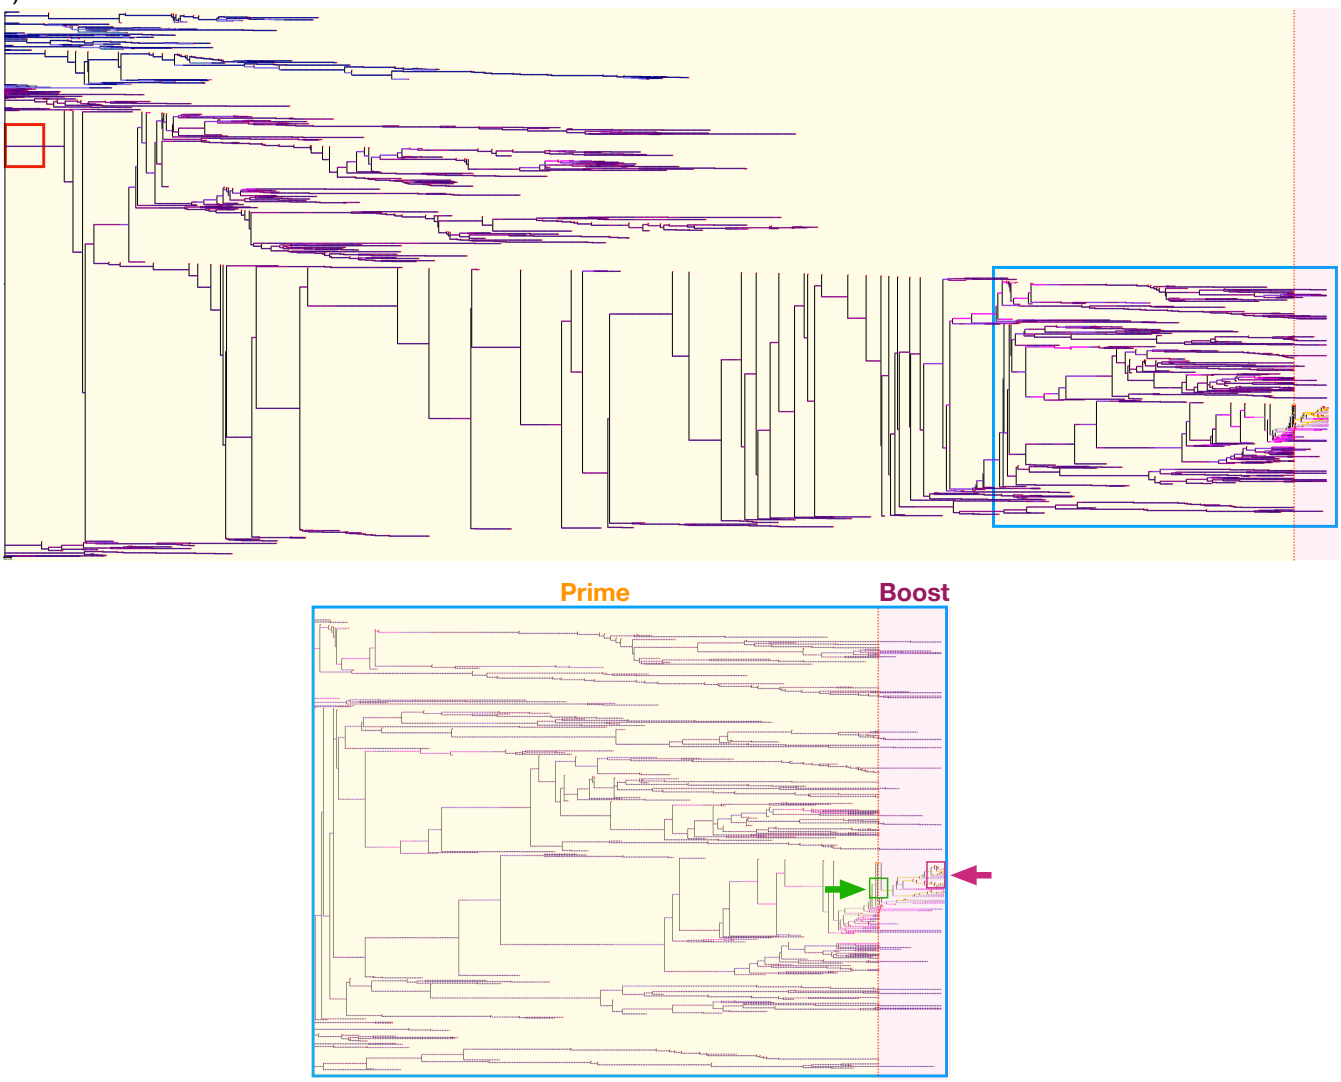

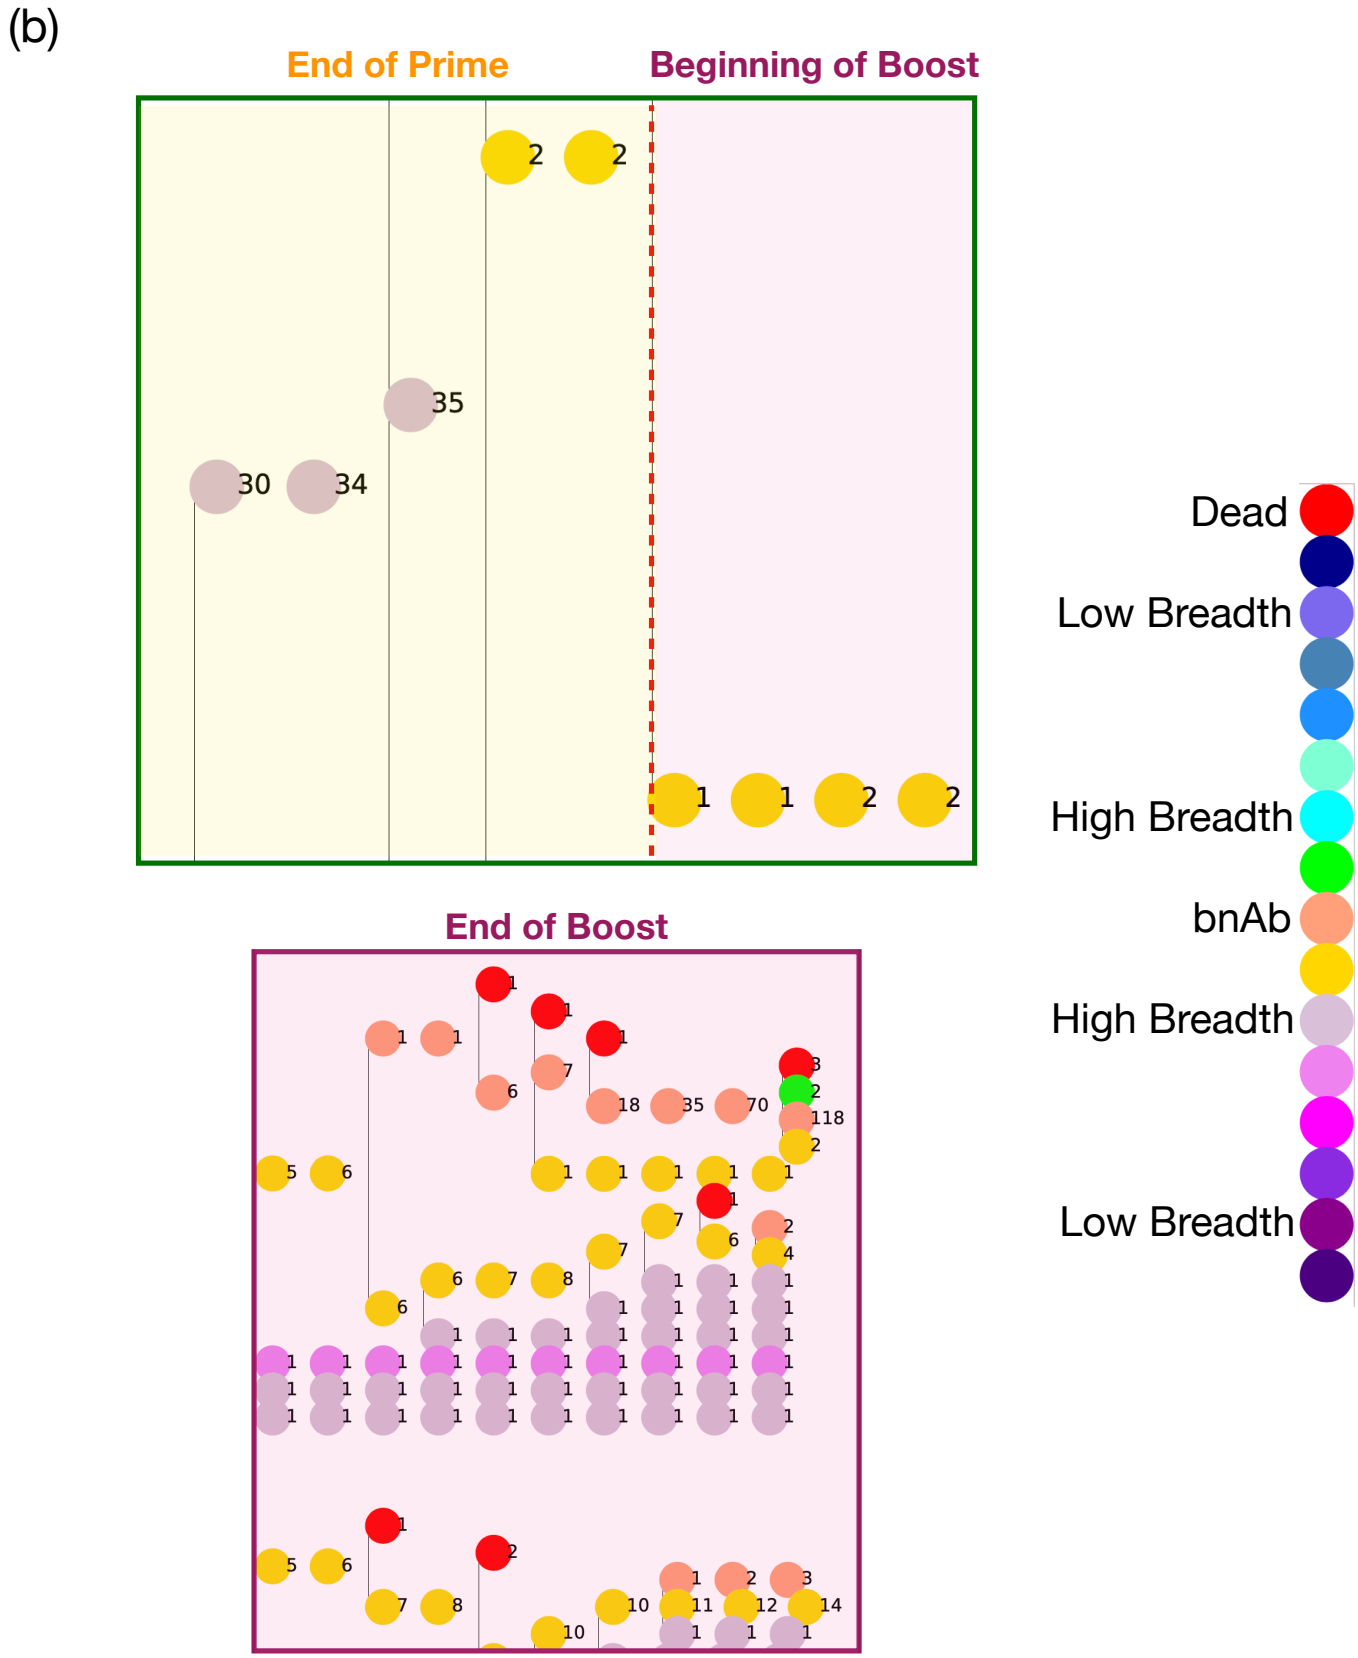

FIG. 9: Phylogeny of evolving B cell lineages for optimal $D(p_0||f_1)$ during prime followed by the optimal boost. (a) An example of a population of B cells evolving in a germinal center. The initial B cells were sampled from the precursor distribution (black, Fig. 2(b)) subjected to $D(p_0||f_1) = 2.76$ and $D(p_1||f_2) = 56.46$. The yellow region corresponds to the prime and the pink region to the boost. (b) The top panel is an expanded view of the green box (green arrow) in (a), which shows high breadth B cells generated at the end of prime. The bottom panel corresponds to the purple box (purple arrow) in (a), which corresponds to the end of boost. Colors of circles depicting B cells denote their breadth states, and the numbers to the right of each circle quantify the number of B cells in that state.