Affinity maturation for an optimal balance between long-term immune coverage and short-term resource constraints

Victor Chardès, Massimo Vergassola, Aleksandra M. Walczak,* and Thierry Mora*
Laboratoire de physique de l'École normale supérieure, CNRS, PSL University, Sorbonne Université, and Université de Paris, 75005 Paris, France

In order to target threatening pathogens, the adaptive immune system performs a continuous reorganization of its lymphocyte repertoire. Following an immune challenge, the B cell repertoire can evolve cells of increased specificity for the encountered strain. This process of affinity maturation generates a memory pool whose diversity and size remain difficult to predict. We assume that the immune system follows a strategy that maximizes the long-term immune coverage and minimizes the short-term metabolic costs associated with affinity maturation. This strategy is defined as an optimal decision process on a finite dimensional phenotypic space, where a pre-existing population of naive cells is sequentially challenged with a neutrally evolving strain. We unveil a trade-off between immune protection against future strains and the necessary reorganization of the repertoire. This plasticity of the repertoire drives the emergence of distinct regimes for the size and diversity of the memory pool, depending on the density of naive cells and on the mutation rate of the strain. The model predicts power-law distributions of clonotype sizes observed in data, and rationalizes antigenic imprinting as a strategy to minimize metabolic costs while keeping good immune protection against future strains.

I. INTRODUCTION

Adaptive immunity relies on populations of lymphocytes expressing diverse antigen-binding receptors on their surface to defend the organism against a wide variety of pathogens. B lymphocytes rely on a two-step process to produce diversity: first a diverse naive pool of cells is generated; upon recognition of a pathogen the process of affinity maturation allows B cells to adapt their B-cell receptor (BCR) to epitopes of the pathogen through somatic hypermutation [1]. This process, which takes place in germinal centers [2], can increase the affinity of naive BCR for the target antigen by up to a thousand fold factor [3]. Through affinity maturation, the immune system generates high-affinity, long-lived plasma cells, providing the organism with humoral immunity to pathogens through the secretion of antibodies—the soluble version of the matured BCR—as well as a pool of memory cells with varying affinity to the antigens [4]. However, the diversity and coverage of the memory pool, as well as the biological constraints that control its generation, have not yet been fully explored.

Analysis of high-throughput BCR sequencing data has revealed long tails in the distribution of clonotype abundances, identifying some very abundant clonotypes as well as many very rare ones [5, 6]. Additionally, many receptors have similar sequences and cluster into phylogenetically related lineages [7–11]. These lineages have been used to locally trace the evolution of antibodies in HIV patients [12, 13] and in influenza vaccinees [14, 15]. Memory B-cell clones are more diverse and less specific to the infecting antigen than antibody-producing plasma cells [16, 17]. This suggests that the immune system is trying to anticipate infections by related pathogens or future escape mutants [18].

Theoretical approaches have attempted to qualitatively describe affinity maturation as a Darwinian co-evolutionary process, and studied optimal affinity maturation schemes [19–21, 21], as well as optimal immunization schedules to stimulate antibodies with large neutralizing capabilities [22–24]. Most of these approaches have been limited to short timescales, often with the goal of understanding the evolution of broadly neutralizing antibodies. Here we propose a mathematical framework to explore the trade-offs that control how the large diversity of memory cells evolves over a lifetime.

Despite long-lasting efforts to describe the co-evolution of pathogens and hosts immune systems [25–29], and recent theoretical work on optimal schemes for using and storing memory in the presence of evolving pathogens [30], few theoretical works have described how the B-cell memory repertoire is modified by successive immunization challenges. Early observations in humans [31] have shown that sequential exposure to antigenically drifted influenza strains was more likely to induce an immune response strongly directed towards the first strain the patients were exposed to [32]. This immune imprinting with viral strains encountered early in life was initially called "original antigenic sin," as it can limit the efficiency of vaccination [33]. This phenomenon has been observed in a variety of animal models and viral strains [34]. Secondary infections with an antigenically diverged influenza strain can reactivate or "backboost" memory cells specific to the primary infecting strain [35]. This response is characterized by lower binding affinity but can still have in-vivo efficiency thanks of cross-reactive antibodies [36]. There is a long-standing debate about how detrimental "original antigenic sin" is [37, 38]. However, the question of under what circumstances an immune response based on memory re-use is favourable has not been addressed.

*Corresponding authors. These authors contributed equally.

We build a theoretical framework of joint virus and repertoire evolution in antigenic space, and investigate how acute infections by evolving pathogens have shaped, over evolutionary timescales, the B-cell repertoire response and re-organization. Pathogens causing acute infections may be encountered multiple times over time scales of years, especially when they show a seasonal periodicity, while the maturation processes in the B-cell repertoire take place over a few weeks. This observation allows us to consider that affinity maturation happens in a negligible time with respect to the reinfection period. Within this approximation, we investigate the optimal immune maturation strategies using a framework of discrete-time decision process. We show the emergence of three regimes—monoclonal memory response, polyclonal memory response, and naive response—as trade-offs between immune coverage and resource constraint. Additionally, we demonstrate that reactivation of already existing memory clonotypes can lead to self-trapping of the immune repertoire to low reactivity clones, opening the way for "original antigenic sin."

II. RESULTS

A. Affinity maturation strategies for recurring infections

B cells recognize pathogens through the binding of their BCR to parts of the pathogen's proteins, called epitopes, which we refer to as "antigens" for simplicity. To model this complex protein-protein interaction problem, we assume that both receptors and antigens may be projected into an effective, d-dimensional antigenic space (Fig. 1), following the "generalized shape space" idea pioneered by Perelson and Oster [39]. Receptor-antigen pairs at close distance in that space bind well, while those that are far away bind poorly. Specifically, we define a cross-reactivity function $0 \leq f \leq 1$ quantifying the binding affinity between antigen a and receptor x, which we model by a stretched exponential, $f(x,a) = e^{-(\|x-a\|/r_0)^q}$. This choice of function is the simplest that allows for introducing a cross-reactivity radius, r_0, while controlling how sharply recognition is abrogated as the distance between antigen and receptor oversteps that radius, through the stretching exponent q.

For simplicity, we focus on a single pathogen represented by its immunodominant antigen, so that each viral strain is represented by a single point a_t in antigenic space (black square), where $t = 1, 2, \ldots$ is a discrete time counting the number of re-infections. It is difficult to estimate the rate of re-infections or exposures to the same pathogen. It can be fairly high in humans, where individuals are exposed to the most common viruses from less than once to several times a year [40]. The numbers of lifetime exposures would then range from a few to a few hundreds.

The B cell repertoire, on the other hand, is represented

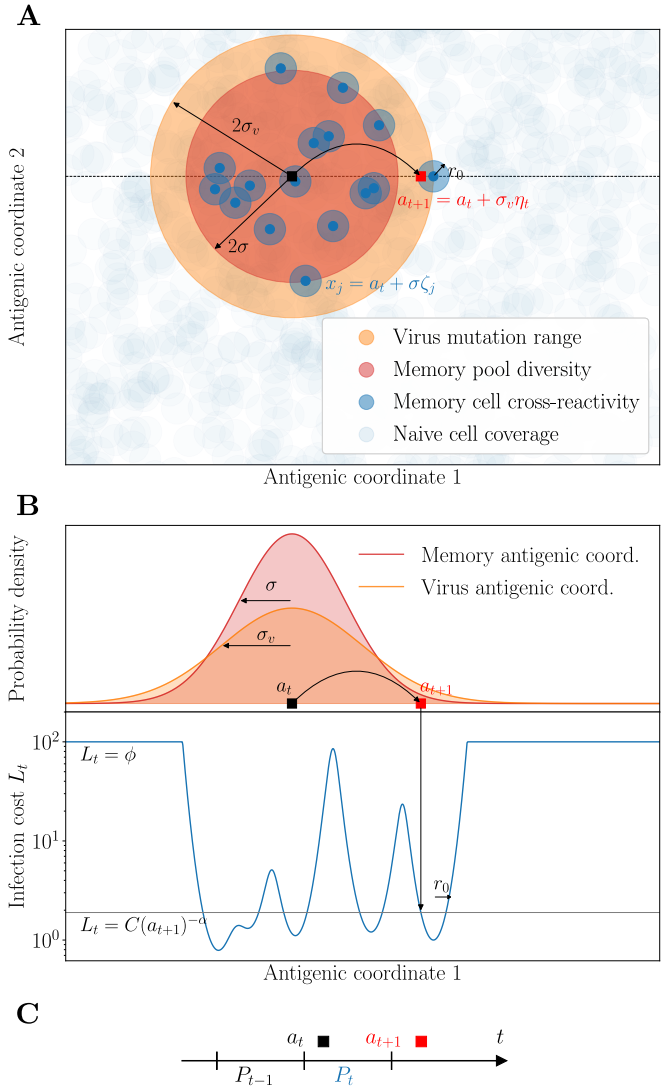

FIG. 1: **Model of sequential affinity maturation. A.** An infecting strain is defined by its position a_n in antigenic space (dark square). In response, the immune system creates m new memory clonotypes x_j (blue points) from a Gaussian distribution of width σ centered in a_t (red area). These new clonotypes create a cost landscape (blue areas) for the next infection, complemented by a uniform background of naive cells (light blue). The next infecting strain (red square) is drawn from a Gaussian distribution of width σ_v centered in a_t (orange area). The position of this strain on the infection landscape is shown with the arrow. Antigenic space is shown in 2 dimensions for illustration purposes, but can have more dimensions in the model. **B.** Cross-section of the distributions of memories and of the next strain, along with the infection cost landscape L_t (in blue). Memories create valleys in the landscape, on a background of baseline naive protection ϕ. **C.** Sequential immunization. Strain a_t modifies the memory repertoire into P_t, which is used to fight the next infection a_{t+1}. P_t is made of all newly created clonotypes (blue points in A) as well as some previously existing ones (not shown). Clonotype abundances are boosted following each infection as a function of the cross-reactivity, and each individual cell survives from one challenge to the other with a probability γ.

by a collection of antigenic coordinates corresponding to each receptor clonotype. We distinguish memory cells (dark blue circles in Fig. 1A), denoted by P_t, which have emerged in response to the presence of the virus, and a dense background of naive cells N (light blue circles) which together provide a uniform but weakly protective coverage of any viral strain (subsumed into the parameter ϕ defined later).

The viral strain evolves randomly in antigenic space, sequentially challenging the existing immune repertoire. This assumption is justified by the fact that for acute infections with a drifting viral strain, such as influenza, the immune pressure exerted on the strain does not happen in hosts but rather at the population level [25]. Viral evolution is not neutral, but it is unpredictable from the point of view of individual immune systems. Specifically, we assume that, upon reinfection, the virus is represented by a new strain, which has moved from the previous antigenic position a_t to the new one a_{t+1} according to a Gaussian distribution with typical jump size σ_v, called "divergence" (see Methods).

Upon infection by a viral strain at a_t, available cross-reactive memory or naive cells will produce antibodies whose affinities determine the severity of the disease. We quantify the efficiency of this early response to the strain a_t with an infection cost I_t:

$$I_t = \min\Big[\phi, \Big(\sum_{x\in P_{t-1}} n_{x,t} f(x,a_t)\Big)^{-\alpha}\Big], \qquad (1)$$

where $\phi > 0$ is a maximal cost corresponding to using naive cells, and where $n_{x,t}$ denotes the size of clonotype x at time t. This infection cost is a decreasing function of the coverage of the virus by the pre-existing memory repertoire, P_{t-1}, $C(a_t) = \sum_{x\in P_{t-1}} n_{x,t} f(x,a_t)$, with a power α governing how sharp that dependence is. Intuitively, the lower the coverage, the longer it will take for memory cells to mount an efficient immune response and clear the virus, incurring larger harm on the organism [41, 42].

When memory coverage is too low, the naive B-cell repertoire as well as the rest of the immune system (including its innate branch as well as T-cell cytotoxic activity) still offers some protection, incurring a maximal cost fixed to ϕ. Memory cells respond more rapidly than naive cells, which is indirectly encoded in our model by the naive cost ϕ being larger than the cost when specific memory cells are present (of order 1 or less). In the Supplementary Text we show how this naive cut-off may be derived in a model where the immune system activates its memory and naive compartments in response to a new infection, when naive clonotypes are very numerous but offer weak protection. In that interpretation, ϕ scales like the inverse density of naive cells. We will refer to ϕ^{-1} as "naive density," although one should keep in mind that this basal protection levels also includes other arms of the immune system. In Fig. 1B we plot an example of the infection cost along a cross-section of the antigenic space.

After this early response, activated memory cells proliferate and undergo affinity maturation to create *de novo* plasma and memory cells targeting the infecting strain. To model this immune repertoire re-organization in response to a new infection a_t, we postulate that its strategy has been adapted over evolutionary timescales to maximize the speed of immune response to subsequent challenges, given the resource constraints imposed by affinity maturation [41]. This strategy dictates the stochastic rules according to which the BCR repertoire evolves from P_{t-1} to P_t as a result of affinity maturation (Fig. 1C).

We consider the following rules inspired by known mechanisms available to the immune system [2]. After the infection by a_t has been tackled by existing receptors, and the infection cost has been paid, new receptors are matured to target future versions of the virus. Their number m_t is distributed according to a Poisson law, whose mean is controlled by the cost of infection, $\bar{m}(I_t)$. This dependence accounts for the feedback of the early immune response on the outcome of affinity maturation, consistent with extensive experimental evidence of the history dependence of the immune response [43]. Each new receptor is roughly located around a_t in antigenic space with some added noise $\sigma(I_t)$, and starts with clonotype size $n_{x,t} = 1$ by convention. The diversification parameter σ can be tuned by the immune system through the permissiveness of selection in germinal centers, through specific regulation factors induced at the early stage of affinity-based selection [44]: $\sigma = 0$ means that affinity maturation only keeps the best binders to the antigens, while $\sigma > 0$ means that selection is weaker.

At the same time, each clonotype $x \in P_{t-1}$ from the previous repertoire may be reactivated and be subsequently duplicated through cell divisions [18], with probability $\mu f(x,a_t)$ (see Methods), proportional to the cross-reactivity, where $0 \le \mu \le 1$ is a proliferation parameter. These previously existing cells and their offspring may then die before the next infection. We denote by γ their survival probability, so that the average lifetime of each cell is $(1-\gamma)^{-1}$. The proliferation and death parameters μ and γ are assumed to be constrained and fixed. The net mean growth is thus given by $\langle n_{x,t}\rangle = (1+\mu f(x,a_t))\gamma n_{x,t-1}$. $\Gamma \equiv (1+\mu)\gamma$ is defined as the maximum growth factor. At the end of the process, the updated repertoire P_t combines the result of this proliferation and death process applied to P_{t-1} with the new receptors obtained from affinity maturation.

To assess the performance of a given strategy, we define an overall cost function at each time step:

$$L_t = I_t + \kappa m_t. \qquad (2)$$

The second term κm_t corresponds to a plasticity cost encoding the resources necessary to generate and maintain new memory clonotypes with affinity maturation. This cost enforces a minimal homeostatic constraint on the memory repertoire. We neglect any dependence of the cost on the diversification σ, which is secondary and

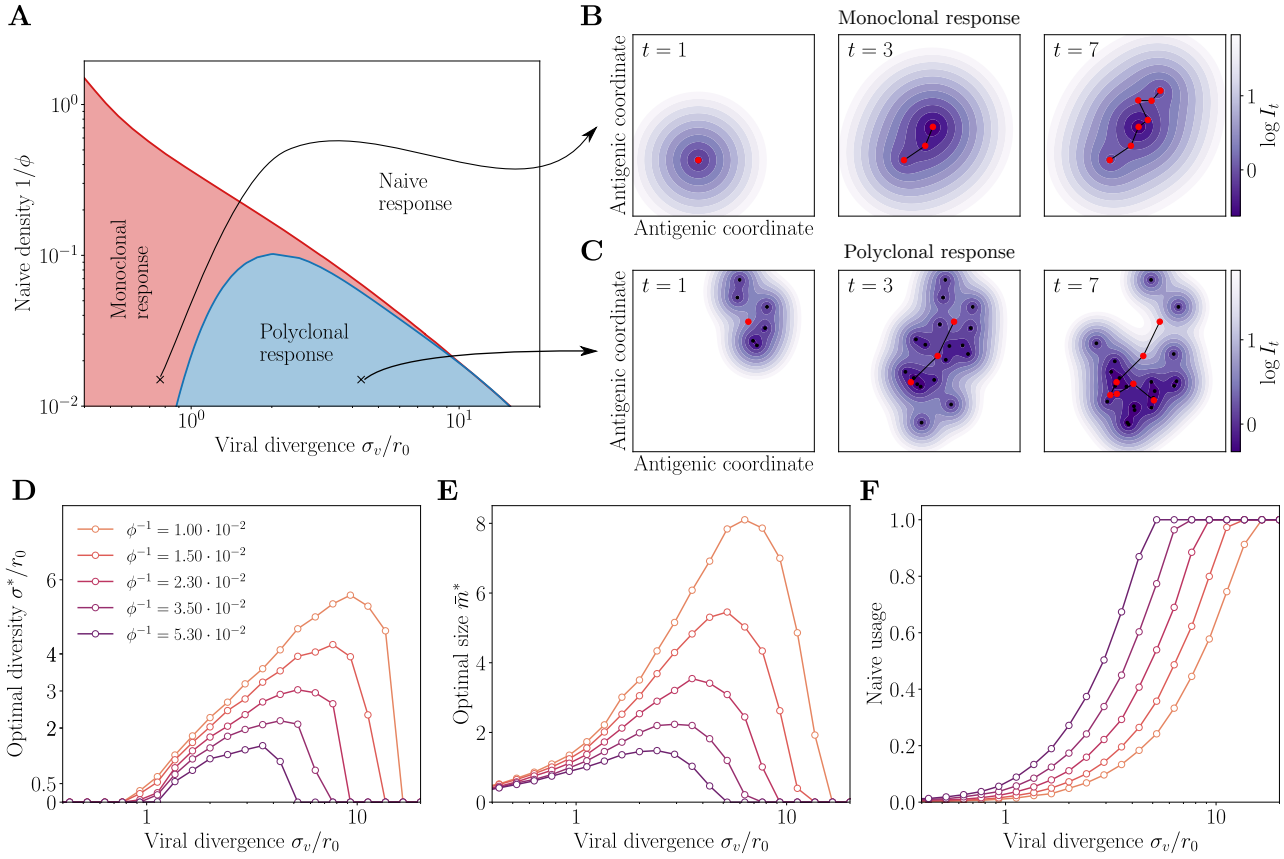

FIG. 2: **Regimes of affinity maturation.** **A.** Phase diagram of the model as a function of the naive coverage $1/\phi$, and viral divergence σ_v, in a two-dimensional antigenic map. Three phases emerge: monoclonal memory (red), polyclonal memory (purple) and naive response (white). **B-C.** Snapshots in antigenic space of the sequential immunization by a viral strain in the (B) monoclonal and (C) polyclonal phases. We show the viral position (red dots), memory clonotypes (black dots), and viral trajectory (black line). The colormap shows the log infection cost. Parameters σ_v and ϕ correspond to the crosses on the phase diagram in A, with their respective optimal σ^*, $\bar{m}^*$ (see arrows). **D.** Diversity σ^*, **E.** optimal size $\bar{m}^*$, and **F.** frequency of naive cell usage in response to an immunization challenge for different naive coverages $1/\phi$. Parameters values: $\kappa = 3.3$, $\alpha = 1$, $q = 2$, $d = 2$, $\gamma = 0.85$, $\mu = 0.5$.

would require adding additional parameters without affecting the qualitative picture. We assume that, over evolutionary timescales, the immune system has minimized the average cumulative cost $\langle L_t \rangle$ over a large number of infections (see Methods). This optimization yields the optimal parameters of the strategy, namely the best functions $m^*(I)$ and $\sigma^*(I)$ describing the extent and diversity of affinity maturation and how they should depend on the strength of the infection I. For the sake of simplicity, in the next three sections we will specialize to the case of constant functions $m(I) \equiv m$ and $\sigma(I) \equiv \sigma$. We will come back to the general case in the last section of the results.

B. Phase diagram of optimal affinity maturation strategies

We obtain optimal constant strategies $\bar{m}(I) = m^*, \sigma(I) = \sigma^*$, by minimizing the simulated long-term cost $\mathcal{L}(\mathcal{S})$ (Eq. 6) in a 2-dimensional antigenic space (see Methods for details of the simulation, optimization procedures, and phase determination). By varying two key parameters, the cost ϕ associated to the use of the naive repertoire, and the virus divergence σ_v, we see a phase diagram emerge with three distinct phases: the naive, monoclonal response, and polyclonal response phases (Fig. 2A). In Figs. 2B-C we show examples of the stochastic evolution of memory repertoires with optimal rules in the two phases (monoclonal and polyclonal responses). Figs. 2D-F show the behaviour of the optimal parameters, as well as the fraction of infections for which the naive repertoire is used (when the maximal infection cost ϕ is paid). The general shape and behaviour of this

phase diagram depends only weakly on the parameter choices (see SI Appendix, Fig. S1).

When the naive repertoire is sufficiently protective (small ϕ), or when the virus mutates too much between infections (large σ_v), the optimal strategy is to produce no memory cells at all ($\bar{m}^* = 0$), and rely entirely on the naive repertoire, always paying a fixed cost $L_t = \mathcal{L} = \phi$ (naive phase).

When the virus divergence σ_v is small relative to the cross-reactivity range r_0, it is beneficial to create memory clonotypes ($\bar{m}^* > 0$), but with no diversity, $\sigma^* = 0$ (monoclonal response). In this case, all newly created clonotypes are invested into a single antigenic position a_t that perfectly recognizes the virus. This strategy is optimal because subsequent infections, typically caused by similar viral strains of the virus, are well recognized by these memory clonotypes.

For larger but still moderate virus divergences σ_v, this perfectly adapted memory is not sufficient to protect from mutated strains: the optimal strategy is rather to generate a polyclonal memory response, with $\bar{m}^* > 0$, $\sigma^* > 0$. In this strategy, the immune system hedges its bet against future infections by creating a large diversity of clonotypes that cover the vicinity of the encountered strain. The created memories are thus less efficient against the *current* infection, which they never will have to deal with. The advantage of this strategy is to anticipate future antigenic mutations of the virus. This diversified pool of cells with moderate affinity is in agreement with recent experimental observations [2, 18, 45, 46]. The diversity of the memory pool is supported by a large number of clonotypes $\bar{m}^*$ (Fig. 2D). As the virus divergence σ_v is increased, the optimal strategy is to scatter memory cells further away from the encountered strain (increasing σ^*, Fig. 2E). However, when σ_v it is too large, both drop to zero as the naive repertoire takes over (Fig. 2F). Increasing the naive density ϕ^{-1} also favors the naive phase. When there is no proliferation on average, i.e. $\Gamma = (1+\mu)\gamma < 1$, there even exists a threshold ϕ_c^{-1} above which the naive strategy is always best (SI Appendix Fig. S1, and Text for estimates of that threshold).

We derived analytical results and scaling laws in a simplified version of the model, where cross-reactivity is ideally sharp, and where memory cells survive only until the next infection. The results are derived in Supplementary Text and illustrated in SI Appendix, Fig. S2. In this simplified setting, the monoclonal to polycolonal transition occurs at $\sigma_v \sim r_0$, consistent with the intuition that diversification occurs when the virus is expected to escaped out of the cross-reactivity radius. The polyclonal-to-naive transition occurs at $(\sigma_v/r_0)^d \sim \phi/\kappa$, when basal coverage by the naive repertoire (and rest of the immune system) outcompetes the protection afforded by memory cells relative to their cost κ.

In summary, the model predicts the two expected regimes of naive and memory use depending on the parameters that set the costs of infections and memory formation. But in addition, it shows a third phase of polyclonal response, where affinity maturation acts as an anticipation mechanism whose role is to generate a large diversity of cells able to respond to future challenges. The prediction of a less focused and thus weaker memory pool observed experimentally is thus rationalized as a result of a bet-hedging strategy.

C. Population dynamics of optimized immune systems

We now turn to the population dynamics of the memory repertoire. When the virus drifts slowly in antigenic space (small σ_v), the same clonotypes get reactivated multiple times, causing their proliferation, provided that $\Gamma = \gamma(1+\mu) > 0$. This reactivation continues until the virus leaves the cross-reactivity range of the original clonotype, at which point the memory clone decays and eventually goes extinct (Fig. 3A). Typical clonotype size trajectories from the model are shown in Fig. 3B. They show large variations in both their maximal size and lifetime. The distribution of clonotype abundances, obtained from a large number of simulations, is indeed very broad, with a power-law tail (Fig. 3C). The lifetime of clonotypes, defined as the time from emergence to extinction, is distributed according to an exponential distribution (Fig. 3D). The exponents governing the tails of these distributions, β and γ, depend on the model parameters, in particular the divergence σ_v.

We can understand the emergence of these distributions using a simple scaling argument, detailed in Supplementary Text. The peak size of a clonotype depends on the number of successive infections by viral strains remaining within a distance $r^* = r_0 \ln[\gamma\mu/(1-\gamma)]^{1/q}$ from the clonotype, under which it continues expanding. This number has a long exponential tail with characteristic time $t_s \sim (r^*/\sigma_v)^2$. One can show that this translates into a power law tail for the distribution of clonotype sizes:

$$p(n^*) \sim \frac{1}{n^{*1+\beta}}, \quad \text{with } \beta \sim \frac{\sigma_v^2}{r^{*2}\ln\Gamma}, \tag{3}$$

and an exponential tail for the lifetime of clonotypes:

$$p(t_l) \sim e^{-\lambda t_l}, \quad \lambda \sim \frac{\sigma_v^2}{r^{*2}}\left(1+\frac{\ln\Gamma}{\ln(1/\gamma)}\right)^{-1}. \tag{4}$$

This simple scaling argument predicts the exponents β and λ fairly well: Figs. 3E-F confirm the validity of the scaling relations (D5)-(D7) against direct evaluation from simulations, for $d = 2$ and $q = 2$. These scalings still hold for different parameter choices (see SI Appendix, Fig. S3).

These scaling relations are valid up to a geometry-dependent prefactor, which is governed by dimensionality and the shape of the cross-reactivity kernel. In the Supplementary Text, we calculate this prefactor in the special case of an all-or-nothing cross-reactivity function, $q = \infty$. Generally, β increases with d, as shown

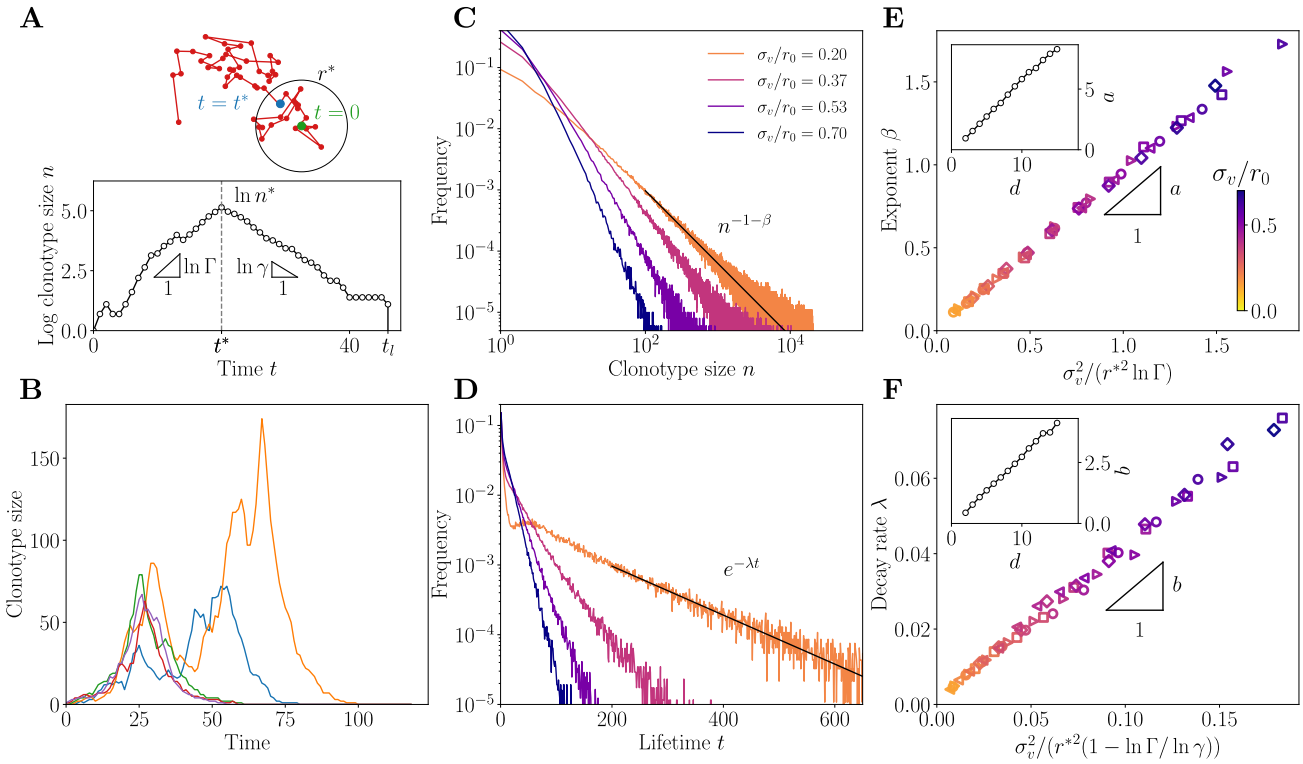

FIG. 3: **Clonotype dynamics and distribution. A.** Sketch of a recall response generated by sequential immunization with a drifting strain. Clonotypes first grow with multiplicative rate $\Gamma = \gamma(1 + \mu)$, until they reach the effective cross-reactivity radius r^*, culminating at n^*, after which they decay with rate γ until extinction at time t_l. **B.** Sample trajectories of clonotypes generated by sequential immunization with a strain of mutability $\sigma_v/r_0 = 0.53$. **C.** Distribution of clonotype size for varying virus mutability σ_v/r_0. **D.** Distribution of the lifetime of a clonotype for varying virus mutability σ_v/r_0. From **B** to **D** the proliferation parameters are set to $\gamma = 0.85$, $\gamma = 0.5$ ie. $\Gamma = 1.275$. **E.** Scaling relation of the power law exponent for varying values of the parameters. Inset: dependence of the proportionality factor a on dimension. **F.** Scaling relation of the decay rate λ for varying σ_v/r_0, with scaling of the proportionality factor b. In both **E-F.**, the different parameters used are $(\gamma = 0.82, \mu = 0.65)$ ie. $\Gamma = 1.353$ (diamonds), $(\gamma = 0.8, \mu = 0.62)$ ie. $\Gamma = 1.296$ (squares), $(\gamma = 0.85, \mu = 0.5)$ ie. $\Gamma = 1.275$ (circles), $(\gamma = 0.87, \mu = 0.4)$ ie. $\Gamma = 1.21$ (triangles >), $(\gamma = 0.9, \mu = 0.35)$ ie. $\Gamma = 1.21$ (triangles <). From **B** to **F**, the strategy was optimized for $\phi = 100$ and $\kappa = 0.5/(1 - \gamma)$. The color code for σ_v/r_0 is consistent across the panels **C** to **F**. In this panel, the other parameters used are $\alpha = 1$, $q = 2$, $d = 2$.

in the insets of Figs. 3E-F for $q = 2$. In higher dimensions, there are more routes to escape the cross-reactivity range, and thus a faster decaying tail of large clonotypes. This effect cannot be explained by having more dimensions in which to mutate, since the antigenic variance is distributed across each dimension, according to σ_v^2/d. Rather, it results from the absence of antigenic back-mutations: in high dimensions, each mutation drifts away from the original strain with a low probability of return, making it easier for the virus to escape, and rarer for memory clonotypes to be recalled upon infections by mutant strains.

D. Comparison to experimental clone-size distributions

The power-law behaviour of the clone-size distribution predicted from the model (Fig. 3E) can be directly compared to existing data on bulk repertoires. While the model makes a prediction for subsets of the repertoire specific to a particular family of pathogens, the same power-law prediction is still valid for the entire repertoire, which is a mixture of such sub-repertoires. Power laws have been widely observed in immune repertoires: from early studies of repertoire sequencing data of BCR in zebrafish [5, 48], to the distribution of clonal family sizes of human IgG BCR [14, 49], as well as in T-cell receptor repertoires [6]. However, these power laws have not yet been reported in the clonotype abundance distribution of human BCR.

To fill this gap, we used publicly available IgG repertoire data of 9 human donors from a recent ultra-deep repertoire profiling study of immunoglobulin heavy-chains (IGH) [47]. The data was downloaded from Sequence Read Archive and processed as in [49]. Repertoires were obtained from the sequencing of IGH mRNA molecules to which unique molecular identifiers (UMI)

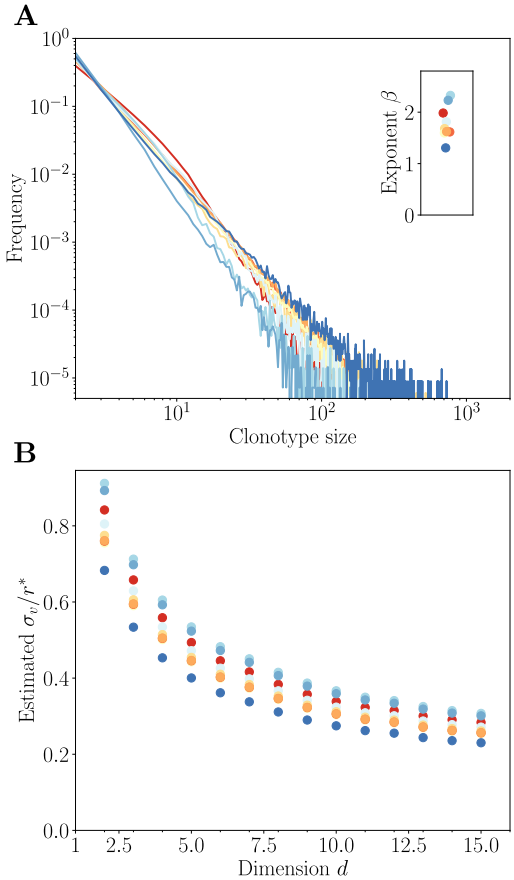

FIG. 4: **Comparison to repertoire data.** **A.** Clonotype abundance distribution of IgG repertoires of healthy donors from [47]. **B.** Estimated mutatibility σ_v in units of the rescaled cross-reactivity r^*, defined as the antigenic distance at which clonotypes stop growing. σ_v is obtained as a function of d by inverting the linear relationship estimated in the inset of Fig. 3E, assuming $q = 2$ and $\Gamma = 1.4$ (estimated from [18]).

were appended. For each IGH nucleotide clonotype found in the dataset, we counted the number of distinct UMI attached to it, as a proxy for the overall abundance of that clonotype in the expressed repertoire. The distributions of these abundances are shown for all 9 donors in Fig. 4A. In agreement with the theory, they display a clear power-law decay $p(n) \sim n^{-1-\beta}$, with $\beta = 1.2$-2.4.

Since the experimental distribution is derived from small subsamples of the blood repertoire, the absolute abundances cannot be directly compared to those of the model. In particular, subsampling means that the experimental distribution focuses on the very largest clonotypes. Thus, comparisons between model and data should be restricted to the tail behavior of the distribution, namely on its power-law exponent β. The bulk repertoire is a mixture of antigen-specific sub-repertoires, each predicted to be a power law with a potentially different exponent. The resulting distribution is still a power-law dominated by the largest exponent.

We used this comparison to predict from the exponent β the virus divergence between infections. To do so, we fit a linear relationship to the inset of Fig. 3E, and invert it for various values of the dimension d to obtain σ_v/r^*. We fixed $\Gamma = 1.4$, which corresponds to a 40% boost of memory B cells upon secondary infection, inferred from a 4-fold boost following 4 sequential immunizations reported in mice [18]. The result is robust to the choice of donor, but decreases substantially with dimension, because higher dimensions mean faster escape, and thus a lower divergence for a given measured exponent (Fig. 4B). The inferred divergence σ_v is always lower than, but of the same order as, the effective cross-reactivity range r^*, suggesting that the operating point of the immune system falls in the transition region between the monoclonal and polyclonal response phases (Fig. 2A).

E. Inhibition of affinity maturation and antigenic imprinting

In this section we come back to general strategies where the process of affinity maturation depends on the immune history through the infection cost I experienced by the system during the early immune response, which controls the number and diversity of newly created memories following that response: $\bar{m}(I)$, $\sigma(I)$. The optimization of the loss function (2) is now carried out with respect to two *functions* of I. To achieve this task, we optimize with respect to discretized functions $(\bar{m}_1, ..., \bar{m}_n)$ and $(\sigma_1, ..., \sigma_n)$ taken at n values of the infection cost I between 0 and ϕ. From this optimization, a clear transition emerges between a regime of complete inhibition of affinity maturation ($\bar{m}^*(I) = 0$) at small infection costs, and a regime of affinity maturation ($\bar{m}^*(I) > 0$) at larger infection costs (Fig. 5A). In the phase where affinity maturation occurs, the optimal diversity $\sigma^*(I)$ is roughly constant (Fig. 5B).

This transition means that when pre-existing protection is good enough, the optimal strategy is not to initiate affinity maturation at all, to save the metabolic cost κm_t. This inhibition of affinity maturation is called "antigenic imprinting," and is linked to the notion of "original antigenic sin," whereby the history of past infections determines the process of memory formation, usually by suppressing it. This phenomenon leads to the paradoxical prediction that a better experienced immune system is less likely to form efficient memory upon new infections. Importantly, in our model this behaviour does not stem from a mechanistic explanation, such as competition for antigen or T-cell help between the early memory response and germinal centers, but rather as a long-term optimal strategy maximizing immune coverage while minimizing the costs of repertoire re-arrangement.

To simplify the investigation of antigenic imprinting, we approximate the optimal strategies in Fig. 5A and B by step functions, with a suppressed phase, $\bar{m}(I) = 0$,

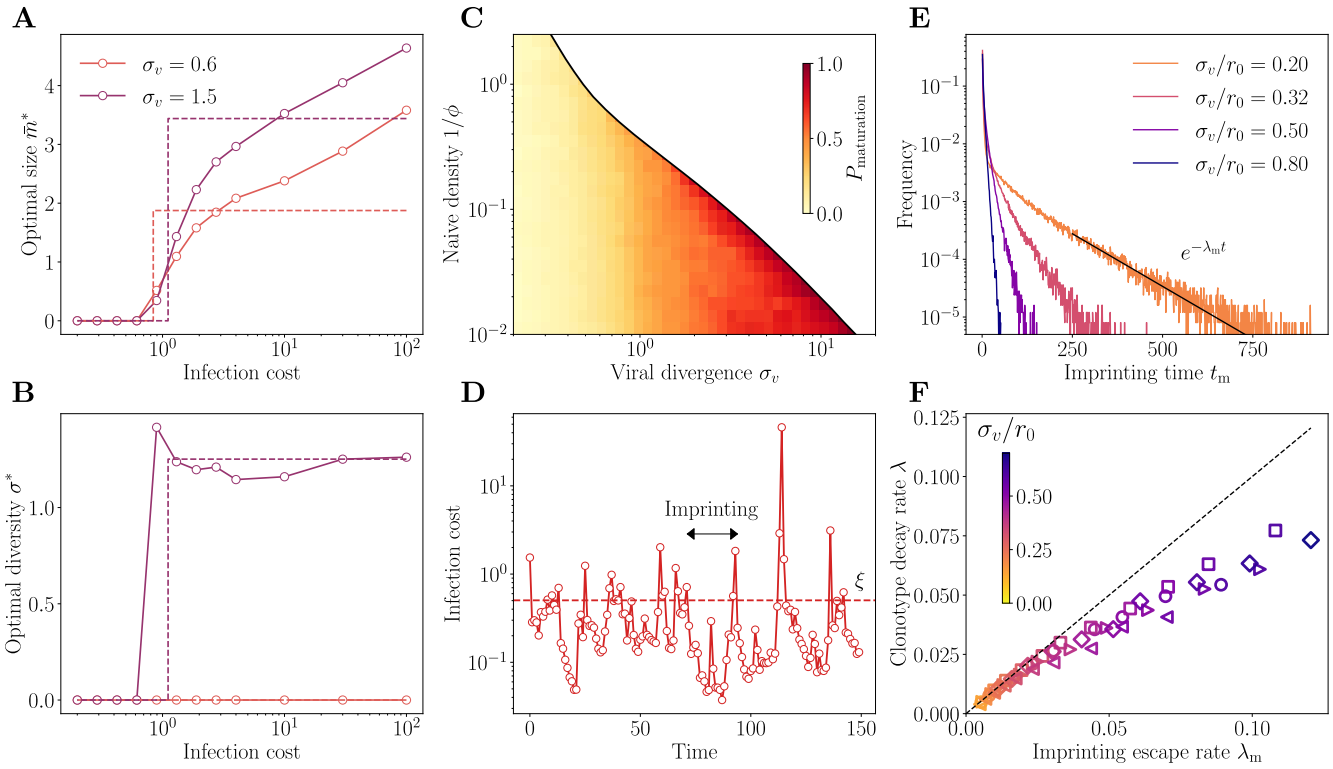

FIG. 5: **Imprinting and backboosting. A,B.** Optimal regulatory functions for (**A**) the number $m^*(I)$ and (**B**) the diversity $\sigma^*(I)$ of new memories as a function of the infection cost I, for two values for the viral divergence. These functions show a sharp transition from no to some memory formation, suggesting to replace them by simpler step functions (dashed lines). This step function approximation is used in the next panels. **C.** Frequency of infections leading to affinity maturation in the optimal strategy. The frequency increases with the virus divergence σ_v, up to the point where the transitions to the naive phase where memory is not used at all. **D.** Typical trajectory of infection cost in sequential infections at $\sigma_v/r_0 = 0.5$. When the cost goes beyond the threshold ξ, affinity maturation is activated, leading to a drop in infection cost. These periods of sub-optimal memory describe an "original antigenic sin," whereby the immune system is frozen in the state imprinted by the last maturation event. **E.** Distribution of imprinting times, i.e. the number of infections between affinity maturation events, decays exponentially with rate λ_m. The proliferation parameters in **A** to **E** are set to $\gamma = 0.85$ and $\mu = 0.5$. **F.** Predicted scaling of λ_m with the clonotype decay rate λ from Figs. 3D and F. In **F**, the different parameters used are ($\gamma = 0.82, \mu = 0.65$) ie. $\Gamma = 1.353$ (diamonds), ($\gamma = 0.8, \mu = 0.62$) ie. $\Gamma = 1.296$ (squares), ($\gamma = 0.85, \mu = 0.5$) ie. $\Gamma = 1.275$ (circles), ($\gamma = 0.87, \mu = 0.4$) ie. $\Gamma = 1.21$ (triangles >), ($\gamma = 0.9, \mu = 0.35$) ie. $\Gamma = 1.21$ (triangles <). The color code for σ_v/r_0 is consistent across the panels **E** and **F**. From **D** to **F**, the strategy was optimized for $\phi = 100$ and $\kappa = 0.5/(1-\gamma)$. In this panel, the other parameters used are $\alpha = 1$, $q = 2$, $d = 2$.

for $I < \xi$, and an active phase, $\bar{m}(I) \equiv \bar{m} > 0$ and $\sigma(I) \equiv \sigma$, for $I > \xi$. The threshold ξ is left as an optimization parameter, in addition to σ and $\bar{m}$. Optimizing with respect to these three parameters, we observe that the frequency of affinity maturation events mostly depends on σ_v (Fig. 5C). While this threshold remains approximately constant, the frequency of affinity maturation events increases as σ_v increases. At small σ_v, the optimal strategy is to extensively backboost existing memory cells; for large σ_v, the growing unpredictability of the next viral move makes it more likely to have recourse to affinity maturation. In other words, when the virus is stable (low σ_v), the immune system is more likely to capitalize on existing clonotypes, and not implement affinity maturation, because savings made on the plasticity cost outweigh the higher infection cost. As the virus drifts away with time, this infection cost also increases, until it reaches the point where affinity maturation becomes worthwile again.

Trajectories of the infection cost show the typical dynamics induced by backboosting, with long episodes where existing memory remains sufficient to keep the cost below ξ (Fig. 5D), interrupted by infections that fall too far away from existing memory, triggering a new episode of affinity maturation and concomitant drop in the infection cost.

We call the time between affinity maturation events $t_{\rm m}$. Its mean $\langle t_{\rm m} \rangle$ is equal to the inverse of the frequency of maturation events, and thus decreases with σ_v. Its distribution, shown Fig. 5E, has an exponential tail with exponent $\lambda_{\rm m}$. The exponential tail of the distribution of $t_{\rm m}$ is dominated by episodes where the viral strain

drifted less than expected. In that case, the originally matured clonotype grows to a large size, offering protection for a long time, even after it has stopped growing and only decays. We therefore expect that in the case of a slowly evolving virus $\sigma_v \ll r_0$, the escape rate from the suppressed phase is given by the clonotype decay rate: $\lambda_{\rm m} \sim \lambda$. We verify this prediction in Fig. 5F. Interestingly, for slowly evolving viruses, the typical clonotype lifetime diverges, leading to a lifelong imprinting by the primary immune challenge. Conversely, as the viral divergence σ_v grows, the imprinting time decays faster than the typical clonotype lifetime and the extent of the imprinting phenomenon is limited.

III. DISCUSSION

Adaptive immunity coordinates multiple components and cell types across entire organisms over many space and time scales. Because it is difficult to empirically characterize the immune system at the system scale, normative theories have been useful to generate hypotheses and fill the gap between observations at the molecular, cellular, and organismal scales [50, 51]. Such approaches include clonal selection theory [52], or early arguments about the optimal size and organization of immune repertoires [39, 53, 54], and of affinity maturation [20, 21]. While these theories do not rely on describing particular mechanisms of immune function, they may still offer quantitative insights and help predict new mechanisms or global rules of operation.

Previous work developed models of repertoire organization as a constrained optimization problem where the expected future harm of infection, or an *ad hoc* utility function, is minimized [30, 41, 42, 55]. In Ref. [42], it was assumed clonotypes specific to all antigens are present at all times in the repertoire; the mechanism of immune memory then merely consists of expanding specific clonotypes at the expense of others. This assumption describes T-cell repertoires well, where there are naive cells with good affinity to essentially any antigen [56]. For B cells the situation is more complex because of affinity maturation. In addition, re-organizing the repertoire through mutation and selection has a cost, and is subject to metabolic and physical constraints.

Our work addresses these challenges by proposing a framework of discrete-time decision process to describe the optimal remodeling of the B-cell repertoire following immunization, through a combination of affinity maturation and backboosting. While similar to [30], our approach retains the minimal amount of mechanistic details and focuses on questions of repertoire remodeling, dynamics, and structure. The specific choices of the cost functions were driven by simplicity, while still retaining the ability to display emergent behaviour. Generalizing the metabolic cost function to include e.g. costs of diversification (through a dependence on σ) or of cell proliferation is not expected to affect our results qualitatively, although it may shift the exact positions of the transition boundaries.

We investigated strategies that maximize long-term protection against subsequent challenges and minimize short-term resource costs due to the affinity maturation processes. Using this model, we observed that optimal strategies may be organized into three main phases as the pathogen divergence and naive coverage are varied. We expect these distinct phases to co-exist in the same immune system, as there exists a wide range of pathogen divergences, depending on their evolutionary speed and typical frequency of recurrence.

For fast recurring or slowly evolving pathogens, the monoclonal response ensures a very specific and targeted memory. This role could be played by long-lived plasma cells. These cells are selected through the last rounds of affinity maturation, meaning that they are very specific to the infecting strain [57]. Yet, despite not being called memory cells, they survive for long times in the bone marrow, providing long-term immunity.

For slow recurring or fast evolving pathogens, the polyclonal response provides a diverse memory to preempt possible antigenic drift of the pathogen. The function could be fulfilled by memory B cells, which are released earlier in the affinity maturation process, implying that they are less specific to the infecting strain, but collectively can cover more immune escape mutations. While affinity maturation may start from both memory or naive B cells during sequential challenges, the relative importance of each is still debated [58–60]. Our model does not commit on this question since we assume that the main benefit of memory is on the infection cost, rather than its re-use in subsequent rounds of affinity maturation.

For simplicity our model assumed random evolution of the virus. However, there is evidence, backed by theoretical arguments, that successive viral strains move in a predominant direction in antigenic space, as a result of immune pressure by the host population [29, 61, 62]. While it is unlikely that the immune system has evolved to learn how to exploit this persistence of antigenic motion in a specific manner, such a bias in the random walk is expected to affect the optimal strategy, as we checked in simulations (Fig. S4). The bias of the motion effectively increases the effective divergence of the virus, favoring the need for more numerous and more diverse memory cells. However, it does not seem to affect the location of the polyclonal-to-naive transition.

The model is focused on acute infections, motivated by the assumption that recurring infections and antigenic drift are the main drivers of affinity maturation evolution. However, much of the model and its results can be reinterpreted for chronic infections. In that context, the sequential challenges of our model would correspond to selective sweeps in the viral population giving rise to new dominant variants. While the separation of time scales between the immune response and the rate of reinfections would no longer hold, we expect some predictions, such as the distribution of clonotype sizes and the emergence

of imprinting, to hold true. Chronic infections also imply that the virus evolves as a function of how the immune system responds. Including this feedback would require a game-theoretic treatment. We speculate that it would drive antigenic motion in a persistent direction, as argued earlier and evaluated in Fig. S4.

We investigated strategies where the outcome of affinity maturation is impacted by the efficiency of the early immune response. It is known that the extrafollicular response can drastically limit antigen availability and T-cell help, decreasing the extent of affinity maturation and the production of new plasma and memory cells [63]. Our general framework allows for but does not presume the existence of such negative feedbacks. Instead, they naturally emerge from our optimization principle. We further predict a sharp transition from no to some affinity maturation as a function of the infection cost. This prediction can be interpreted as the phenomenon of antigenic imprinting widely described in sequential immunization assays [35], or "original antigenic sin" [34]. It implies that having been exposed to previous strains of the virus is detrimental to mounting the best possible immune response. Importantly, while antigenic imprinting has been widely described in the literature, no evolutionary justification was ever provided for its existence. Our model explains it as a long term optimal strategy for the immune system, maximizing immune coverage while minimizing repertoire re-arrangements (encoded in the cost κm).

We believe this framework can be generalized to investigate interactions between slow and fast varying epitopes, which are known to be at the core of the low effectiveness of influenza vaccines [32]. When during sequential challenges only one of multiple epitopes changes at a time, it may be optimal for the immune system to rely on its protections against the invariant epitopes. Only after all epitopes have escaped immunity does affinity maturation get re-activated concomitantly to a spike of infection harm, similar to our result for a single antigen.

Our model can explain previously reported power laws in the distributions of abundances of B-cell receptor clonotypes. However, there exist alternative explanations to such power laws [64, 65] that do not require antigenically drifting antigens. Our model predictions could be further tested in a mouse model, by measuring the B-cell recall response to successive challenges [35], but with epitopes carefully designed to drift in a controlled manner, to check the transition predicted in Fig. 5A. While not directly included in our model, our result also suggest that the size of the inoculum, which would affect the infection cost, should also affect backboosting. This effect could also be tested in mouse experiments. The predicted relationships between viral divergence and the exponents of the power law and clonotype lifetimes (Figs. 3E and F) could be tested in longitudinal human samples, by sequencing sub-repertoires specific to pathogens with different rates of antigenic evolution. This would require to computationally predict what B-cell receptors are specific to what pathogen, which in general is difficult.

We only considered a single pathogenic species at a time, with the assumption that pathogens are antigenically independent, so that the costs relative to different pathogens are additive. Possible generalizations could include secondary infections, as well as antigenically related pathogens showing patterns of cross-immunity (such as cowpox and smallpox, or different groups of influenza), which could help us shed light on complex immune interactions between diseases and serotypes, such as negative interference between different serotypes of the Dengue fever leading to hemorrhagic fever, or of the human Bocavirus affecting different body sites or [34].

IV. MATERIALS AND METHODS

A. Mathematical model

The viral strain is modeled by its antigenic position, which follows a discrete random walk:

$$a_{t+1} = a_t + \sigma_v \eta_{t+1}, \tag{5}$$

where η_t is a normally distributed d-dimensional variable with $\langle \eta_t \rangle = 0$ and $\langle \eta_t \cdot \eta_{t'} \rangle = \delta_{tt'}$.

The positions of newly created memory receptors are drawn at random according to: $x_j = a_t + \sigma \xi_j$, $j = 1, \ldots, m_t$, where ξ_j is normally distributed with $\langle \xi_j \rangle = 0$ and $\langle \xi_j^2 \rangle = 1$. Their initial sizes are set to $n_{x_j,t} = 1$.

Upon further stimulation, the new size $n'_{x,t}$ of a pre-existing clonotype right after proliferation is given by $n'_{x,t} - n_{x,t-1} \sim \text{Binom}(n_{x,t-1}, \mu f(x, a_t))$, where $f(x, a) = e^{-(\|x-a\|/r_0)^q}$ is the cross-reactivity Kernel. After proliferation, each memory may die with probability γ, so that the final clonotype size after an infection cycle is given by $n_{x,t} \sim \text{Binom}(n'_{x,t}, \gamma)$.

The objective to be minimized is formally defined as a long-term average:

$$\mathcal{L}(m, \sigma) = \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} L_t. \tag{6}$$

The optimal strategy is defined as:

$$(\bar{m}^*, \sigma^*) = \arg\min_{(m,\sigma)} \mathcal{L}(m, \sigma), \tag{7}$$

where m^* and σ^* are, in the general case, full functions of the infection cost I, $m^*(I)$ and $\sigma^*(I)$. In all results except in the last section of the Results, we use the Ansatz of constant functions, $m(I) \equiv m$, $\sigma(I) \equiv \sigma$. In the last section of the Results, we first perform optimization over discretized functions $m = (m_1, \ldots, m_n)$, $\sigma = (\sigma_1, \ldots, \sigma_n)$ defined over n chosen values of $I = (I_1, \ldots, I_n)$. Then, we parametrize the functions as step functions: $\sigma(I) = 0$ and $m(I) = 0$ for $I < \xi$, and $\sigma(I) = \sigma$ and $m(I) = m$ for $I > \xi$, and optimize over the 3 parameters σ, m, ξ.

B. Monte-Carlo estimation of the optimal strategies

The average cumulative cost $\mathcal{L}$ in Eq. 6 is approximated by a Monte-Carlo method. To ensure the simulated repertoire reaches stationarity, we start from a naive repertoire and discard an arbitrary number of initial viral encounters. Because the process is ergodic, simulating a viral-immune trajectory over a long time horizon is equivalent to simulating M independent trajectories of smaller length T. To ensure the independence of our random realizations across our M parallel streams we use the random number generators method *split* provided in Tina's RNG library [66]. The cumulative cost function $\mathcal{L}$ is convex for the range of parameters tested. To optimize $\mathcal{L}$ under positivity constraints for the objective variables σ, $\bar{m}$ and ξ, we use Py-BOBYQA [67], a coordinate descent algorithm supporting noisy objective functions.

The polyclonal to monoclonal (red curve) and memory to naive (blue curve) boundaries of the phase diagrams in SI Appendix, Fig. S2 and Fig. S3 are obtained by respectively solving $\partial\mathcal{L}/\partial\sigma = 0$ in the monoclonal phase and $\partial\mathcal{L}/\partial\bar{m} = 0$ in the naive phase. Both these derivatives can be approximated by finite differences with arbitrary tolerances on σ and $\bar{m}$. We fix the tolerance on σ to 0.2 and the tolerance on $\bar{m}$ to 0.01. To obtain the root of these difference functions, we use a bisection algorithm. In order to further decrease the noise level, we compute the difference functions across pairs of simulations, each pair using an independent sequence of pathogens a_t of length $L = 400$. The number of independent pairs of simulations used for each value of σ_v and ϕ is $M \sim 10^5$.

Acknowledgements

The study was supported by the European Research Council COG 724208 and ANR-19-CE45-0018 "RESP-REP" from the Agence Nationale de la Recherche and DFG grant CRC 1310 "Predictability in Evolution". The authors thank Natanael Spisak for pre-processing the raw data from [47], and for useful discussions and suggestions.

[1] Nieuwenhuis P, Opstelten D (1984) Functional anatomy of germinal centers. *American Journal of Anatomy* 170:421–435.

[2] Victora GD, Nussenzweig MC (2012) Germinal Centers. *Annual Review of Immunology* 30:429–457.

[3] Eisen HN, Siskind GW (1964) Variations in Affinities of Antibodies during the Immune Response *. *Biochemistry* 3:996–1008.

[4] Weisel F, Shlomchik M (2017) Memory B Cells of Mice and Humans. *Annual Review of Immunology* 35:255–284.

[5] Weinstein JA, Jiang N, White RA, Fisher DS, Quake SR (2009) High-Throughput Sequencing of the Zebrafish Antibody Repertoire. *Science* 324:807–810.

[6] Mora T, Walczak AM (2019) in *Systems Immunology* (CRC Press).

[7] Kocks C, Rajewsky K (1988) Stepwise intraclonal maturation of antibody affinity through somatic hypermutation. *Proceedings of the National Academy of Sciences* 85:8206–8210.

[8] Kleinstein SH, Louzoun Y, Shlomchik MJ (2003) Estimating Hypermutation Rates from Clonal Tree Data. *The Journal of Immunology* 171:4639–4649.

[9] Kepler TB (2013) Reconstructing a B-cell clonal lineage. I. Statistical inference of unobserved ancestors. *F1000Research* 2:103.

[10] Yaari G, Kleinstein SH (2015) Practical guidelines for B-cell receptor repertoire sequencing analysis. *Genome Medicine* 7:121.

[11] Ralph DK, Iv FAM (2016) Likelihood-Based Inference of B Cell Clonal Families. *PLOS Computational Biology* 12:e1005086.

[12] Liao HX, et al. (2013) Co-evolution of a broadly neutralizing HIV-1 antibody and founder virus. *Nature* 496:469–476.

[13] Nourmohammad A, Otwinowski J, Łuksza M, Mora T, Walczak AM (2019) Fierce Selection and Interference in B-Cell Repertoire Response to Chronic HIV-1. *Molecular Biology and Evolution* 36:2184–2194.

[14] Jiang N, et al. (2013) Lineage Structure of the Human Antibody Repertoire in Response to Influenza Vaccination. *Science Translational Medicine* 5:171ra19–171ra19.

[15] Horns F, Vollmers C, Dekker CL, Quake SR (2019) Signatures of selection in the human antibody repertoire: Selective sweeps, competing subclones, and neutral drift. *Proceedings of the National Academy of Sciences* 116:1261–1266.

[16] Smith KGC, Light A, Nossal GJV, Tarlinton DM (1997) The extent of affinity maturation differs between the memory and antibody-forming cell compartments in the primary immune response. *The EMBO Journal* 16:2996–3006.

[17] Weisel FJ, Zuccarino-Catania GV, Chikina M, Shlomchik MJ (2016) A Temporal Switch in the Germinal Center Determines Differential Output of Memory B and Plasma Cells. *Immunity* 44:116–130.

[18] Viant C, et al. (2020) Antibody Affinity Shapes the Choice between Memory and Germinal Center B Cell Fates. *Cell* 183:1298–1311.e11.

[19] Oprea M, Perelson AS (1997) Somatic mutation leads to efficient affinity maturation when centrocytes recycle back to centroblasts. *The Journal of Immunology* 158:5155–5162.

[20] Oprea M, Van Nimwegen E, Perelson AS (2000) Dynamics of one-pass germinal center models: Implications for affinity maturation. *Bulletin of Mathematical Biology* 62:121–153.

[21] Kepler TB, Perelson AS (1993) Somatic Hypermutation in B Cells: An Optimal Control Treatment. *Journal of Theoretical Biology* 164:37–64.

[22] Wang S, et al. (2015) Manipulating the Selection Forces

during Affinity Maturation to Generate Cross-Reactive HIV Antibodies. *Cell* 160:785–797.

[23] Sachdeva V, Husain K, Sheng J, Wang S, Murugan A (2020) Tuning environmental timescales to evolve and maintain generalists. *Proceedings of the National Academy of Sciences* 117:12693–12699.

[24] Molari M, Eyer K, Baudry J, Cocco S, Monasson R (2020) Quantitative modeling of the effect of antigen dosage on B-cell affinity distributions in maturating germinal centers. *eLife* 9:e55678.

[25] Grenfell BT (2004) Unifying the Epidemiological and Evolutionary Dynamics of Pathogens. *Science* 303:327–332.

[26] Blanquart F, Gandon S (2013) Time-shift experiments and patterns of adaptation across time and space. *Ecology Letters* 16:31–38.

[27] Cobey S, Wilson P, Matsen FA (2015) The evolution within us. *Philosophical Transactions of the Royal Society B: Biological Sciences* 370:20140235.

[28] Koelle K, Cobey S, Grenfell B, Pascual M (2006) Epochal Evolution Shapes the Phylodynamics of Interpandemic Influenza A (H3N2) in Humans. *Science* 314:1898–1903.

[29] Marchi J, Lässig M, Walczak AM, Mora T (2021) Antigenic waves of virus–immune coevolution. *Proceedings of the National Academy of Sciences* 118:e2103398118.

[30] Schnaack OH, Nourmohammad A (2021) Optimal evolutionary decision-making to store immune memory. *eLife* 10:e61346.

[31] Francis T (1960) On the Doctrine of Original Antigenic Sin. *Proceedings of the American Philosophical Society* 104:572–578.

[32] Cobey S, Hensley SE (2017) Immune history and influenza virus susceptibility. *Current Opinion in Virology* 22:105–111.

[33] Hoskins TW, Davies J, Smith AJ, Miller C, Allchin A (1979) Assessment of inactivated influenza-A vaccine after three outbreaks of influenza A at Christ's Hospital. *The Lancet* 313:33–35.

[34] Vatti A, et al. (2017) Original antigenic sin: A comprehensive review. *Journal of Autoimmunity* 83:12–21.

[35] Kim JH, Skountzou I, Compans R, Jacob J (2009) Original Antigenic Sin Responses to Influenza Viruses. *The Journal of Immunology* 183:3294–3301.

[36] Linderman SL, Hensley SE (2016) Antibodies with 'Original Antigenic Sin' Properties Are Valuable Components of Secondary Immune Responses to Influenza Viruses. *PLOS Pathogens* 12:e1005806.

[37] Yewdell JW, Santos JJS (2020) Original Antigenic Sin: How Original? how Sinful? *Cold Spring Harbor Perspectives in Medicine* p a038786.

[38] Worobey M, Plotkin S, Hensley SE (2020) Influenza Vaccines Delivered in Early Childhood Could Turn Antigenic Sin into Antigenic Blessings. *Cold Spring Harbor Perspectives in Medicine* p a038471.

[39] Perelson AS, Oster GF (1979) Theoretical studies of clonal selection: Minimal antibody repertoire size and reliability of self-non-self discrimination. *Journal of Theoretical Biology* 81:645–670.

[40] Cohen C, et al. (2021) Asymptomatic transmission and high community burden of seasonal influenza in an urban and a rural community in South Africa, 2017–18 (PHIRST): A population cohort study. *The Lancet Global Health* 9:e863–e874.

[41] Mayer A, Balasubramanian V, Mora T, Walczak AM (2015) How a well-adapted immune system is organized. *Proceedings of the National Academy of Sciences* 112:5950–5955.

[42] Mayer A, Balasubramanian V, Walczak AM, Mora T (2019) How a well-adapting immune system remembers. *Proceedings of the National Academy of Sciences* 116:8815–8823.

[43] Oidtman RJ, et al. (2021) Influenza immune escape under heterogeneous host immune histories. *Trends in Microbiology* 29:1072–1082.

[44] Nakagawa R, et al. (2021) Permissive selection followed by affinity-based proliferation of GC light zone B cells dictates cell fate and ensures clonal breadth. *Proceedings of the National Academy of Sciences* 118.

[45] Tas JMJ, et al. (2016) Visualizing antibody affinity maturation in germinal centers. *Science* 351:1048–1054.

[46] Kuraoka M, et al. (2016) Complex Antigens Drive Permissive Clonal Selection in Germinal Centers. *Immunity* 44:542–552.

[47] Briney B, Inderbitzin A, Joyce C, Burton DR (2019) Commonality despite exceptional diversity in the baseline human antibody repertoire. *Nature* 566:393–397.

[48] Mora T, Walczak AM, Bialek W, Callan CG (2010) Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences* 107:5405–5410.

[49] Spisak N, Walczak AM, Mora T (2020) Learning the heterogeneous hypermutation landscape of immunoglobulins from high-throughput repertoire data. *Nucleic Acids Research* 48:10702–10712.

[50] Chakraborty AK (2017) A Perspective on the Role of Computational Models in Immunology. *Annual Review of Immunology* 35:403–439.

[51] Altan-Bonnet G, Mora T, Walczak AM (2020) Quantitative Immunology for Physicists. *Physics Reports* 849:1–83.

[52] Burnet FM (1957) A modification of Jerne's theory of antibody production using the concept of clonal selection. *The Australian Journal of Science* 20:67–69.

[53] Perelson AS, Mirmirani M, Oster GF (1976) Optimal strategies in immunology. I. B-cell differentiation and proliferation. *Journal of Mathematical Biology* 3:325–367.

[54] Perelson AS, Mirmirani M, Oster GF (1978) Optimal strategies in immunology. II. B memory cell production. *Journal of Mathematical Biology* 5:213–256.

[55] Marsland R, Howell O, Mayer A, Mehta P (2021) Tregs self-organize into a computing ecosystem and implement a sophisticated optimization algorithm for mediating immune response. *Proceedings of the National Academy of Sciences* 118:e2011709118.

[56] Moon JJ, et al. (2007) Naive CD4+ T Cell Frequency Varies for Different Epitopes and Predicts Repertoire Diversity and Response Magnitude. *Immunity* 27:203–213.

[57] Akkaya M, Kwak K, Pierce SK (2020) B cell memory: Building two walls of protection against pathogens. *Nature Reviews Immunology* 20:229–238.

[58] Turner JS, et al. (2020) Human germinal centres engage memory and naive B cells after influenza vaccination. *Nature* 586:127–132.

[59] Mesin L, et al. (2020) Restricted Clonality and Limited Germinal Center Reentry Characterize Memory B Cell Reactivation by Boosting. *Cell* 180:92–106.e11.

[60] Wong R, et al. (2020) Affinity-Restricted Memory B Cells Dominate Recall Responses to Heterologous Flaviviruses.

Immunity 53:1078–1094.e7.
[61] Bedford T, Rambaut A, Pascual M (2012) Canalization of the evolutionary trajectory of the human influenza virus. *BMC Biology* 10:38.
[62] Bedford T, et al. (2014) Integrating influenza antigenic dynamics with molecular evolution. *eLife* 3:e01914.
[63] Arulraj T, Binder SC, Robert PA, Meyer-Hermann M (2021) Germinal Centre Shutdown. *Frontiers in Immunology* 12:2730.
[64] Desponds J, Mora T, Walczak AM (2016) Fluctuating fitness shapes the clone-size distribution of immune repertoires. *Proceedings of the National Academy of Sciences* 113:274–279.
[65] Gaimann MU, Nguyen M, Desponds J, Mayer A (2020) Early life imprints the hierarchy of T cell clone sizes. *eLife* 9:e61639.
[66] Bauke H, Mertens S (2007) Random numbers for large-scale distributed Monte Carlo simulations. *Physical Review E* 75:066701.
[67] Cartis C, Fiala J, Marteau B, Roberts L (2019) Improving the Flexibility and Robustness of Model-based Derivative-free Optimization Solvers. *ACM Transactions on Mathematical Software* 45:32:1–32:41.
[68] Redner S (2001) *A Guide to First-Passage Processes* (Cambridge University Press, Cambridge).

Supplementary information

Appendix A: Mean-field naive coverage

Here we show how the infection cost function defined in the main text,

$$I_t = \min\left[\phi, \left(\sum_{x\in P_{t-1}} c_{x,t} f(x,a_t)\right)^{-\alpha}\right], \tag{A1}$$

may be derived as the mean-field limit of a repertoire with memory and naive compartments.

In addition to the evolving memory repertoire P_t already described in the main text, we define a naive repertoire made of random receptors $\mathcal{N}$, distributed uniformly with density ρ. Viruses may be recognized by either the memory or naive clonotypes. The naive coverage is defined as:

$$C_{\text{naive}}(a_t) = \sum_{x\in\mathcal{N}} f(x,a_t), \tag{A2}$$

and the memory coverage as before:

$$C(a_t) = \sum_{x\in P_{t-1}} c_{x,t} f(x,a_t). \tag{A3}$$

(In this convention, each naive clonotype has size one in arbitrary units.)

Depending on the values of these coverages, the system will choose to use either the naive repertoire, or an existing memory. In this decision, we factor in the fact that using the naive repertoire is more costly, which we account for using a prefactor $\beta < 1$. The cost is then defined as:

$$L_t = \max\left[\beta C_{\text{naive}}(a_t), C(a_t)\right]^{-\alpha}. \tag{A4}$$

We can simplify this expression in the limit where naive clonotypes are very numerous, but each offer weak coverage. In the limit of high density of naive cells, $\rho \to \infty$, the coverage self-averages to its mean value:

$$C_{\text{naive}} \approx \langle C_{\text{naive}}\rangle = \rho \int d^d x\, f(x,a_t) = \rho U_d(q) r_0^d, \tag{A5}$$

with

$$U_d(q) = \int d^d y\, e^{-\|y\|^q} = S_d \int r^{d-1} dr\, e^{-r^q}, \tag{A6}$$

where we have done the change of variable $x = a_t + y r_0$, and where $S_d = 2\pi^{d/2}/\Gamma(d/2)$ is the surface area of the unit sphere.

Taking the $\rho \to \infty$ and $\beta \to 0$ limits, while keeping $\beta\rho$ finite, corresponds to a dense naive repertoire but where each naive cell weakly covers the antigenic space. In this limit we recover the model of the main text

$$I_t = \min\left[\phi, C(a_t)^{-\alpha}\right], \tag{A7}$$

with $\phi = (\rho\beta U_d(q) r_0^d)^{-\alpha}$.

Appendix B: Transition from monoclonal to naive phase at $\sigma_v = 0$

Here we derive an expression for the phase boundary between the monoclonal and polyclonal phases in the limit $\sigma_v = 0$, where the virus does not move.

In the special case where $\mu = 0$, clonotypes cannot multiply. At each time step, a number m_n of new clonotypes at created at $a_n = \text{const}$, distributed according to a Poisson law of mean $\bar{m}$. This number is added to existing clonotypes, of which a random fraction γ survives. If the previous number of clonotypes, M_n, is Poisson distributed with mean $\bar{M}_n$, the number of surviving ones M'_n is also Poisson distributed with mean $\gamma\bar{M}_n$ (since subsampling a Poisson-distributed number still gives a Poisson law). Then, the new number of clonotypes, $M_{n+1} = M'_n + m_n$, is also Poisson distributed, with the recurrence relation:

$$\bar{M}_{n+1} = \gamma\bar{M}_n + \bar{m}. \tag{B1}$$

At steady state, we have

$$\bar{M}_n \to \frac{\bar{m}}{1-\gamma}. \tag{B2}$$

Since all clonotypes are at $x = a_n$, the coverage is $C(a_n) = M$, so that the expected cost reads:

$$\mathcal{L} = \phi \exp\left(-\frac{\bar{m}}{1-\gamma}\right) + \sum_{M=1}^{+\infty} \frac{1}{M^\alpha} \exp\left(-\frac{\bar{m}}{1-\gamma}\right) \frac{1}{m!} \left(\frac{\bar{m}}{1-\gamma}\right)^M + \kappa\bar{m}. \tag{B3}$$

To find the transition from monoclonal to naive, $\bar{m} = 0$, we need to find the value of ϕ for which $\partial\mathcal{L}/\partial\bar{m}$ changes sign at $\bar{m} = 0$: if this derivative is positive, it is better to have $\bar{m} = 0$ (since the function is convex); if it is negative, there is benefit to be gained by increasing $\bar{m} > 0$. The condition:

$$\left.\frac{\partial\mathcal{L}}{\partial\bar{m}}\right|_{\bar{m}=0} = \kappa - \frac{\phi}{1-\gamma} + \frac{1}{1-\gamma} \tag{B4}$$

gives the transition point

$$\phi_c = 1 + \kappa(1-\gamma) \tag{B5}$$

For $\mu > 0$, we redefine M_n as the sum of all clonotype sizes, which is equal to the coverage, $C = M_n = \sum_{x\in P_n} c_{x,n}$. The recurrence relation is replaced by:

$$\bar{M}_{n+1} = \gamma(1+\mu)\bar{M}_n + \bar{m}. \tag{B6}$$

For $\gamma(1+\mu) > 1$, this number explodes, so that M is infinite, reducing the infection cost to 0 regardless of $\bar{m}$. The transition point is then

$$\phi_c = 0. \tag{B7}$$

For $\gamma(1+\mu) < 1$, $\bar{M}_n$ reaches a steady state value,

$$\bar{M}_n \to \frac{\bar{m}}{1-\gamma(1+\mu)}. \tag{B8}$$

Although M_n is not strictly distributed according to a Poisson law, it is still a good approximation, so that we can repeat the same argument as with $\mu = 0$,

$$\phi_c \approx 1 + \kappa(1-\gamma(1+\mu)). \tag{B9}$$

Appendix C: Analytical results in a solvable model

1. Model definition and main results

To gain insight into the transitions observed in the phase diagram of Fig. 2, we can make a series of simplifications and approximations about the model that allow for analytical progress. We assume a step function for the cross-reactivity function $f(x,a) = 1$ for $\|x-a\| \leq r_0$, and 0 otherwise, corresponding the limit $q = \infty$. Likewise, we assume

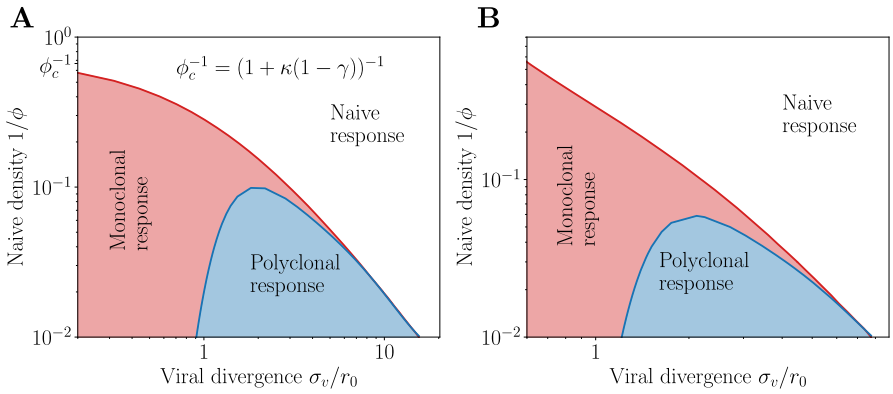

FIG. S1: **Phase diagram for various parameters. A.** Phase diagram for parameters $\mu = 0$, $\gamma = 0.85$ and $d = 2$. We observe the upper transition point $\phi_c = (1+\kappa(1-\gamma))^{-1}$ at $\sigma_v = 0$. **B.** Phase diagram for parameters $\mu = 0.5$, $\gamma = 0.85$ and $d = 3$. Since $\gamma(1+\mu) > 1$ the transition point $\phi_c = \infty$. In both **A** and **C** we observe that the phase diagram retains the same shape. In this panel $\alpha = 1$, $q = 2$ and $\kappa = 0.5/(1-\gamma)$.

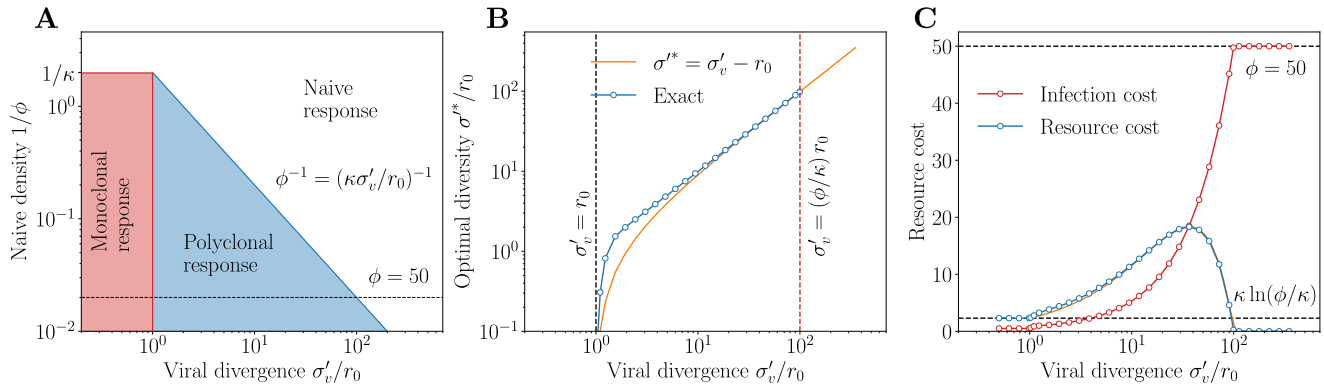

FIG. S2: **Analytical solution of a tractable model. A.** Exact phase diagram in $d = 1$ for the simplified model ($q = \infty$, $\gamma = 0$, and all-or-nothing infection cost). The boundary between monoclonal is given by $\sigma_v' = r_0$ and the boundary between polyclonal by $\phi^{-1} = (\kappa\sigma_v'/r_0)^{-d}$. **B.** Optimal memory diversity $\sigma'^* \approx \sigma_v'$ and **C.** optimal infection and plasticity costs for as a function of σ_v' for $\phi = 50$. σ' and σ_v' are rescaled versions of the diversity and divergence to match the variances of the original model.

a uniform distribution of viral antigenic mutations $a_{t+1} = a_t + \sigma_v' \eta_{t+1}'$, where η_t' is a random point of the d-dimensional unit ball, with $\sigma_v' = \sigma_v\sqrt{1+2/d}$ (so that the variance is the same as in the Gaussian case), and similarly for memory diversification, with new clonotypes drawn from a uniform distribution is a ball of radius $\sigma' = \sigma\sqrt{1+2/d}$. The infection cost is approximated by an all-or-nothing function, with $I_t = 0$ if there is any coverage $C(a_t) > 0$, and $I_t = \phi$ if $C(a_t) = 0$. We further assume $\gamma = 0$: all clonotypes are discarded at each time step, so that memory may only be used once.

In this simplified version of the model, the phase diagram and optimal parameters can be computed analytically. We first summarize the main results below. In the subsequent paragraphs, we provide detailed derivations in the case of arbitrary dimensions, and also provide additional exact formulas in the one-dimensional case.

One can show (see next paragraphs) that the transition from monoclonal to polyclonal response occurs exactly when the radius of the ball within which viral mutations occur reaches the cross-reactivity radius r_0:

$$\sigma_v' = r_0. \tag{C1}$$

Below this transition ($\sigma_v' < r_0$), the optimal strategy is to have no diversity at all and perfectly target the recognized antigen a_t, $\sigma'^* = 0$, as any memory cell at a_t will recognize the next infection. In this case the optimal mean number of memories $\bar{m}^* = \ln(\phi/\kappa)$ results in a trade off between the cost of new memories with the risk of not developing any memory at all by minimizing $\phi e^{-\bar{m}} + \kappa\bar{m}$. The transition from the monoclonal response to naive phases is then given by $\phi = \kappa$, where $\bar{m}^* = 0$.

The polyclonal-to-naive transition may also be understood analytically. In the polyclonal response phase, the optimal strategy is, in either of the limits $\sigma'_v \gg r_0$ or $\bar{m}^* \ll 1$ (see next paragraphs):

$$\sigma'^* \approx \sigma'_v - r_0 \tag{C2}$$

$$\bar{m}^* \approx \frac{\sigma'^d_v}{r_0^d} \ln\left(\frac{\phi}{\kappa}\frac{r_0^d}{\sigma'^d_v}\right). \tag{C3}$$

In particular this result becomes exact at the transition from polyclonal to naive, where $\bar{m}^* = 0$. The transition is thus given by:

$$\phi^{-1} = \frac{r_0^d}{\kappa \sigma'^d_v}. \tag{C4}$$

The polyclonal response is outcompeted by the naive one when the density of naive cell (ϕ^{-1}) becomes larger than the probability density of new strains falling within the cross-reactivity radius ($r_0^d/{\sigma'_v}^d$), rescaled by the memory cost coefficient κ^{-1}.

Fig. S2 shows the resulting phase diagram, as well as the optimal diversity σ'^* and predicted costs for a fixed ϕ and $d = 1$. These predictions reproduce the main features of the full model, in particular the scaling of the immune diversity σ with σ_v (Fig. 2D vs. Fig. S2B) and the general shape of the optimal memory size $\bar{m}^*$ (Fig. 2E vs. Fig. S2C), which first increases as the virus becomes more divergent, to later drop to zero as memory becomes too costly to maintain and the system falls into the naive phase.

2. General formulation

Define $P_{\text{hit}}(\sigma', r_0, r)$ as the probability that a random memory will recognize the next infection at distance r, i.e. the probability that a random point in the ball of radius σ' and a point at distance r from its center are at distance $\leq r_0$. The probability that none of m clonotypes recognize the virus, where m is drawn from a Poisson distribution of mean $\bar{m}$, reads:

$$P_{\text{miss}}(\sigma', r_0, r) = \sum_m e^{-\bar{m}} \frac{\bar{m}^m}{m!} (1 - P_{\text{hit}}(\sigma', r_0, r))^m = e^{-\bar{m} P_{\text{hit}}(\sigma', r_0, r)}. \tag{C5}$$

The best strategy maximizes this probability, averaged over the location of the next infection, over σ':

$$\bar{P}_{\text{miss}}(\sigma', \bar{m}, r_0, \sigma'_v) = \left\langle e^{-\bar{m} P_{\text{hit}}(\sigma', r_0, r)} \right\rangle_{\mathcal{B}(\sigma'_v)} = \frac{1}{\sigma'^d_v V_d} \int_0^{\sigma'_v} S_d r^{d-1} dr \, e^{-\bar{m} P_{\text{hit}}(\sigma', r_0, r)} \tag{C6}$$

where $\mathcal{B}_{\sigma'_v}$ is the ball of radius σ'_v, $V_d = \pi^{d/2}/\Gamma(d/2+1)$ is the volume of a unit ball, and $S_d = 2\pi^{d/2}/\Gamma(d/2)$ the area of the unit sphere.

Then the expected overall cost reads:

$$\mathcal{L} = \phi \bar{P}_{\text{miss}}(\sigma', \bar{m}, r_0, \sigma_v) + \kappa \bar{m} = \phi \left\langle e^{-\bar{m} P_{\text{hit}}(\sigma', r_0, r)} \right\rangle_{\mathcal{B}(\sigma'_v)} + \kappa \bar{m}. \tag{C7}$$

3. Exact location of the phase transitions, and approximate solution in the polyclonal phase

The location of the optimal σ' may be rigorously bounded from above and below. If $\sigma' < \sigma'_v - r_0$, then only part of the future positions of the virus are covered, so increasing σ' can bring no harm. Likewise, for $\sigma' > \sigma'_v + r_0$, memory covers parts of the antigenic space that have no chance of harboring the next virus, so that decreasing σ' is also always advantageous. Thus, the optimum σ'^* must satisfy:

$$\sigma'_v - r_0 \leq \sigma'^* \leq \sigma'_v + r_0. \tag{C8}$$

As already argued in the main text, when $\sigma'_v < r_0$, there is clearly no benefit to having $\sigma' > 0$, so the optimum is reached at $\sigma' = 0$. (C8) further shows that if $\sigma'_v > r_0$, then $\sigma'^* > 0$, so that a polyclonal phase is optimal. As a consequence, the transition from the monoclonal to polyclonal phases happens exactly at

$$\text{Monoclonal to polyclonal:} \quad \sigma'_v = r_0. \tag{C9}$$

In the monoclonal phase, memory always recognizes the next virus. The only risk of paying ϕ is when no memory is created, which happens with probability $e^{-\bar{m}}$, so that the cost reads:

$$\mathcal{L} = \phi e^{-\bar{m}} + \kappa \bar{m}. \tag{C10}$$

The optimal $\bar{m}^* = \ln(\phi/\kappa)$ cancels at the monoclonal-to-naive transition:

$$\text{Monoclonal to naive:} \quad \phi = \kappa. \tag{C11}$$

In the polyclonal phase, we could not find a general analytical solution, but there are two limits in which the solution may be calculated. The first limit is when $\sigma_v \gg r_0$. In that case, (C8) implies $\sigma'^* \approx \sigma'_v$, and

$$P_{\text{hit}}(\sigma'^*, r_0, r) \approx \frac{r_0^d}{\sigma_v'^d}, \tag{C12}$$

which doesn't depend on r. Then, minimizing

$$\mathcal{L} \approx \phi e^{-\bar{m} P_{\text{hit}}} + \kappa \bar{m} \tag{C13}$$

with respect to $\bar{m}$ yields:

$$\bar{m}^* = \frac{1}{P_{\text{hit}}} \ln\left(\frac{\phi}{\kappa} P_{\text{hit}}\right) = \frac{\sigma_v'^d}{r_0^d} \ln\left(\frac{\phi}{\kappa} \frac{r_0^d}{\sigma_v'^d}\right). \tag{C14}$$

The second limit in which things simplify is for small $\bar{m}$. Then the exponential in (C6) may be expanded at first order, yielding:

$$\mathcal{L} = \phi \left(1 - \bar{m} \langle P_{\text{hit}}(\sigma', r_0, r)\rangle_{\mathcal{B}(\sigma'_v)}\right) + \kappa \bar{m}, \tag{C15}$$

where $\langle \cdot \rangle_r$ is the mean of over the ball of radius σ'_v. Minimizing with respect to σ' is equivalent to maximizing $\langle P_{\text{hit}}(\sigma', r_0, r)\rangle_r$, which is the probability that a random point in the ball of radius σ' and a random point in the ball of radius σ'_v are separated by less than r_0. This probability is maximized for any $\sigma' \leq \sigma_v - r_0$, where it is equal to $(r_0/\sigma'_v)^d$. Increasing σ' beyond $\sigma_v - r_0$ can only lower the probability of recognition. Thus:

$$\min_{\sigma'} \mathcal{L} = \phi + \bar{m}\left(\kappa - \phi \frac{r_0^d}{\sigma_v'^d}\right). \tag{C16}$$

This gives us the condition for the transition from polyclonal to naive, where $\bar{m}^* = 0$. This happens when

$$\text{Polyclonal to naive:} \quad \phi = \kappa \frac{\sigma_v'^d}{r_0^d}. \tag{C17}$$

This condition gives us an exact expression for the location of the transition.

4. Exact solution in dimension 1

For $d = 1$, the cost $\mathcal{L}$ in (C7) may be calculated analytically, by using exact expressions of $P_{\text{hit}}(\sigma', r_0, r)$. When $\sigma'_v \leq r_0$, the optimal σ' is zero as explained in the main text. When $\sigma'_v > r_0$, we distinguish two cases: $r_0 < \sigma'_v \leq 2r_0$, and $\sigma'_v > 2r_0$.

a. Case $r_0 < \sigma'_v \leq 2r_0$. Since we know that the optimal σ' is between $\sigma'_v - r_0$ and $\sigma'_v + r_0$, we focus on that range. Then there are two subcases for σ'.

If $\sigma'_v - r_0 < \sigma' \leq r_0$, there are two contributions to the integral of $\bar{P}_{\text{miss}}$ over the position of the virus r. Either $r \leq r_0 - \sigma'$, then all memories recognize the virus, $P_{\text{hit}} = 1$; or $r_0 - \sigma' < r < \sigma'_v < r_0 + \sigma'$, in which case the recognition probability is given by the normalized intersection of two balls at distance r of radii σ' and r_0,

$$P_{\text{hit}} = \frac{\sigma' + r_0 - r}{2\sigma'}. \tag{C18}$$

Thus we obtain doing the integral over r in (C6):

$$\bar{P}_{\text{miss}}(\sigma', \bar{m}, r_0, \sigma'_v) = \frac{1}{\sigma'_v}\left[(r_0 - \sigma')e^{-\bar{m}} + \int_{r_0-\sigma'}^{\sigma'_v} \exp\left(-\bar{m}\frac{\sigma' + r_0 - r}{2\sigma'}\right) dr\right] \text{ if } \sigma'_v - r_0 < \sigma' \leq r_0. \tag{C19}$$

If $r_0 < \sigma' \leq \sigma'_v + r_0$, there are also two contributions. Either $r \leq \sigma' - r_0$, in which case there is no boundary effect, and the recognition probability is just $P_{\rm hit} = r_0/\sigma'_v$; or $\sigma' - r_0 < r \leq \sigma'_v \leq \sigma' + r_0$, in which case we have again (C18). Performing the integration in (C6) we obtain:

$$\bar{P}_{\rm miss}(\sigma', \bar{m}, r_0, \sigma'_v) = \frac{1}{\sigma'_v}\left[\exp\left(-\bar{m}\frac{r_0}{\sigma'}\right)(\sigma' - r_0) + \int_{\sigma'-r_0}^{\sigma'_v} \exp\left(-\bar{m}\frac{\sigma' + r_0 - r}{2\sigma'}\right) dr\right] \text{ if } r_0 < \sigma' \leq \sigma'_v + r_0. \quad \text{(C20)}$$

Numerical analysis shows that (C19) admits a minimum as a function of σ' in its interval of validity, $\sigma'_v - r_0 < \sigma'^* \leq r_0$, while (C20) is always increasing.

b. Case $\sigma'_v > 2r_0$. In this case, there is only a single subcase in the range of interest $\sigma'_v - r_0$ and $\sigma'_v + r_0$. This case is the same as the previous one considered, and the result is given by the same formula (C20). However, for $\sigma'_v > 2r_0$, this expression now admits a minimum $\sigma'_v - r_0 < \sigma'^* \leq \sigma'_v + r_0$.

We recover that in both cases (a and b), in the limit $\bar{m} \to 0$, this minimum is reached at $\sigma'_v - r_0$.

Appendix D: Population dynamics in sequential immunization

1. Clonotype growth and decay as a first-passage problem

We now want to study clonotype proliferation induced by a recall response. We focus on the limit of small mutation rates $\sigma_v \ll r_0$. Within this regime, the system is in the monoclonal phase with $\sigma^* = 0$. We can therefore focus on the case of a single clonotype at position $x = 0$ on the phenotypic space, and ask how successive challenges will modify its size. (Different initial conditions will only change the prefactor in front of the exponential modes in the distribution of first passage times, so the large time behavior of this probability distribution will be the same as discussed below.)

The clonotype has an initial size $c = 1$, and the virus drifts away from $x = 0$ with viral divergence σ_v. In the general model, cells have probability γ to survive from one challenge to the other. Proliferation is taken to be proportional to the cross reactivity radius, $\mu e^{-(r/r_0)^q}$. The population dynamics is thus given by the approximate recursion:

$$n_{t+1} \approx n_t \gamma \left[1 + \mu e^{-(r/r_0)^q}\right], \quad \text{(D1)}$$

where we have neglected birth-death noise. We can further simplify this equation to $n_{t+1} = \gamma(1 + \mu\Theta(r - r^*))n_t$, where $r^* = r_0 \ln(\gamma\mu/(1-\gamma))^{1/q}$ is defined as the radius at which the net fold-change factor crosses 1, i.e. when birth is exactly compensated by death. This means that, as long as the virus is within distance r^*, the clonotype grows with fold-change factor $\sim \Gamma$. As soon as it reaches r^*, and neglecting possible returns below r^* (which happen with probability 1 for $d \leq 2$, but with a frequency that does not affect the overall decay), it will decay with fold-change factor $\sim \gamma$. The problem is thus reduced to determining the first-passage time of the viral antigenic location at radius r^*.

We use a continuous approximation corresponding to a slowly evolving strain, $\sigma_v \ll r_0$:

$$a(0) = 0, \qquad da = \frac{\sigma_v}{\sqrt{d}} dW, \quad \text{(D2)}$$

where W is a Wiener process. The radius, given by $r(t) = |a(t)|$, behaves on average as:

$$\langle r(t)^2 \rangle = t\sigma_v^2. \quad \text{(D3)}$$

The time it takes for $r(t)$ to reach r^*, denoted by t^*, is approximately given by $\langle t^* \rangle \sim (r^*/\sigma_v)^2$.

The tail of the distribution for this first-passage time is dominated by rare events when the virus mutates less than expected between infections, leading to larger episodes of growth. We will show in the next two sections that the distribution of these exceptionally long t^* has an exponential tail:

$$P(t^* > t) \sim e^{-t/t_s}, \quad t_s \sim \langle t^* \rangle \sim \frac{r^{*2}}{\sigma_v^2}. \quad \text{(D4)}$$

This translates into a power-law tail for the peak clonotype abundance,

$$p(n^*) \sim \frac{1}{n^{*1+\beta}}, \quad \text{with } \beta \sim \frac{\sigma_v^2}{r^{*2}\ln\Gamma}. \quad \text{(D5)}$$

The same scaling holds for the distribution of all abundances, since the peak determines the rest of the trajectory.

Within the same simplified picture, the lifetime t_l of a clonotype is the sum of the time it takes to reach the peak, t^*, and the decay time until extinction, which is approximately $\ln(n^*)/\ln(1/\gamma)$:

$$t_l = t^* + \frac{\ln(n^*)}{\ln(1/\gamma)} = \left(1 + \frac{\ln\Gamma}{\ln(1/\gamma)}\right) t^*. \tag{D6}$$

Thus, t_l is proportional to t^*, and therefore also exponentially distributed:

$$p(t_l) \sim e^{-\lambda t_l}, \quad \lambda \sim \frac{\sigma_v^2}{r^{*2}}\left(1 + \frac{\ln\Gamma}{\ln(1/\gamma)}\right)^{-1}. \tag{D7}$$

Next we derive in detail the distributions of the first passage time of $r(t)$ to r^* to obtain Eq. D4.

2. First passage time in $d = 1$

The distribution of first passage time, $p(t)$, can be computed solving diffusion with a box of size $2r^*$ in $d = 1$ [68]:

$$p(t) = \sum_{n=0}^{+\infty} \frac{(2n+1)\sigma_v^2\pi}{2r^{*2}}(-1)^n \exp\left(-\frac{(2n+1)^2\pi^2\sigma_v^2 t}{8r^{*2}}\right) \tag{D8}$$

The dominant term ($n = 0$) at long times gives an exponential decay:

$$p(t) \approx \frac{\sigma_v^2\pi}{2r^{*2}} \exp\left(-\frac{\pi^2\sigma_v^2 t}{8r^{*2}}\right). \tag{D9}$$

3. First passage in higher dimensions

We define $f(r,t)$ as the probability density that the virus has not yet reached r^* at time t, and is at radius r.

This probability density is solution to the diffusion equation with spherical symmetry and absorbing boundary conditions in arbitrary dimension $d > 1$:

$$\frac{\partial f}{\partial t} = \frac{\sigma_v^2}{2d}\left[\frac{\partial^2 f}{\partial r^2} + \frac{d-1}{r}\frac{\partial f}{\partial r}\right], \qquad f(r^*, t) = 0. \tag{D10}$$

Assuming separation of variables, $f(r,t) = T(t)R(r)$, we have:

$$\frac{T'(t)}{T(t)} = \frac{\sigma_v^2}{2d}\frac{R''(r) + \frac{d-1}{r}R'(r)}{R(r)} \equiv -\frac{\sigma_v^2}{2dr^{*2}}\lambda, \tag{D11}$$

where λ is to be determined later. This implies $T(t) = Ce^{-\lambda\frac{\sigma_v^2}{2dr^{*2}}t}$ where C is a constant. The radial part $R(r)$ is solution to:

$$R''(r) + \frac{d-1}{r}R'(r) = -\frac{\lambda}{r^{*2}}R(r). \tag{D12}$$

For $d = 1$ this equation reduces to a harmonic equation and we recover the above solution in 1D. Using the change of variable $R(r) = r^{1-d/2}g(r)$ we derive the following equation:

$$r^2 g''(r) + rg'(r) + \left(\lambda\frac{r^2}{r^{*2}} - \left(\frac{d}{2} - 1\right)^2\right)g(r) = 0. \tag{D13}$$

Changing the variable $x = \sqrt{\lambda}r/r^*$, the function $\tilde{g}(x) = g(xr^*/\sqrt{\lambda})$ is solution to the Bessel differential equation of order $d/2-1$. It can therefore be written as a superposition of a Bessel function of the first kind and a Bessel function of the second kind, both of order $d/2 - 1$. The Bessel function of the second kind having a singularity at $x = 0$, our solution is only given by the Bessel function of the first kind $g(xr^*/\sqrt{\lambda}) = BJ_{d/2-1}(x)$. The radial function R now reads:

$$R(r) = Br^{1-d/2}J_{d/2-1}(\sqrt{\lambda}r/r^*). \tag{D14}$$

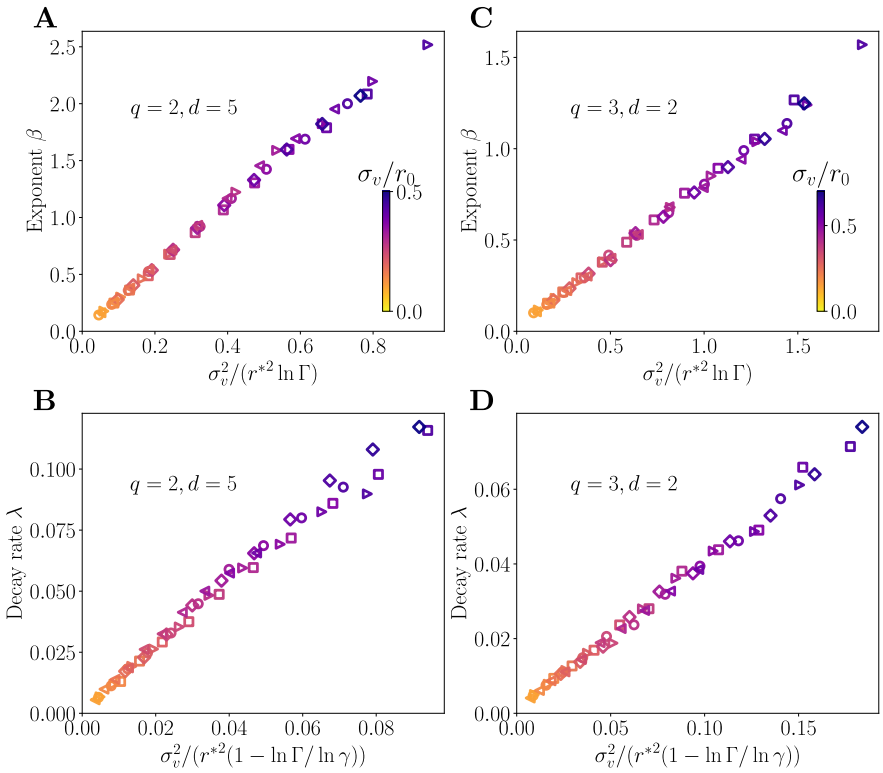

FIG. S3: **Scaling relations for various parameters. A-B.** Power law exponent β and lifetime decay rate λ in dimension $d = 5$ with a Gaussian cross-reactivity kernel with $q = 2$. **C-D.** Power law exponent β and lifetime decay rate λ in dimension $d = 2$ with a cross-reactivity kernel with $q = 3$. The different parameters used are $(\gamma = 0.82, \mu = 0.65)$ i.e. $\Gamma = 1.353$ (diamonds), $(\gamma = 0.8, \mu = 0.62)$ i.e. $\Gamma = 1.296$ (squares), $(\gamma = 0.85, \mu = 0.5)$ i.e. $\Gamma = 1.275$ (circles), $(\gamma = 0.87, \mu = 0.4)$ i.e. $\Gamma = 1.21$ (triangles $>$), $(\gamma = 0.9, \mu = 0.35)$ ie. $\Gamma = 1.21$ (triangles $<$). The strategy is optimized with $\phi = 100$, $\kappa = 0.5/(1-\gamma)$. We used $\alpha = 1$ throughout.

The absorbing boundary condition at $r = r^*$ gives us the condition $J_{d/2-1}(\sqrt{\lambda}) = 0$, which has an infinite number of solutions $j_{0,d/2-1}, ..., j_{n,d/2-1}, ...$, so that λ can take values

$$\lambda_n = j^2_{n,d/2-1}. \tag{D15}$$

The general solution to (D10) is given as a linear combination of all possible modes, with coefficients C_n determined from boundary conditions and the Dirac delta initial condition, $f(r, 0) = \delta(r)$:

$$f(r,t) = \sum_{n=0}^{+\infty} C_n r^{1-d/2} J_{d/2-1}\left(\frac{j_{n,d/2-1}}{r^*} r\right) \exp\left(-\frac{j^2_{n,d/2-1}}{2dr^{*2}}\sigma_v^2 t\right). \tag{D16}$$

The distribution of first passage times asymptotically follows the largest mode of this series, $n = 0$, so that:

$$p(t) \sim \exp\left(-\frac{j^2_{0,d/2-1}}{2dr^{*2}}\sigma_v^2 t\right). \tag{D17}$$

For instance for $d = 2, 3, 4$ we have $j_{0,0} \approx 2.40483$, $j_{0,1/2} = \pi$, $j_{0,1} \approx 3.83171$.

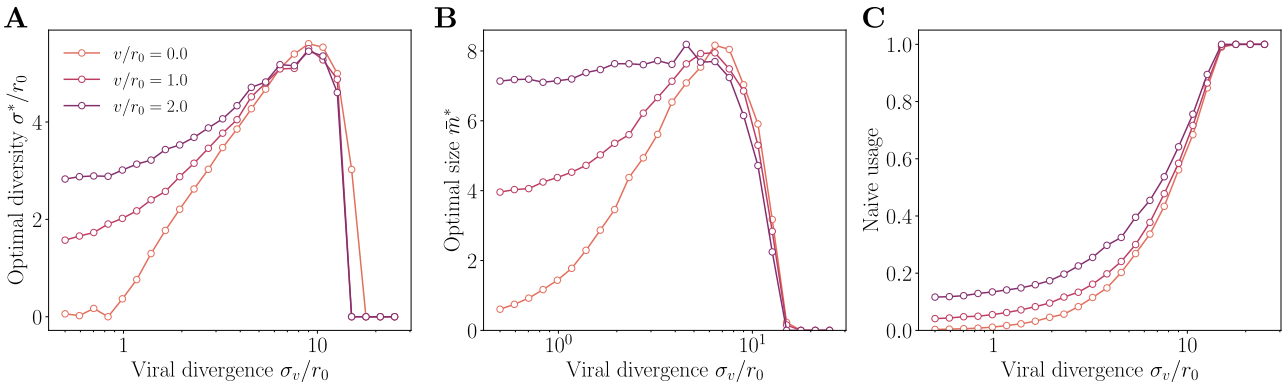

FIG. S4: **Optimal strategy in presence of drift. A.** Optimal diversity **B.** size and **C.** frequency of naive usage in response to an immunization challenge for different strain drift v/r_0. The strain follows a random walk with drift v in a fixed direction e_0: $a_{t+1} = a_t + ve_0 + \sigma_v \eta_t$. Parameters values: $\mu = 0.5$, $\gamma = 0.85$, $\kappa = 3.3$, $\alpha = 1$, $q = 2$, $d = 2$.