Optimal evolutionary decision-making to store immune memory

Oskar H Schnaack[1] and Armita Nourmohammad[1,2,3*]
1. *Max Planck Institute for Dynamics and Self-organization, Am Faßberg 17, 37077 Göttingen, Germany*
2. *Department of Physics, University of Washington, 3910 15th Ave Northeast, Seattle, WA 98195*
3. *Fred Hutchinson Cancer Research Center, 1100 Fairview ave N, Seattle, WA 98109*

The adaptive immune system in vertebrates consists of highly diverse immune receptors to mount specific responses against a multitude of pathogens. A central feature of the adaptive immune system is the ability to form a memory to act more efficiently in future encounters with similar pathogens. However, memory formation especially in B-cells is one of the least understood cell fate decisions in the immune system. Here, we present a framework to characterize optimal strategies to store memory in order to maximize the utility of immune response to counter evolving pathogens throughout an organism's lifetime. To do so, we have incorporated the kinetics and energetics of memory response as ingredients of non-equilibrium decision-making between an adaptive exploration to mount a specific and novel response or exploitation of existing memory that can be activated rapidly yet with a reduced specificity against evolved pathogens. To achieve a long-term benefit for the host, we show that memory generation should be actively regulated and dependent on immune receptors' affinity, with a preference for cross-reactive receptors with a moderate affinity against pathogens as opposed to high affinity receptors— a recipe that is consistent with recent experimental findings [1, 2]. Moreover, we show that the specificity of memory should depend on the organism's lifespan, and shorter-lived organisms with fewer pathogenic encounters throughout their lifetime should store more cross-reactive memory. Overall, our framework provides a baseline to gauge the efficacy of immune memory formation in light of an organism's coevolutionary history with pathogens.

I. INTRODUCTION

Adaptive immunity in vertebrates develops during the lifetime of an organism to battle a multitude of evolving pathogens. The central actors in our adaptive immune system are diverse B- and T-cells, whose unique surface receptors are generated through genomic rearrangement, mutation, and selection [3]. The diversity of receptors allows the immune system to mount specific responses against diverse pathogens. B-cell receptors (BCRs) in particular can specialize through a process of affinity maturation, which is a form of *somatic Darwinian evolution* within an individual to enhance the affinity of BCRs to pathogens. During affinity maturation, B-cells aggregate in lymph nodes to form germinal centers, where they undergo somatic hypermutation, compete for resources like T-cell help, and are selected based on their binding affinity to pathogens. Several rounds of mutation and selection can increase binding affinities of receptors up to 10-1000 fold [4, 5].

Beside receptor diversity, immune cells also differentiate and specialize to take on different roles. Importantly, both B- and T-cells can differentiate into different forms of memory cells (e.g. IgG or IgM memory in B-cells, and CD4 or CD8 memory in T-cells) [6, 7]. Memory is central to adaptive immunity through which our immune system can be trained during an infection or by vaccination against a pathogen to respond more readily upon re-exposure [3]. Memory responses are highly efficient since memory cells can be reactivated faster than naïve cells and can mount a more robust response to an infection [8–11].

Memory generation is a form of cell fate decision in the immune system, which can occur at different stages of an immune response. In B-cells, activated naïve cells can differentiate into antibody-secreting long-lived plasma cells or a T-cell independent un-hypermutated memory cell, or they can initiate a germinal center [12]. B-cells that enter germinal centers differentiate during affinity maturation into high affinity plasma cells, which act as antibody factories, or T-cell dependent long-lived memory cells that circulate in the blood for antigen surveillance; see schematic Fig. 1.

The basis for differentiation of B-cells into memory, especially during affinity maturation, is among the least understood in cell fate decision-making in the immune system [12]. A long-standing view was that memory is continuously produced during affinity maturation [13]. Memory receptors often have lower affinity compared to plasma cells [14], and therefore, if memory were to be generated continuously it should be able to proliferate without strong affinity dependent selection [4, 12]. However, recent experiments indicate that memory differentiation is highly regulated [1, 2, 15–17], reflecting a temporal switch in germinal centers that preferentially pro-

*correspondence should be addressed to Armita Nourmohammad: armita@uw.edu

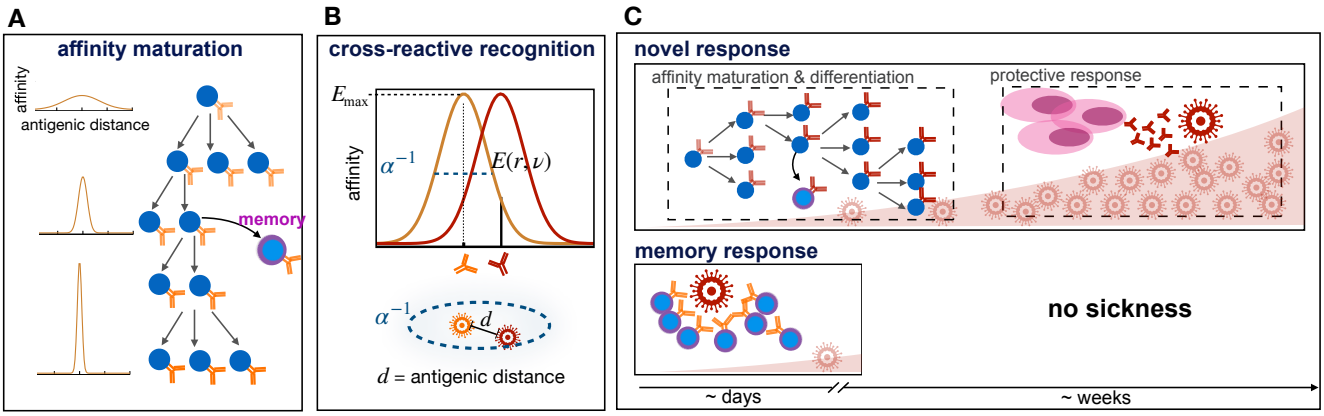

FIG. 1: **Immune memory or naïve response upon infection.** **(A)** Schematic shows affinity maturation in germinal centers(right), where B-cell receptors acquire mutations and undergo selection, resulting in an increase in their affinity to an antigen (from light to dark receptors), indicated by the sharpening of receptors' affinity profiles (on left). **(B)** Affinity profile $E_{\alpha,\theta}(r, \upsilon) \sim \alpha \exp[-(\alpha d)^{\theta}]$ of a receptor r is shown in orange as a function of the distance $d = \upsilon_r^* - \upsilon$ in the antigenic shape space, between the receptor's cognate antigen υ_r^* (orange) and another target υ (red). The affinity of a receptor decays with increasing distance between targets and its cognate antigen. The antigenic range over which a receptor is reactive inversely depends on its specificity α. The shape of the binding profile is tuned by the factor θ, here shown for $\theta = 2$. **(C)** Upon infection, the immune system can initiate a novel response (top) or a memory response (bottom). A novel B-cell response could involve affinity maturation to generate memory or high affinity plasma cells (pink) that can secrete antibodies to battle the pathogen. A novel response can take 1-2 weeks, during which pathogen can replicate within a host and a patient can show symptoms from the disease (top). If the host carries memory from a previous infection or vaccination (bottom), the immune system can robustly and rapidly activate a memory response to battle the infection. Cross-reactive memory can be activated against slightly evolved forms of an antigen (dark red), drifted from to the primary infection (orange).

duces memory at early stages and plasma at later stages of affinity maturation [16]. Importantly, it has been suggested that this regulation is done through a T-cell dependent pathway in germinal centers, in which T-cell help suppresses the expression level of the transcription factor Bach2, whose abundance predisposes B-cells to become memory [1, 2]. This active regulation introduces an affinity-dependent cell fate decision, as B-cells with lower affinities attract less T-cell help and hence, maintain a higher level of Bach2, leading to a preferential selection of low-affinity cells to the memory compartment.

Low affinity memory may be at a disadvantage in mounting a protective immune response since immune-pathogen recognition is largely determined by the binding affinity between an immune receptor and antigenic epitopes. On the other hand, as in most molecular interactions, immune-pathogen recognition is cross-reactive, which would allow memory receptors to recognize slightly evolved forms of the antigen, in response to which they were originally generated. Therefore, the long-term utility of a memory cell depends on balancing the receptor's affinity and cross-reactivity to protect against moderately evolved forms of a pathogen [18, 19].

The program for differentiation of immune cells to memory should be viewed in light of the utility of memory during an organism's lifetime. Here, we present a theoretical framework that incorporates the kinetics and energetics of memory responses as ingredients of a non-equilibrium decision-making process to choose a memory or a novel response upon re-exposure to evolving pathogens. We quantify the characteristics of memory, including its relative cross-reactivity and speed of activation, that would confer the highest long-term utility for the immune system in countering pathogens with different evolutionary rates. We propose that the hard-wired affinity-dependent regulatory measures during differentiation of memory could be understood as a way to optimize the long-term utility of immune memory against evolving pathogens.

Indeed, previous work on theoretical modeling of cellular differentiation together with experiments has been instrumental in understanding immune memory generation; e.g. see reviewed work in [20, 21]. For example, mechanistic models have indicated the importance of signal integration at the cellular level [22] and the relevance of stochastic effects at the population level [23], to explain heterogeneous cell fate decisions for the generation of memory. Our statistical framework, on the other hand, aims to characterize high-level features for an optimal memory strategy, without relying on mechanistic details of the underlying process, some of which are at least partially unknown [24, 25]. In the case of the immune system, statistical models have provided an intuition for how an immune repertoire should be organized to optimally counter diverse pathogens [19, 26–29]. In a similar fashion, optimal memory strategies identified by our model provide a baseline to gauge the performance of real immune systems in storing and utilizing memory.

II. MODEL

Kinetics of naïve and memory immune response. Upon encountering a pathogen, the adaptive immune system mounts a response by activating the naïve repertoire (i.e., a novel response) and/or by triggering previously stored immune receptors in the memory compartment. A memory receptor often shows a reduced affinity in interacting with an evolved form of the pathogen. Nonetheless, memory plays a central role in protecting against re-infections since even a suboptimal memory can be kinetically more efficient than a naïve response, both in B-cells [10] and T-cells [30, 31]. First, memory cells are fast responders and initiate cell division about $\tau_0 \approx 1-2$ days before naïve cells [9, 10, 32]. Second, the number of memory cells that are recruited to proliferate and differentiate to effector cells is $b \approx 2-3$ times larger than the number of naïve cells [9, 10]. Once recruited however, memory and naïve cells have approximately a similar doubling time of about $t_{1/2} \approx 0.5-2$ days [9, 33]. Putting these kinetic factors together, we can define an effective deliberation time τ for the naïve population to reach an activity level (i.e., a population size) comparable to the memory. Assuming an exponential growth during early stages of memory and naïve proliferation, the deliberation time can be estimated in terms of the kinetic factors by $\tau = \tau_0 + t_{1/2} \ln b / \ln 2$ and it is within a range of $\tau \approx 1.5-5$ days; see Fig. 1.

Energetics of immune recognition. Recognition of a pathogen by an immune receptor is mediated by the affinity of the molecular interactions $E(r, \upsilon)$ between the immune receptor r and the antigenic epitope υ from the pathogen. As in most molecular interactions, immune-pathogen recognition is *cross-reactive*: each immune receptor can interact with several antigens and each antigen can be recognized by more than one immune receptor; see Fig. 1. We describe immune-pathogen recognition in an immune *shape space* [26], which reflects the intuition that specific protein-protein interactions are determined by the 3D structure of the interface, i.e., their complementary shapes. Receptors located near each other in shape space can recognize similar antigens, and in the complementary space, antigens that are close to each other can be recognized by the same immune receptor. We assume that each immune receptor r has a cognate antigen υ_r^* against which it has the highest affinity. We express the binding affinity between a receptor r and an arbitrary target antigen υ in terms of the antigenic distance $d_r(\upsilon) = \|\upsilon - \upsilon_r^*\|$ between the receptor's cognate antigen υ_r^* and the target υ: $E(r, \upsilon) \equiv E(d_r(\upsilon))$. This distance-dependent binding affinity is measured with respect to the affinity of unspecific antigen-receptor interactions, sufficient to trigger a generic naïve response.

Physico-chemical constraints in protein structures can introduce a tradeoff between immune receptors' affinity and cross-reactivity (i.e., ability to equally react to multiple targets). Although we lack a systematic understanding of these structural constraints, affinity-specificity tradeoffs have been reported repeatedly for B-cells and antibodies. Prior to affinity maturation, the structure of naïve receptors is relatively flexible whereas hypermutations often reconfigure the active sites of a receptor and make them more specific so that they match their target antigens like a lock and key [34, 35]. As a result, the IgM class of antibodies, which are the first line of defense in B-cell response, often have low affinities, yet they are cross-reactive and can recognize mutated forms of the same epitope. On the other hand, the high affinity IgG class of antibodies, which are the late outcomes of affinity maturation in germinal centers, have higher affinities but bind very specifically to their cognate antigen [35]. Broadly neutralizing antibodies (bNAbs) are exceptions to this rule since they often have high potency and can react to a broad range of viral strains. However, bNAbs often react to vulnerable regions of a virus where escape mutations are very deleterious [36]. In other words, the majority of bNAbs are not cross-reactive per se, but they are exceptionally successful in targeting conserved epitopes in otherwise diverse viral strains. Nevertheless, an affinity-specificity tradeoff has been reported for a bNAb against the hemagglutinin epitope of influenza [37].

We use a simple functional form to qualitatively capture the tradeoff between cross-reactivity and affinity of antigen-receptor binding interactions: We assume that the binding affinity of a receptor r to an antigen υ depends on the antigenic distance $d_r(\upsilon) = \|\upsilon - \upsilon_r^*\|$ through a kernel with a specificity factor α and a shape factor θ such that, $E(r, \upsilon) \equiv E_{\alpha,\theta}(d_r(\upsilon)) \sim \alpha \exp[-(\alpha\|\upsilon - \upsilon_r^*\|)^\theta]$, with $\theta \geq 0$. This affinity function defines a receptor's binding profile over the space of antigens. The width of this binding profile is set by the inverse of the specificity factor $1/\alpha$ (Fig. 1). As specificity α increases (or cross-reactivity $1/\alpha$ decays), the binding affinity profile sharpens and binding becomes more restrictive to antigens closer to the receptor's cognate antigen. Moreover, the absolute strength of binding to the cognate antigen (i.e., a receptor's maximum affinity) increases with specificity α, resulting in a tradeoff between affinity and cross-reactivity. The parameter θ tunes the shape of the receptor's binding profile $E_{\alpha,\theta}(d_r(\upsilon))$, resulting in a flat function (i.e., no tradeoff) for $\theta = 0$, a double-sided exponential function for $\theta = 1$, a Gaussian (bell-curve) function for $\theta = 2$, and top-hat functions for $\theta \gg 2$.

Structural constraints and molecular features of protein receptors define a bound on the minimum cross-reactivity or equivalently, a maximum specificity $\alpha_{\max}$, achievable by a receptor. Using this bound, we define rescaled specificity $\hat{\alpha} \equiv \alpha/\alpha_{\max}$ to characterize the energetics of an immune response in a dimensionless form.

Immune response to evolving pathogens. Upon primary infection (i.e., an encounter with a novel pathogen)

naïve immune receptors with moderate affinity are activated to develop a specific response through affinity maturation (Fig. 1). Since the naïve repertoire is diverse enough to contain receptors of moderate affinity against different antigens, we assume that the affinity of responsive naïve receptors, and hence, the strength of a primary immune response to be approximately the same for all pathogens. This simplification becomes less accurate as the immune system ages and the supply of effective receptors become more scarce.

Following a naïve response to a primary infection and the subsequent affinity maturation, the immune system stores memory cells with an enhanced affinity to use them against future infections [3]; see Fig. 1. Therefore, the cognate antigen $\upsilon^*_{r_m}$ for a given memory receptor r_m is an epitope derived from the primary infection that led to the formation of memory, which we denote by υ_0 with a subscript that indicates round of infection. Thus, the binding profile $E_{\alpha,\theta}(r_m,\upsilon)$ of the memory receptor r_m is peaked around the primary antigenic epitope $\upsilon^*_{r_m} = \upsilon_0$ (Fig. 1). As pathogens evolve globally to escape the immune challenge, drugs or vaccination, they drift away from the primary antigen in antigenic space. We model this antigenic shift as a diffusion in shape space whereby a reinfecting pathogen at the ith round of infection υ_i is *on average* at a distance $\delta = \sqrt{\langle \|\upsilon_i - \upsilon_{i-1}\|^2\rangle}$ from the previous infection υ_{i-1}. This antigenic shift is proportional to the rate of pathogen evolution ζ_υ and the average time between infections $\Delta t = t_i - t_{i-1}$, such that $\delta \propto \zeta_\upsilon\sqrt{\Delta t}$. A cross-reactive memory can mount a response to an evolved antigen, yet with a reduced affinity that decays with antigenic shift; see Fig. 1. It should be noted that the minimum level of receptor's cross-reactivity (or maximum specificity) $(\alpha_{\max})^{-1}$ defines a natural scale against which we can measure antigenic divergence δ and hence, form a dimensionless measure of antigenic divergence $\hat{\delta} \equiv \delta/(\alpha_{\max})^{-1}$.

Immune-pathogen recognition depends both on the binding affinity $E_{\alpha,\theta}(r,\upsilon)$ and the encounter rate $\gamma_\upsilon(t)$ between an immune receptor r and the antigen υ at a given time t. The encounter rate $\gamma_\upsilon(t)$ depends on the abundance of the antigen and the immune receptor, and hence, can vary during an infection within a host. The probability that a receptor r encounters and binds to an antigen υ in a short time interval $[t, t+dt]$ can be expressed by, $\rho(r,\upsilon,t)\mathrm{d}t = \gamma_\upsilon(t)E_{\alpha,\theta}(r,\upsilon)\mathrm{d}t$; a similar notion of encounter rate has been previously used in ref. [19]. A memory response in an individual is triggered through recognition of an antigen by a circulating memory receptor. If no such recognition occurs during the deliberation time $\tau \approx 1.5 - 5$ days, the immune system initiates a naïve response. Therefore, the probability that an antigen is recognized through a novel naïve response $P^{(0)}_{\text{recog.}}$ can be expressed as the probability of the antigen not being recognized $1 - P^{(m)}_{\text{recog.}}$ by an available memory receptor r_m over the deliberation period τ,

$$\begin{aligned} P^{(0)}_{\text{recog.}}(\upsilon) &= 1 - P^{(m)}_{\text{recog.}}(r_m,\upsilon) \\ &= e^{-\int_0^\tau \rho(\upsilon,t)\mathrm{d}t} = e^{-E_{\alpha,\theta}(r_m,\upsilon)\Gamma(\upsilon,\tau)} \end{aligned} \quad (1)$$

where $\Gamma(\upsilon,\tau) = \int_0^\tau \gamma_\upsilon(t)dt$ is the expected number of pathogenic encounters over the deliberation time τ and depends on the accumulated pathogenic load, as pathogens proliferate in the absence of an effective memory prior to a naïve response. Here, we have assumed that the affinity of the memory receptor does not change over the response time, which is a simplification since memory receptor can undergo limited affinity maturation [38, 39]. To further simplify, we also assume that the accumulated pathogenic load is independent of the type of the pathogen $\Gamma(\upsilon,\tau) \equiv \Gamma(\tau)$. As pathogens evolve away from the primary infector, the binding affinity $E_{\alpha,\theta}(r_m,\upsilon)$ of the stored memory receptor r_m, and hence, the probability to mount a memory response $P^{(m)}_{\text{recog.}}(r_m,\upsilon,\tau)$ decays.

The deliberation time prior to a novel response provides a window for memory to react with an antigen and mount an immune response by initiating an irreversible cascade of downstream events. Although initiation of this pathogenic recognition can be modeled as an equilibrium process, the resulting immune response is a non-equilibrium and an irreversible process, the details of which are not included in our model.

Decision-making to mount a memory or naïve response. The decision to mount a naïve or a memory response depends on the energetics and the kinetics of the immune machinery, including the cross-reactivity of memory to recognize evolved pathogens and the deliberation time to mount a naïve response upon infection — we refer to these choices as *memory strategies*. We expect that the biochemical machinery involved in making this decision upon an infection has been fine-tuned and selected over evolutionary time scales in order to utilize immune memory and mount an effective response against evolving pathogens. Here, we introduce a decision-making framework to identify optimal memory strategies for the adaptive immune system to achieve its highest *long-term utility*. This framework does not assume that the immune system, as we know it, should be optimized, rather it defines a baseline against which memory strategies of the actual immune system could be gauged.

In the theory of decision-making, a rational decision-maker chooses between two possible actions $a \in$ {naïve, memory} each contributing a utility U_a. The action of a rational decision-maker should follow an optimal distribution $Q(a)$, which maximizes the expected utility while satisfying the constraints in processing new information, e.g. due to prior preferences for either of the actions $Q_0(a)$ [40, 41]. The optimal solution for a rational yet constrained decision follows (see SI),

$$Q(a) = \frac{1}{Z}Q_0(a)e^{\beta U_a} \quad (2)$$

where $Z = \sum_a Q_0(a)e^{\beta U_a}$ is a normalization factor and β is the efficacy of information processing. As β increases, a rational decision-maker more readily chooses the action with the highest utility. Moreover, if the prior distribution is uniform across actions (i.e., no prior preference), rational decisions follow a maximum entropy distribution $Q(a) \sim \exp[\beta U_a]$ [42], which we will use in the rest of our analysis.

A decision to mount a memory or naïve response $Q(a)$ based on their respective utilities (eq. 2) should be consistent with the biophysical description of the immune response through recognition of an antigen by either of these cell types (eq. 1). By equating these two descriptions of an immune response (eqs. 1, 2) we can specify the utility gain associated with mounting a memory or a naïve response in terms of the biophysics and kinetic of receptor-antigen interactions. Importantly, in the regime that memory is efficient and being used (i.e., $P^{(m)}_{\text{recog.}} \gg 0$), the utility gain by a memory or a naïve response to an evolved antigen υ_i at an antigenic distance $d_i = \|\upsilon_i - \upsilon_0\|$ from the memory receptor's cognate antigen $\upsilon^*_{r_m} \equiv \upsilon_0$ follows (see SI),

$$\begin{aligned} U_{\text{mem}}(\|\upsilon_i - \upsilon_0\|; \alpha, \theta) &= U_{\text{naïve}} + E_{\alpha,\theta}(r_m, \upsilon_i) \\ &= -\Omega(\Gamma_\tau) + E_{\alpha,\theta}(\|\upsilon_i - \upsilon_0\|) \end{aligned} \tag{3}$$

Here, we introduce the cost for deliberation $\Omega(\Gamma_\tau)$ as the negative utility of the nïve response $U_{\text{naïve}}$. Deliberation cost $\Omega(\Gamma_\tau)$ is a monotonically increasing function of the cumulative pathogen load Γ_τ and reflects the damage (cost) incurred by pathogens as they proliferate during the deliberation time τ prior to activation of the novel naïve response; see Fig. 1. It is important to note that the difference in the memory and the naïve utility $\Delta U = U_{\text{mem}} - U_{\text{naïve}}$ determines the decision to mount either of these responses.

The same consistency criteria between decision-making (eq. 2) and cellular recognition (eq. 1) indicates that the information processing factor β in eq. 2 should be equal to the accumulated pathogenic load $\Gamma(\tau)$ during the deliberation period τ: $\beta = \Gamma(\tau)$. A longer deliberation, which on one hand leads to the accumulation of pathogens, would allow the immune system to exploit the utility of a usable memory (i.e., process information), even if the memory has only a slight advantage over a responsive naïve receptor. As a result, we refer to β as the *deliberation factor*. Moreover, this analogy relates the efficacy of information processing β, which plays the role of inverse temperature in thermodynamics, and the total accumulated pathogenic load $\Gamma(\upsilon, \tau)$, which acts as the sample size for memory receptors as they encounter and accumulate information about pathogens. Interestingly, previous work has drawn a similar correspondence between the inverse temperature in thermodynamics and the effect of sample size on statistical inference [43].

The deliberation factor in the immune system should be bounded $\beta \leq \beta_{\max}$ in order for the organism to survive new infections by mounting a novel response that can suppress an exponentially replicating pathogen before it overwhelms the host. Using this bound, we define rescaled deliberation factor $\hat{\beta} \equiv \beta/\beta_{\max} \leq 1$ to characterize the kinetics of an immune response in a dimensionless fashion.

If the immune system decides to mount a memory response against an evolved antigen υ_i, the binding profile of memory against the target pathogen remains unchanged and equal to the profile $E_{\alpha,\theta}(r_{\upsilon_0}, \upsilon)$ against the primary infection υ_0. However, if the immune system mounts a naïve response, a new memory receptor r_{υ_i} would be generated with a binding profile $E_{\alpha,\theta}(r_{\upsilon_i}, \upsilon)$, centered around the latest infection υ_i. As a result, the expected binding profile $\hat{E}^{(i)}_{\alpha,\theta}(\upsilon)$ at the ith round of infection is an interpolation between the profiles associated with memory and naïve response, weighted by the likelihood of each decision (eq. 1),

$$\hat{E}^{(i)}_{\alpha,\theta}(\upsilon) = P^{(m)}_{\text{recog.}}(r_{\upsilon_0}, \upsilon_i)E_{\alpha,\theta}(r_{\upsilon_0}, \upsilon) + P^{(0)}_{\text{recog.}}(\upsilon_i)E_{\alpha,\theta}(r_{\upsilon_i}, \upsilon) \tag{4}$$

The expected binding profile at the ith round of infection $\hat{E}^{(i)}_{\alpha,\theta}(\upsilon)$ (eq. 4) deviates from the optimal profile centered around the infecting pathogen $E_{\alpha,\theta}(r_{\upsilon_i}, \upsilon)$ (i.e., memory profile stored following a novel response); see Fig. 1. This deviation arises because an energetically sub-optimal memory response can still be favorable when time is of an essence and the decision has to be made on the fly with short deliberation. This tradeoff between the kinetics and the energetics of immune response results in a *non-equilibrium decision-making* [44] by the immune system. In analogy to non-equilibrium thermodynamics, we express this deviation as a dissipative cost of memory response $K_{\text{diss}}(t_i; \alpha, \theta)$ at the ith round of infection (time point t_i), which we quantify by the Kullback-Leibler distance between the expected and the optimal binding profiles, in units of the deliberation factor β,

$$\begin{aligned} K_{\text{diss}}(t_i; \alpha, \theta) &= \frac{1}{\beta} D_{KL}\left(\hat{E}^{(i)}_{\alpha,\theta}(\upsilon) || E_{\alpha,\theta}(r_{\upsilon_i}, \upsilon)\right) \\ &= \frac{1}{\beta} \sum_{\text{antigens: } \upsilon} \hat{E}^{(i)}_{\alpha,\theta}(\upsilon) \log\left[\frac{\hat{E}^{(i)}_{\alpha,\theta}(\upsilon)}{E_{\alpha,\theta}(r_{\upsilon_i}, \upsilon)}\right] \end{aligned} \tag{5}$$

where we ensure that binding profiles are normalized over the space of antigens. The dissipation K_{diss} measures the sub-optimality (cost) of the mounted response through non-equilibrium decision-making and quantifies deviation from an equilibrium immune response [44].

An optimal memory strategy should be chosen such that it maximizes the expected utility of the immune response $\langle U \rangle = U_{\text{mem}} P^{(m)}_{\text{recog.}} + U_{\text{naïve}} P^{(0)}_{\text{recog.}}$, while minimizing the dissipation cost due to the non-equilibrium response K_{diss}, over the lifetime of an organism. To infer an optimal strategy, we introduce net utility U_{net} that accounts for the tradeoff between the expected utility and

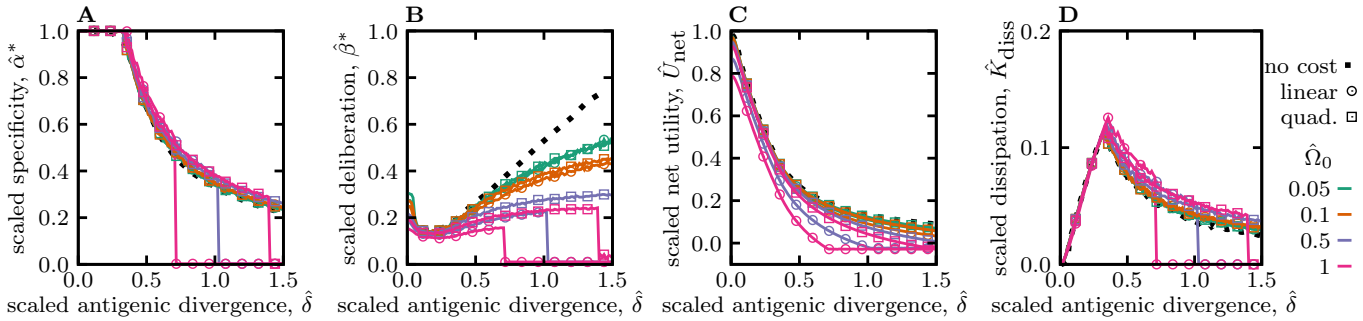

FIG. 2: **Optimal memory strategies against evolving pathogens. (A)** and **(B)** show the optimal rescaled specificity $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\max}$ and rescaled deliberation factor $\hat{\beta}^* \equiv \beta^*/\beta_{\max}$ as a function of the rescaled antigenic divergence per infection $\hat{\delta} = \delta/(\alpha_{\max}^{-1})$, respectively. **(C)** and **(D)** show the scaled net utility $\hat{U}_{\rm net} \equiv U_{\rm net}/E_{\max}$ (eq. 7) and dissipation $\hat{K}_{\rm diss} \equiv K_{\rm diss}/E_{\max}$ (eq. 5) per round of infection as a function of the antigenic divergence $\hat{\delta}$. Rescaling by $E_{\max}$ sets the magnitude of net utility for a response to conserved antigens (with $\hat{\delta} = 0$), and in the limit of zero deliberation cost $\Omega \to 0$, to 1. Colors / markers indicate different naïve cost functions for deliberation, including no-cost $\hat{\Omega} \equiv \Omega/E_{\max} = 0$, linear cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}$, and quadratic cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}^2$, with varying amplitudes. Simulation parameters: $\alpha_{\max} = 4$, $\beta_{\max} = 10$, and $\theta = 2$. Results for other shape parameters θ and specificity thresholds $\alpha_{\max}$ are shown in Figs. S1, S2, respectively.

dissipation at a given round of infection at time point t_i,

$$U_{\rm net}(t_i;\alpha,\beta,\theta) = \langle U_{\alpha,\beta,\theta}(t_i)\rangle - K_{diss}(t_i;\alpha,\theta) \qquad (6)$$

Net utility can be interpreted as the extracted (information theoretical) work of a rational decision-maker that acts in a limited time, and hence, is constantly kept out of equilibrium [44]. We infer the optimal memory protocol (i.e., the optimal memory specificity α^* and deliberation factor β^*) by maximizing the total net utility of memory responses throughout the lifetime of an organism,

$$(\alpha^*,\beta^*) = \underset{\alpha,\beta}{\rm argmax} \sum_{i:\ \text{infections}} U_{\rm net}(t_i;\alpha,\beta,\theta). \qquad (7)$$

III. RESULTS

Efficient immune memory balances specificity and speed. An optimal immune protocol balances the specificity of memory and deliberation time that it takes to mount a novel response in a host. On the one hand, a longer deliberation would allow the immune system to exploit a stored cross-reactive memory, even with a low affinity, to recognize an evolved pathogen (eq. 1). On the other hand, if memory is inefficient, long deliberations would allow pathogens to further proliferate, incurring a larger cost Ω to a host prior to activation of a novel response.

The extent of cross-reactivity and deliberation needed for the memory to react to pathogens should be set by the amount of pathogenic evolution and more specifically, the antigenic divergence $\delta \equiv \sqrt{\langle \|\upsilon_i - \upsilon_{i-1}\|^2\rangle}$ that a pathogen traces between two infections. Therefore, we set to find an optimal immune protocol (i.e., specificity α^* and deliberation β^*) by maximizing the net utility $U_{\rm net}$ of an immune system (eq. 7) that is trained to counter pathogens with a given (scaled) antigenic divergence $\hat{\delta} \equiv \delta/(\alpha_{\max})^{-1}$ over rounds of infection; see SI for details on the optimization procedure.

We can describe immune strategies in three different regimes, corresponding to optimal strategies against slowly evolving pathogens ($\hat{\delta} \leq 20\%$), moderately evolving pathogens ($\hat{\delta} \simeq 20\% - 80\%$) and extremely rapidly evolving pathogen ($\hat{\delta} > 80\%$).

To battle slowly evolving pathogens ($\hat{\delta} \leq 20\%$) an optimal immune system stores highly specific memory receptors, with a specificity that approaches the upper bound $\alpha_{\max}$; see Fig. 2A and Figs. S1, S2. Importantly, the dependency of optimal specificity on antigenic divergence is insensitive to the cost of deliberation Ω prior to mounting a naïve response (Fig. 2A), the shape factor θ for the specificity profile (Fig. S1), and the specificity threshold $\alpha_{\max}$ (Fig. S2). For relatively conserved pathogens ($\hat{\delta} \simeq 0$), the highly specific memory (with $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\max} \simeq 1$) stored from a previous infection remains perfect and therefore, immune system maintains a moderate level of deliberation to exploit this efficient memory during infections. However, as antigenic divergence grows, specific memory becomes less effective against future infections and therefore, the immune system reduces the deliberation factor to allow a timely novel response, once memory becomes inefficient (Figs. 2B, S1, S2). The magnitude of deliberation decays as the cost of deliberation Ω increases but its overall dependency on antigenic divergence remains comparable for different cost functions (shown in Fig. 2B for zero cost, and cost functions that grow linearly and quadratically with deliberation factor β). Overall, the net utility of the stored memory in response to slowly evolving pathogens is high (Figs. 2C, S1, S2), while its dissipation remains small $K_{\rm diss} \simeq 0$ (Figs. 2D, S1, S2). Therefore, in analogy

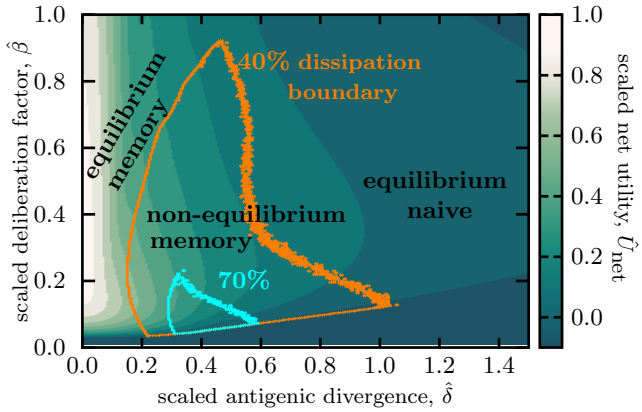

FIG. 3: **Antigenic divergence determines modes of immune strategies.** The heat map shows the expected net utility $U_{\rm net}$ (eq. 7) per round of infection for an immune system with an optimal specificity $\hat{\alpha}^*$, as a function of rescaled antigenic divergence (per encounter) $\hat{\delta}$ and deliberation factor $\hat{\beta}$. Boundaries indicate different levels of dissipation, with the orange and blue boundaries encompassing a region with $\geq 40\%$ and $\geq 70\%$ of the maximum dissipation $K_{\rm max}$, respectively. The magnitudes of dissipation and net utility of an optimal memory classifies immune response to into three indicated modes: (i) equilibrium memory, (ii) non-equilibrium memory and (iii) equilibrium naïve. Simulation parameters: linear deliberation cost function $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}$ with $\hat{\Omega}_0 = 0.1$, $\alpha_{\rm max} = 4$, $\beta_{\rm max} = 10$, and $\theta = 2$.

to thermodynamics, we term this immune strategy with low dissipation as *equilibrium memory response*; Fig. 3.

To battle moderately evolving pathogens (with $\hat{\delta} \simeq 20\% - 80\%$), an optimal immune system stores cross-reactive memory (i.e., with a lower specificity $\hat{\alpha}$) that can recognize moderately evolved form of the primary antigen (Figs. 2A, S1, S2). However, cross-reactive receptors tend to have lower affinities [34, 35], which could lead to deficient responses against antigens. Importantly, activation of energetically sub-optimal yet cross-reactive memory could be detrimental as it may hinder a stronger novel response without providing protective immunity to the host— a deficiency known as the original antigenic sin [45, 46]. An optimal immune system can mitigate this problem by using kinetic optimization to tune the deliberation factor β in order to avoid an elongated memory engagement prior to a naïve response. This optimization results in a smaller deliberation factor β (i.e., a faster naïve response) compared to the scenario with slowly evolving pathogens, yet a long enough deliberation to allow the energetically suboptimal memory to react to an infection, whenever feasible (Figs. 2B, S1, S2). With this kinetic optimization, the immune system can utilize cross-reactive memories through multiple rounds of infection (Fig. S3), yet with a declining net utility as pathogens evolve away from the primary infection (Figs. 2C, S1, S2). The prominent memory response to moderately evolving pathogens is dissipative with $K_{\rm diss} \gg 0$ (Figs. 2D, S1, S2), and in analogy with thermodynamics, we term this dissipative immune strategy as *non-equilibrium memory response*; Fig. 3.

For extremely rapidly evolving pathogens ($\hat{\delta} > 80\%$), the immune system would not be able to store an efficient memory to battle future encounters, and hence, each infection would trigger a novel naïve response — the reduced net utility of memory and the decay of memory usage in this regime are shown in Figs. 2C, S1, S2, and Fig. S3, respectively. Without a protective memory, a novel response is triggered to counter each infection and it maturates specifically around the infecting pathogen, resulting in a non-dissipative naïve-dominated immune response with $K_{\rm diss} \simeq 0$, which we term *equilibrium naïve response*; Fig. 3.

It should be noted that when the cost of deliberation Ω is very high, utilizing memory against pathogens with relatively high evolutionary rates becomes highly unfavorable. In this extreme case, the immune system switches into a state where it invariably mounts a novel response upon an infection (Fig. S3), and it assures that memory is not utilized by setting the parameters for specificity α and deliberation β to zero (Fig. 2A, B).

Optimal memory production is regulated to store receptors with intermediate affinity and cross-reactivity. The results shown in Figs. 2, 3 indicate that a rational decision to become a memory or a plasma cell during an immune response should depend on the affinity of a cell's receptors and it should not be a stochastic choice with a constant rate throughout affinity maturation. Indeed, cell fate decision for B-cells during affinity maturation is highly regulated and dependent on receptors' affinity [1, 2, 16, 47–49]. Recent experiments have demonstrated that memory generation is highly correlated with the activity of the transcription factor *Bach2* whose expression level is negatively regulated with the abundance of helper CD4$^+$ T-cells [1, 2, 48]. As the affinity of B-cell receptors increases during affinity maturation, more CD4$^+$ T-cells are recruited to germinal centers, resulting in suppression of *Bach2* and a hence, a decline in production of memory cells [1, 2, 48]. In other words, our adaptive immune system has encoded a negative feedback mechanism to store memory with intermediate affinity and cross-reactivity to suppress the production of highly specific memory, which is likely to be impotent against evolved pathogens in future infections.

In addition, Figs. 2, 3 also show that an optimal immune system should provide a deliberation time for a cross-reactive memory to mount an immune response, before activation of novel response. This prediction is also in accordance with the kinetics of immune response, where if available, memory can react and become effective by about $\tau \approx 1.5 - 5$ days faster than a novel adaptive immune response [9, 10, 32]. This

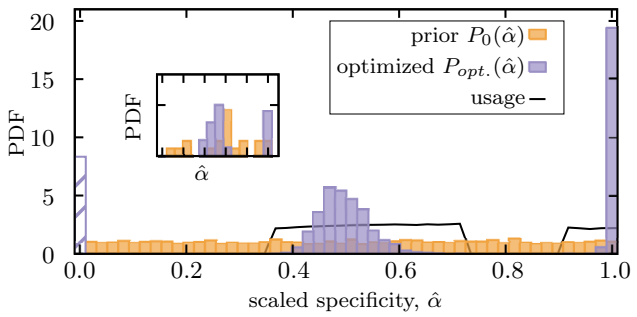

FIG. 4: **Mixed memory strategy against pathogens with distinct evolutionary rates.** Distribution of scaled optimized specificities $\hat{\alpha}^*$ for functional memory (purple) is shown for an immune system with a fixed deliberation factor $\hat{\beta} = 0.2$, in which a mixture strategy with a bimodal distribution of specificities $P(\hat{\alpha})$ is established to counter pathogens with a broad range of antigenic divergences. The dashed bar indicates stored memory with specificity $\alpha = 0$, which is not further used in response to infections. The solid line indicates the probability P_{usage} that a stored memory with a given specificity is utilized in future infections (SI). Optimization is done by maximizing the net utility of immune response averaged over encounters with 1000 independently evolving antigens with (scaled) antigenic divergences drawn uniformly from a range $\hat{\delta} \in (0, 1.6)$ (SI). The distribution shows the ensemble statistics of functional memory accumulated from 200 independent optimizations, each starting from a flat prior for specificities (orange). The insert shows the optimized mixture strategy for one optimization with 3000 steps. Simulation parameters: $\alpha_{\max} = 4$, $\beta_{\max} = 10$, and $\theta = 2$.

deliberation time provides a window for cross-reactive (and lower affinity) memory to recognize an antigen and initiate a downstream cascade of immune reactions. Indeed, if such deliberation were not programmed in the immune system, memory storage would become futile after only a few rounds of encounters with an evolved pathogen. The fact that the adaptive immune system has set up machinery with intermediate specificity $\hat{\alpha}$ and deliberation factor $\hat{\beta}$ suggests that this immune strategy has undergone evolutionary optimization to battle evolving pathogens.

A mixture of memory isotypes is necessary to counter pathogens with a broad range of evolutionary rates. The decision to trigger an equilibrium or a non-equilibrium memory response depends on the extent of antigenic divergence that an immune system is trained to cope with (Figs. 2, 3, S1, S2). Equilibrium memory is highly effective (i.e., it has high net utility) against relatively conserved pathogens, however, it fails to counter evolving pathogens (Fig. 2C). On the other hand, cross-reactive non-equilibrium memory is more versatile and can counter a broader range of evolved pathogens but at a cost of reduced net utility in immune response; Figs. 2C, 3, S1, S2.

An optimal immune system should have memory strategies to counter pathogens with varying evolutionary rates, ranging from relatively conserved pathogens like chickenpox to rapidly evolving viruses like influenza. We use our optimization protocol to find such memory strategies that maximize the net utility of an immune system that encounters evolving pathogens with antigenic divergences uniformly drawn from a broad range of $\hat{\delta} \in [0\ 1.6]$; see SI. This optimization results in a bimodal distribution of optimal specificity for functional memory receptors $P(\alpha)$, with separated peaks corresponding to equilibrium ($\hat{\alpha} \sim 1$) and non-equilibrium ($\hat{\alpha} \sim 0.5$) memory (Figs. 4, S4). This result suggests that specific and cross-reactive memory strategies are complementary modes of immune response that cannot substitute each other. Moreover, non-equilibrium memory tends to be flexible and moderate values of cross-reactivity $1/\hat{\alpha}$ can counter a broad range of antigenic divergences, without a need for fine-tuning. Therefore, upon production of memory, an optimal immune system should harvest both specific equilibrium memory and cross-reactive non-equilibrium memory, as it does not have a priori knowledge about the evolutionary rate of the infecting pathogen.

Interestingly, our adaptive immune system stores a mixture of IgM and class-switched IgG isotypes of B-cell memory that show different levels of specificity. IgM memory is an earlier product of affinity maturation with higher cross-reactivity and a lower affinity to antigens, reflecting a non-equilibrium memory that can counter evolving pathogens. On the other hand, memory from class-switched IgG isotype is produced during later stages of affinity maturation and is highly specific to the infecting pathogen, reflecting equilibrium memory that is effective against relatively conserved pathogens [16]. Storing a mixture of IgM and class-switched IgG memory is consistent with our recipe for optimal immune strategies to counter pathogens with a broad range of evolutionary rates.

Cross-reactive memory dominates immune response in organisms that encounter fewer pathogens over a shorter lifetime. So far, our analysis has focused on maximizing the net utility of immune response, assuming that organisms encounter many such infections throughout their lifetime. This optimization provides a recipe for optimal immune strategies in response to commonly infecting pathogens. However, the expected frequency of infections is also an important factor that can inform immune strategies. For example, imagine the extreme case that an immune system expects to encounter a pathogen at most only once during an organism's lifetime, e.g. in short-lived organisms. In this case, there is no benefit in keeping a memory even to counter extremely conserved pathogens, for which memory would be otherwise very beneficial.

To study the impact of infection frequency on immune strategies, we use our optimization procedure to maximize the net utility of immune response, while setting a bound on the number of infections throughout an organ-

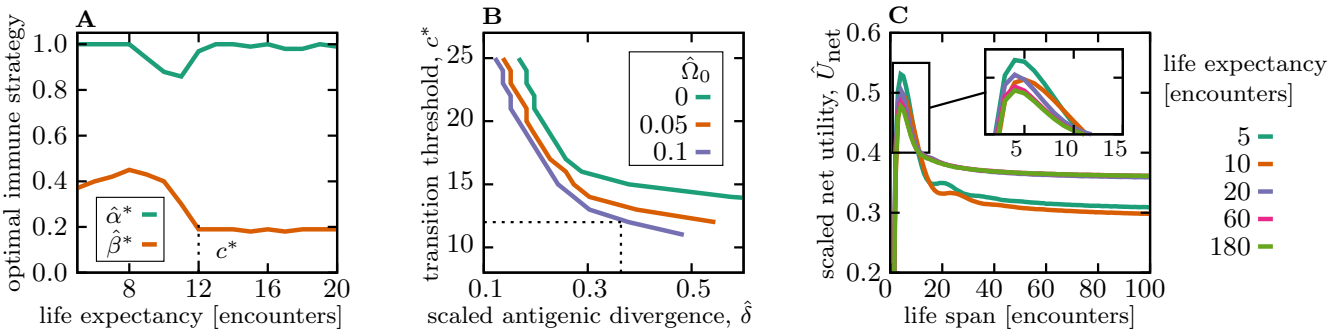

FIG. 5: **Life expectancy influences the specificity of optimal memory. (A)** Memory strategies, i.e., optimal rescaled specificity $\hat{\alpha}^*$ (green) and deliberation factor $\hat{\beta}^*$ (orange) are shown as function of the organism's life expectancy. Memory is highly specific in organisms with unrealistically short lifetimes, which encounter only a few pathogens. Immune strategies show a sharp transition between cross-reactive memory at (realistic) short lifetimes (here at ~ 12 pathogenic encounters) and specific memory at longer lifetime, indicated by the dotted line at the encounter threshold c^*. **(B)** The encounter threshold c^* decays as a function of the antigenic divergence $\hat{\delta}$ and the amplitude of the naive cost $\hat{\Omega}_0$ (colors). **(C)** Scaled net utility $\hat{U}_{\rm net}$ is shown as a function of organism's life span, whose immune strategies ($\hat{\alpha}^*$, $\hat{\beta}^*$) that are optimized for a specified life expectancy (colors as indicated in the legend). Life span and life expectancy are measured in units of the number of pathogenic encounters during lifetime. Simulation parameters: scaled antigenic divergence $\hat{\delta} = 0.35$ (A,C), linear deliberation cost function $\Omega = \Omega_0\hat{\beta}$ with an amplitude $\hat{\Omega}_0 = 0.1$ (A,C), $\alpha_{\rm max} = 4$, $\beta_{\rm max} = 10$, and $\theta = 2$.

ism's lifetime (see SI).

Organisms with an *unrealistically* very short lifetime (measured in units of the number of infections) experience only a few infections, and therefore, a small antigenic drift from the primary infection (lifetime < 8 in units of the number of infections, in Fig 5A). In this case, it would be sufficient for an optimal immune system to generate specific memory ($\hat{\alpha} \approx 1$), which can mount an effective response with only an intermediate deliberation ($\hat{\beta} \sim 0.4$) upon reinfection, even against pathogens with a moderate evolutionary rate (Fig 5A). Organisms with moderately short lifetime experience evolutionary divergence of reinfecting antigens. In this regime, the immune system stores cross-reactive memory (smaller $\hat{\alpha}$) and uses a larger deliberation factor $\hat{\beta}$ such that this lower-affinity and often off-centered memory can mount an effective response to evolved infections (Fig. 5A). Since the organism is relatively short lived, such cross reactive memory could be sufficient throughout the whole lifetime of the organism, without a need for renewal. As the number of encounters passes a threshold c^* (longer lifetime), an optimal immune strategy switches from storing and utilizing cross-reactive memory to generating more specific memory receptors, which would be continually renewed during organisms' lifetime (Fig. 5A). In this regime, the deliberation factor also decreases to facilitate novel responses against antigens that are not readily recognized by memory (Fig 5A). The transition threshold c^* depends on the expected antigenic divergence δ during pathogenic evolution and the details of the immune machinery, and specifically the cost of deliberation $\Omega(\tau)$ due to an elevated level of pathogenic proliferation prior to a novel response (see Fig 5B).

The results in Fig. 5 predict that organisms with few pathogenic encounters or a shorter life-span should generate more cross-reactive and lower affinity (i.e., a naïve-type) memory receptors. Indeed, consistent with our prediction, analysis of immune repertoire data indicates that sequence features of memory and naïve B-cell receptors tend to be more similar to each other in mouse compared to humans that enjoy a longer life expectancy [50]. Nonetheless, more comprehensive data on cross-species comparison of immune strategies is needed to test our predictions.

With the increase in human life expectancy, a pressing question is how well our immune system could cope with a larger number of pathogenic challenges that we are now encountering throughout our lifetimes? Aging has many implications for our immune machinery and the history of infections throughout lifetime leaves a complex mark on immune memory that can have long-lasting consequences [51], which has also been studied through theoretical modeling [28]. In our framework, we can study one aspect of this problem and ask how an immune strategy optimized to battle a given number of infections would perform if the organism were to live longer or equivalently, to encounter pathogens more frequently. Fig. 5C shows that cross-reactive memory generated by an immune system optimized to counter few infections (short life expectancy) becomes highly inefficient (i.e., with a lower net utility $U_{\rm net}$) as the number of encounters increases beyond the organism's expectation (long life span) — an effect that may be in part responsible for the observed decline in the efficacy of our adaptive immunity as we age.

IV. DISCUSSION

Memory is central to our adaptive immunity by providing a robust and preventive response to reinfecting pathogens. In the presence of continually evolving pathogens, immune memory is only beneficial if receptors can recognize evolved antigens by cross-reactivity. As receptors undergo affinity maturation, their structures become more rigid and less cross-reactive, while affinity increases [34, 35]. Consistent with recent experiments [1, 2, 16, 17], we show that memory differentiation should be regulated to preferentially produce lower affinity receptors, which can allow cross-reactive recognition of evolved pathogens. To overcome the resulting energetic impediment of these memory receptors, we infer that the immune system should tune the kinetics of the immune response and allocate a longer deliberation time for memory to react before initiating a novel response — a feature that is also in accordance with observations [9, 10, 32]. Co-optimizing kinetics and energetics of memory ensures an effective response against evolving pathogens, throughout an organism's lifetime.

Optimal cross-reactive immune memory provides a long-term advantage to an organism, yet it may seem energetically sub-optimal over short time scales (Fig. 1). One important consequence of a sub-optimal memory response is known as original antigenic sin, where cross-reactive memory from primary infections could interfere with and suppress a protective novel response [45, 46] — a deficiency that is highly relevant for immunity against rapidly evolving viruses like influenza. The impact of such immune deficiency can even be more pronounced due to changes in an organism's life expectancy. Importantly, we show that immune strategies optimized to benefit short-lived organisms produce highly cross-reactive memory (Fig. 5). If an organism's life-expectancy increases, which is the case for humans, it would be likely for individuals to encounter evolved forms of a pathogen at antigenic distances larger than expected by their immune systems. In this case, cross-reactive memory optimized for a shorter lifetime, could still be activated but with lower efficacy and it could suppress a protective novel response, consistent with original antigenic sin. It is therefore important to consider sub-optimality of immune strategies in the face of extensive elongation of the human lifespan as one of the plausible factors responsible for immune deficiencies brought by aging.

One characteristic of memory B-cells, which is currently missing from our model, is their ability to seed secondary germinal centers and undergo further affinity maturation upon reinfection. Evolvability of memory B-cells can allow cross-reactive memory to further specialize against evolved pathogens, without a need to start a germinal center reaction from an un-mutated naïve receptor. Interestingly, different experiments suggest that the capacity of memory to re-diversify depends on various factors including the memory isotype (IgM vs. class-switch receptors), the type of antigenic target (viruses vs. others) and the extent of memory maturation [38, 39]. Therefore, it is interesting to extend our model to study how evolvability of memory can influence its longterm utility to respond to evolving pathogens, and especially viruses. Understanding the limits of memory re-diversification is instrumental in designing successive vaccination protocols to drive extensive affinity maturation of BCR lineages to elicit broadly neutralizing antibodies, which are the current hopes for universal vaccines against rapidly evolving viruses like HIV.

Acknowledgements

This work has been supported by the DFG grant (SFB1310) for Predictability in Evolution and the MPRG funding through the Max Planck Society. O.S also acknowledges funding from Georg-August University School of Science (GAUSS).

[1] R. Shinnakasu, T. Inoue, K. Kometani, S. Moriyama, Y. Adachi, M. Nakayama, Y. Takahashi, H. Fukuyama, T. Okada, and T. Kurosaki, Nat. Immunol. **17**, 861 (2016).
[2] R. Shinnakasu and T. Kurosaki, Curr. Opin. Immunol. **45**, 126 (2017).
[3] C. A. Janeway, P. Travers, M. Walport, and M. Shlomchik, *Immunobiology: the immune system in health and disease (Garland Science, New York)* (2005).
[4] G. D. Victora and M. C. Nussenzweig, Annu. Rev. Immunol. **30**, 429 (2012).
[5] M. Meyer-Hermann, E. Mohr, N. Pelletier, Y. Zhang, G. D. Victora, and K.-M. Toellner, Cell Rep. **2**, 162 (2012).
[6] S. M. Kaech and E. J. Wherry, Immunity **27**, 393 (2007).
[7] T. Dörner and A. Radbruch, Immunity **27**, 384 (2007).
[8] L. J. McHeyzer-Williams, M. Cool, and M. G. McHeyzer-Williams, The Journal of experimental medicine **191**, 1149 (2000).
[9] S. G. Tangye, D. T. Avery, E. K. Deenick, and P. D. Hodgkin, J. Immunol. **170**, 686 (2003).
[10] S. G. Tangye and P. D. Hodgkin, Immunology **112**, 509 (2004).
[11] L. Moens, A. Kane, and S. G. Tangye, Immunol. Cell Biol. **94**, 774 (2016).
[12] C. C. Goodnow, C. G. Vinuesa, K. L. Randall, F. Mackay, and R. Brink, Nat. Immunol. **11**, 681 (2010).
[13] E. J. Blink, A. Light, A. Kallies, S. L. Nutt, P. D. Hodgkin, and D. M. Tarlinton, The Journal of experimental medicine **201**, 545 (2005).
[14] K. G. Smith, A. Light, G. J. Nossal, and D. M. Tarlinton, EMBO J. **16**, 2996 (1997).
[15] D. Paus, T. G. Phan, T. D. Chan, S. Gardam, A. Basten, and R. Brink, The Journal of experimental medicine **203**, 1081 (2006).
[16] F. J. Weisel, G. V. Zuccarino-Catania, M. Chikina, and M. J. Shlomchik, Immunity **44**, 116 (2016).

[17] T. Recaldin and D. J. Fear, Clin. Exp. Immunol. **183**, 65 (2016).

[18] R. Antia, S. S. Pilyugin, and R. Ahmed, Proc. Natl. Acad. Sci. U.S.A. **95**, 14926 (1998).

[19] A. Mayer, T. Mora, O. Rivoire, and A. M. Walczak, Proc. Natl. Acad. Sci. U.S.A. **113**, 8630 (2016).

[20] A. Perelson and G. Weisbuch, Rev. Mod. Phys. **69**, 1219 (1997).

[21] G. Altan-Bonnet, T. Mora, and A. M. Walczak, Physics Reports **849**, 1 (2020).

[22] B. Laffleur, S. M. Bardet, A. Garot, M. Brousse, A. Baylet, and M. Cogné, Genes Immun. **15**, 341 (2014).

[23] E. D. Hawkins, M. L. Turner, M. R. Dowling, C. van Gend, and P. D. Hodgkin, Proc. Natl. Acad. Sci. U.S.A. **104**, 5032 (2007).

[24] W. S. Bialek, *Biophysics : searching for principles* (Princeton, NJ : Princeton University Press, 2012).

[25] A. Nourmohammad, T. Held, and M. Lässig, Curr Opin Genet Dev **23**, 684 (2013).

[26] A. S. Perelson and G. F. Oster, J. Theor. Biol. (1979).

[27] A. Mayer, V. Balasubramanian, T. Mora, and A. M. Walczak, Proc. Natl. Acad. Sci. U.S.A. **112**, 5950 (2015).

[28] A. Mayer, V. Balasubramanian, A. M. Walczak, and T. Mora, Proc. Natl. Acad. Sci. U.S.A. **116**, 8815 (2019).

[29] S. Bradde, A. Nourmohammad, S. Goyal, and V. Balasubramanian, Proc. Natl. Acad. Sci. U.S.A. **15**, 201903666 (2020).

[30] J. K. Whitmire, B. Eam, and J. L. Whitton, PLoS Pathog. **4**, e1000041 (2008).

[31] M. D. Martin, S. A. Condotta, J. T. Harty, and V. P. Badovinac, J. Immunol. **188**, 1255 (2012).

[32] G. Blanchard-Rohner, A. S. Pulickal, C. M. Jol-van der Zijde, M. D. Snape, and A. J. Pollard, Blood **114**, 4998 (2009).

[33] D. C. Macallan, D. L. Wallace, Y. Zhang, H. Ghattas, B. Asquith, C. de Lara, A. Worth, G. Panayiotakopoulos, G. E. Griffin, D. F. Tough, et al., Blood **105**, 3633 (2005).

[34] G. J. Wedemayer, P. A. Patten, L. H. Wang, P. G. Schultz, and R. C. Stevens, Science **276**, 1665 (1997).

[35] S. A. Frank, *Immunology and Evolution of Infectious Disease* (Princeton University Press, Princeton (NJ), 2002).

[36] J. R. Mascola and B. F. Haynes, Immunol. Rev. (2013).

[37] N. C. Wu, G. Grande, H. L. Turner, A. B. Ward, J. Xie, R. A. Lerner, and I. A. Wilson, Nat Commun **8**, 15371 (2017).

[38] M. J. Shlomchik, Cold Spring Harb Perspect Biol **10**, a029405 (2018).

[39] L. J. McHeyzer-Williams, C. Dufaud, and M. G. McHeyzer-Williams, Cold Spring Harb Perspect Biol **10**, a028878 (2018).

[40] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton University Press, Princeton, 1944).

[41] P. A. Ortega and D. A. Braun, Proc. R. Soc. A **469**, 20120683 (2013).

[42] E. T. Jaynes, Physical Review **106**, 620 (1957).

[43] C. H. LaMont and P. A. Wiggins, Phys Rev E **99**, 052140 (2019).

[44] J. Grau-Moya, M. Krüger, and D. A. Braun, Entropy 2018, Vol. 20, Page 1 **20**, 1 (2018).

[45] T. Francis JR. and 1960, Proc. Am. Philos. Soc. **104**, 572 (1960).

[46] A. Vatti, D. M. Monsalve, Y. Pacheco, C. Chang, J.-M. Anaya, and M. E. Gershwin, J. Autoimmun. **83**, 12 (2017).

[47] K. L. Good-Jacobson and M. J. Shlomchik, J. Immunol. **185**, 3117 (2010).

[48] K. Kometani, R. Nakagawa, R. Shinnakasu, T. Kaji, A. Rybouchkin, S. Moriyama, K. Furukawa, H. Koseki, T. Takemori, and T. Kurosaki, Immunity **39**, 136 (2013).

[49] M. J. Shlomchik, W. Luo, and F. Weisel, Immunol. Rev. **288**, 49 (2019).

[50] Z. Sethna, Y. Elhanati, C. S. Dudgeon, C. G. Callan, A. J. Levine, T. Mora, and A. M. Walczak, Proc. Natl. Acad. Sci. U.S.A. **114**, 2253 (2017).

[51] P. Saule, J. Trauet, V. Dutriez, V. Lekeux, J.-P. Dessaint, and M. Labalette, Mech. Ageing Dev. **127**, 274 (2006).

Supplementary Information

i. Numerical optimization

Numerical optimization is performed on ensembles of immune systems that encounter evolving pathogens. Recognition of an evolved pathogen at the ith round of infection υ_i by a memory that was stored in response to a primary infection υ_0 (0th round) depends on the antigenic distance $d_i = \|\upsilon_i - \upsilon_0\|$. We model pathogenic evolution as diffusion in the antigenic shape space. In this model, the expected antigenic distance between the primary infection υ_0 and the evolved antigen υ_i can be characterized as, $\langle d_i^2 \rangle \equiv \langle \|\upsilon_i - \upsilon_0\|^2 \rangle = \zeta^2 (t_i - t_0) = i\,\delta^2$, where ζ is the diffusion coefficient (i.e., the evolutionary rate) and δ is the (averaged) antigenic divergence per round of infection. Importantly, this relationship does not depend on the dimensionality of the antigenic shape space, which in general, is difficult to characterize. We simulate pathogenic evolution relative to a primary infection by drawing the corresponding antigenic distance d_i of the ith round of infection from a normal distribution with mean $\delta\sqrt{i}$ and standard deviation $0.05\delta\sqrt{i}$. The width of this normal distribution characterizes the fluctuations in the mean divergence between infections and reflects how the evolutionary trajectory of a pathogen samples the multi-dimensional shape space surrounding the antigen from the primary infection. Nonetheless, our results are insensitive to the exact choice of this width.

To characterize optimal specificity α^* and deliberation factor β^* (Figs. 2,3,5), we simulate ensembles of immune systems with different immune strategies (α, β), chosen uniformly from the range $\alpha \in [0, \alpha_{max}]$ and $\beta \in [0, \beta_{max}]$, with 500 increments in both parameters. Each immune system experiences successive rounds of infection with an evolving pathogen with a given antigenic divergence δ. During each encounter, the immune system chooses between utilizing an existing memory or initiating a novel response according to eq. 2. The net utility of each encounter is calculated according to eq. 6. We estimate the expected net utility per encounter over a lifetime of 60 total encounters and repeat this experiment across 10^5 independent ensembles to find the optimal immune strategies (α^*, β^*) with the highest net utility. As shown in Fig. 5, simulating up to 60 encounters is sufficient for inference of optimal strategies in the asymptotic regime (i.e., a long lifetime).

To characterize optimal immune strategies against mixture of pathogens with distinct levels of antigenic divergences, we define the mixture immune strategy by a set of specificities $\vec{\alpha} = \{\alpha_i\} = (\text{with, } i = 1, \ldots, N_m)$, where each α_i is a degree of specificity that a stored memory receptor can potentially have, and N_m is the number of possible specificity strategies that an immune system can choose from. The probability that an immune system with the mixture strategy $\vec{\alpha}$ recognizes a pathogen υ through a memory response follows from an extension of eq. 1,

$$\begin{aligned} P^{(m)}_{\text{recog.}}(\vec{\alpha}, \upsilon) &= 1 - \prod_{\text{specificity: } \alpha_i} \left(1 - P^{(m)}_{\text{recog.}}(r_m^{\alpha_i}, \upsilon)\right) \\ &= 1 - \prod_{\text{specificity: } \alpha_i} e^{-E_\theta(r_m^{\alpha_i}, \upsilon)\Gamma(\tau)} = 1 - e^{-\sum_{\alpha_i} E_\theta(r_m^{\alpha_i}, \upsilon)\Gamma(\tau)} \equiv 1 - e^{-\tilde{\beta}\,\overline{E}_\theta(\upsilon)} \end{aligned} \tag{S1}$$

where $\overline{E}_\theta(\upsilon) = \frac{1}{N_m} \sum_{r_m^\alpha} E_\theta(r_m^{\alpha_i}, \upsilon)$ is the expected affinity of memory (with distinct specificities) against antigen υ in an immune repertoire and $\tilde{\beta} \equiv N_m \beta$ is an effective deliberation factor for all choices of specificity. It should be noted that this effective deliberation factor $\tilde{\beta}$ is an extensive quantity with respect to the number of specificity strategies that an immune system can choose from, and therefore, is comparable across immune systems with different numbers of strategies.

We set out to characterize the mixture strategy as the probability $P_\beta(\alpha)$ based on which an immune system with a given effective deliberation factor $\tilde{\beta}$ should store a memory receptor with specificity α, in order to optimally counter infecting pathogens with distinct antigenic divergences, drawn from a distribution $P(\delta)$. We start our optimization by defining a uniform mixture strategy, where the elements of the immune specificity vector $\vec{\alpha} = \{\alpha_i\}$ (of size $N_m = 20$), are drawn uniformly from the range $[0, \alpha_{\max}]$. Each optimization step aims to improve the specificity vector $\vec{\alpha}$ to maximize the net utility (per encounter) of the mixture immune response $U_{\text{net}}(\vec{\alpha}^k)$ against 1000 independently evolving antigens whose (scaled) antigenic divergences are drawn uniformly from the range $\hat{\delta} = [0, \hat{\delta}_{\max}]$. We use stochastic simulations to estimate the net utility of the mixture strategy $U_{\text{net}}(\vec{\alpha}^k)$, whereby the relative affinity of memory receptors (with varying specificities), $E_\theta(r_m^{\alpha_i}, \upsilon)/\overline{E}_\theta(\upsilon)$, determines the stochastic rate of their response to the infecting antigen υ. The net utility (per encounter) of the immune response against each of the 1000 independently evolving antigens is estimated by averaging over a host's lifetime with 200 rounds of pathogenic encounters. We update the mixture strategy over 3000 steps, using local gradient ascent by sampling 100 points in the space of specificity vectors at each step to maximize net utility,

$$\vec{\alpha}^{k+1} = \vec{\alpha}^k + \epsilon \nabla U_{\text{net}}(\vec{\alpha}^k) \tag{S2}$$

Here, k indicates the optimization step and $\epsilon = 0.1$ is a hyper-parameter for gradient ascent. We repeat the optimization process starting from 200 independently drawn initial uniform mixture strategies $\vec{\alpha}^0$ to characterize the ensemble of optimal memory strategies $P_\beta(\alpha)$ against pathogens with distinct antigenic divergences drawn uniformly from a given range $\hat{\delta} = [0, \hat{\delta}_{\max}]$, as shown in Fig. 4. We also characterize the probability that a stored memory with a given specificity is utilized against future infections (solid line in Fig. 4). To do so, we test the optimized ensemble of specificities $P_\beta(\alpha)$ against 5000 independent pathogens with antigenic divergences drawn uniformly from the range $\hat{\delta} = [0, \hat{\delta}_{\max}]$. We evaluate the usage of a memory with a given specificity α (solid line in Fig. 4) as the conditional probability $P_\beta(\text{use } \alpha|\text{produce } \alpha)$ for using that memory given that it is produced (i.e., drawn from the distribution $P_\beta(\alpha)$).

ii. Decision-making for immune response

In the theory of decision-making, a rational decision-maker chooses between two possible actions $a \in \{\text{naïve}, \text{memory}\}$ each contributing a utility U_a. If the decision-maker has different prior preference for each action, which we denote by the prior probability distribution $Q_0(a)$, its decisions could be swayed by this knowledge. As a result, the constrained decision-maker should choose actions according to an optimized probability density $Q(a)$, which maximizes the expected utility while satisfying constraints due to the prior assumption [40, 41],

$$Q(a) = \underset{Q(a)}{\operatorname{argmax}} \left(\sum_a U_a Q(a) - \frac{1}{\beta} D_{KL} \left(Q(a) || Q_0(a) \right) \right) \tag{S3}$$

Here, $D_{KL}(Q(a)||Q_0(a)) = \sum_a Q(a) \log \left(Q(a)/Q_0(a) \right)$ is the Kullback-Leibler distance between the rational distribution $Q(a)$ and the prior distribution $Q_0(a)$ and $1/\beta$ is a Lagrange multiplier that constrains the efficacy of a decision-maker to process new information and deviate from its prior assumption. The optimal solution for a rational yet constrained decision follows,

$$Q(a) = \frac{1}{Z} Q_0(a) e^{\beta U_a} \tag{S4}$$

where $Z = \sum_a Q_0(a) e^{\beta U_a}$ is a normalization factor. If information processing is highly efficient (i.e., the bias factor $1/\beta \to 0$) the rational decision-maker deterministically chooses the action with the highest utility. On the other hand, if the prior is strong (i.e., $1/\beta \to \infty$), the decision-maker hardly changes its opinion and acts according to its prior belief (i.e., $Q(a) = Q_0(a)$). Moreover, if the prior distribution is uniform across actions (i.e., no prior preference), rational decision maximizes the entropy of the system [42], resulting in the probability of actions $Q(a) \sim \exp[\beta U_a]$. In our analysis we consider the case of unbiased maximum entropy solution for decision-making. As a result the probability to utilize memory $Q_{\text{mem.}}$ or naïve $Q_{\text{naïve}}$ follows,

$$Q_{\text{mem.}} = 1 - Q_{\text{naïve}} = \frac{e^{\beta U_{\text{mem}}}}{e^{\beta U_{\text{mem}}} + e^{\beta U_{\text{naïve}}}} \tag{S5}$$

which is a sigmoidal function, dependent on the utility of each action.

We can draw an analogy between the two description of memory response that are based on (i) recognition of a pathogen by considering the energetics and kinetics of the immune response (eq. 1) and (ii) the rational decision-making by considering the utility and the efficacy of decision-making (eq. S5). By equating these two descriptions we have,

$$\begin{aligned} Q_{\text{mem.}} = P^{(m)}_{\text{recog.}}(r_m, \upsilon) \;&\longrightarrow\; \frac{e^{\beta U_{\text{mem}}}}{e^{\beta U_{\text{mem}}} + e^{\beta U_{\text{naïve}}}} = 1 - e^{-E_{\alpha,\theta}(r_m,\upsilon)\Gamma(\upsilon,\tau)} \\ &\longrightarrow\; \beta(U_{\text{mem.}} - U_{\text{naïve}}) = \log\left[e^{E_{\alpha,\theta}(r_m,\upsilon)\Gamma(\upsilon,\tau)} - 1 \right] \end{aligned} \tag{S6}$$

In the regime that memory is efficient and is being utilized to mount a response (i.e., a low chance for naïve recognition: $P^{(0)}_{\text{recog.}} = e^{-E(\upsilon)\Gamma(\upsilon,\tau)} \ll 1$), the sigmoid form for decision to use memory (eq. S5) is dominated by an exponential factor, resulting in a simple analogy: $\beta \simeq \Gamma(\upsilon, \tau)$, and $U_{\text{mem.}} - U_{\text{naïve}} \simeq E_{\alpha,\theta}(r_m, \upsilon)$.

Previous work has drawn a correspondence between the inverse temperature in thermodynamics and the effect of sample size on statistical inference [43]. Interestingly, we see a similar analogy between the efficacy of information

processing in decision-making (i.e., deliberation factor) β, which plays the role of the inverse temperature in thermodynamics, and the total accumulated pathogenic load $\Gamma(\upsilon, \tau)$, which acts as the sample size for memory receptors as they encounter and accumulate information about pathogens.

It should be noted that our decision-making formalism assumes that if memory is available, it can be utilized much more efficiently and robustly than a naïve response. Therefore, we do not consider scenarios where memory and naïve responses are equally involved in countering an infection— a possibility that could play a role in real immune responses. Nonetheless, since such mixed responses are relatively rare, we expect that including them in our model would only result in a slightly different interpretation of the deliberation factor β and should not qualitatively impact our results.

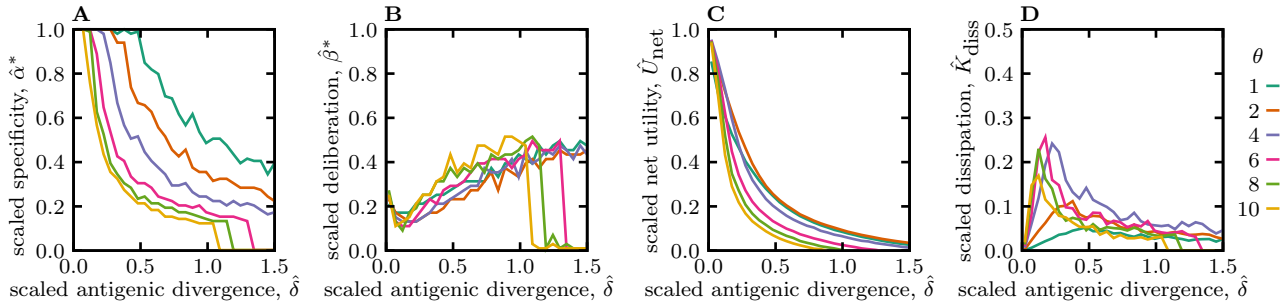

FIG. S1: **Optimal memory strategies for different specificity shape factors θ. (A)** Scaled specificity $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\rm max}$, **(B)** scaled deliberation factor $\hat{\beta}^* \equiv \beta^*/\beta_{\rm max}$, **(C)** scaled net utility $\hat{U}_{\rm net} \equiv U_{\rm net}/E_{\rm max}$, and **(D)** scaled dissipation are shown as a function of the scaled antigenic divergence per infection $\hat{\delta} = \delta/(\alpha_{\rm max}^{-1})$ (similar to Fig. 2). Colors indicate different shape factors θ of the specificity function, ranging from a double-sided exponential ($\theta = 1$), to Gaussian for $\theta = 2$ (as in Fig. 2), and top-hat functions $\theta > 2$. The dependence of memory strategies on antigenic divergence is qualitatively insensitive to the shape factor of the specificity function. Simulation parameters: linear deliberation cost function $\Omega = \hat{\Omega}_0\hat{\beta}$ with $\hat{\Omega}_0 = 0.1$, $\alpha_{\rm max} = 4$, and $\beta_{\rm max} = 10$.

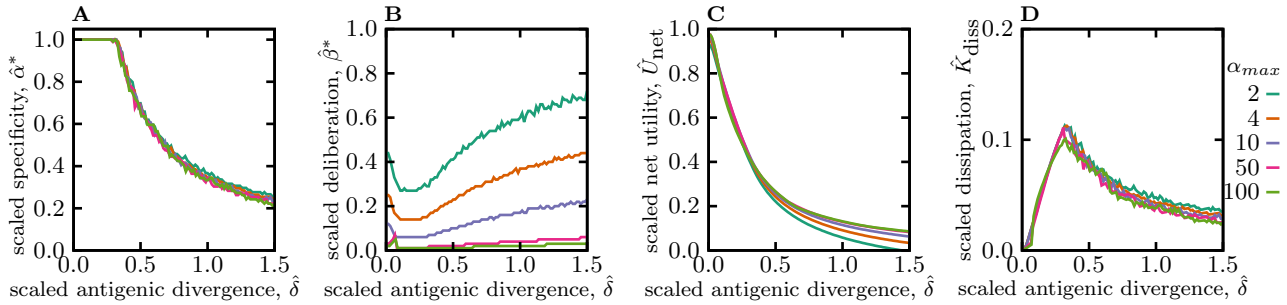

FIG. S2: **Optimal memory strategies for different specificity thresholds $\alpha_{\rm max}$. (A)** Scaled specificity $\hat{\alpha}^* \equiv \alpha^*/\alpha_{\rm max}$, **(B)** scaled deliberation factor $\hat{\beta}^* \equiv \beta^*/\beta_{\rm max}$, **(C)** scaled net utility $\hat{U}_{\rm net} \equiv U_{\rm net}/E_{\rm max}$, and **(D)** scaled dissipation are shown as a function of the scaled antigenic divergence per infection $\hat{\delta} = \delta/(\alpha_{\rm max}^{-1})$ (similar to Fig. 2). Colors indicate different specificity thresholds $\alpha_{\rm max}$. Memory strategies are qualitatively insensitive to the specificity threshold. Simulation parameters: linear deliberation cost function $\Omega = \hat{\Omega}_0\hat{\beta}$ with $\hat{\Omega}_0 = 0.1$ and $\beta_{\rm max} = 10$.

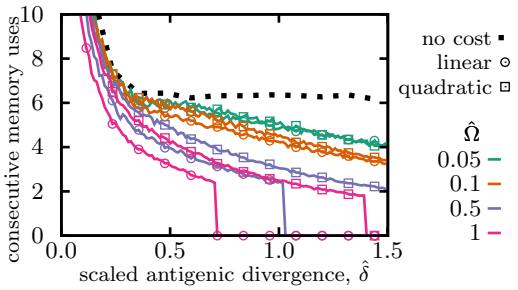

FIG. S3: **Rounds of memory usage depends on antigenic divergence.** The expected number of rounds that a memory receptor is utilized prior to a novel response in an optimal system is shown to decay as the antigenic divergence $\hat{\delta}$ increases. The results are evaluated for immune systems with optimized strategies ($\hat{\alpha}^*$, $\hat{\beta}^*$) against pathogens with a given scaled antigenic divergence $\hat{\delta}$; the corresponding strategies are shown in Fig. 2. Colors / markers indicate different naïve cost functions for deliberation, including no-cost $\hat{\Omega} \equiv \Omega/E_{\max} = 0$, linear cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}$, and quadratic cost $\hat{\Omega} = \hat{\Omega}_0\hat{\beta}^2$, with varying amplitudes $\hat{\Omega}_0$. Simulation parameters: $\alpha_{\max} = 4$, $\beta_{\max} = 10$, and $\theta = 2$.

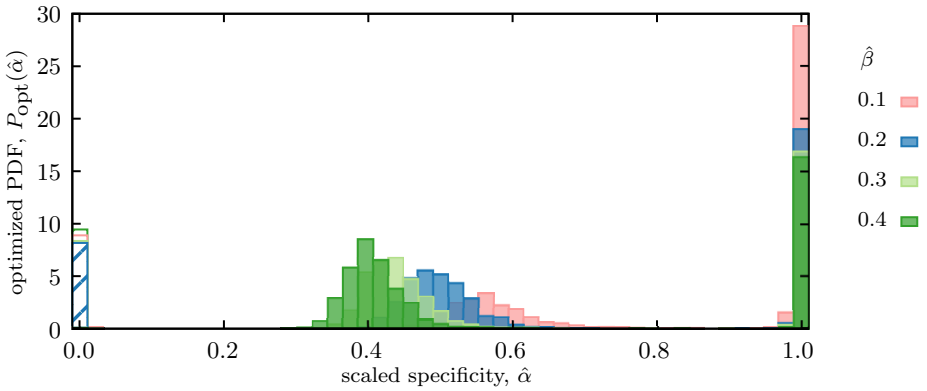

FIG. S4: **Mixed memory strategy against pathogens for different deliberation factors $\hat{\beta}$.** Distribution of scaled optimized specificities $\hat{\alpha}^*$ of functional memories is shown for an immune system with a fixed deliberation factor $\hat{\beta} = 0.2$, in which a mixture strategy with a bimodal distribution of specificities $P(\hat{\alpha})$ is established to counter pathogens with a broad range of antigenic divergences, drawn uniformly from a range $\hat{\delta} \in (0, 1.6)$ (similar to Fig 4). The dashed bars indicate stored memory with specificity $\alpha = 0$, which is not further used in response to infections. Colors indicate different deliberation factors. Simulation parameters: $\alpha_{\max} = 4$, and $\beta_{\max} = 10$.