# Tuning environmental timescales to evolve and maintain generalists
**SI**

Vedant Sachdeva, Kabir Husain, Jiming Sheng, Shenshen Wang, Arvind Murugan

## I. ENTROPICALLY DISFAVORED GENERALISTS

### A. Model

To construct landscapes with entropically disfavored generalists, we model antibodies and antigens in a manner similar to the one described by by Wang et. al[1]. In this model, the sequence of each antibody, $\mathbf{x}$ is a sequence of length $L$ with entries $\pm 1$. Each antigen, indexed by $\eta$, is assumed to have an epitope of sequence, $\mathbf{h}^\eta$. Each epitope is length $L$ with entries $\pm 1$. The binding energy of a given antigen to an antibody is given by an additive sum-over-sites model:

$$E\left(\mathbf{x}, \mathbf{h}^\eta\right) = -\frac{1}{L}\sum_i^L h_i^\eta x_i \tag{1}$$

The fitness of each antibody $\mathbf{x}$ in the presence with antigen $\eta$ is given by thresholding its binding affinity, as follows:

$$F^{(\eta)}\left(\mathbf{x}\right) = s\epsilon\Theta\left(-\left(E(\mathbf{x}, \mathbf{h}^\eta) + \frac{T}{L}\right)\right) - s \tag{2}$$

Here, $\Theta(x)$ is the Heaviside function. $\frac{T}{L}$ is the binding energy threshold that an antibody must overcome before reaping a fitness benefit. The binding energy threshold, $\frac{T}{L}$ can be translated into minimum number of binding interactions needed to reap a fitness benefit, $T_{\text{sites}}$, by $T_{\text{sites}} = (L+T)/2$. By construction, all fit individuals have the same fitness $s(\epsilon - 1) > 0$ and all unfit individuals are equally unfit to an extent $-s$, resulting in a degeneracy of fit and unfit genotypes.

### B. Specialists and Generalists

Here, we demonstrate the fraction of antibodies that are generalists for two antigens. Biological constraints impose that parts of an antigen's epitope is conserved, while other parts are variable as the viral strain evolves. As such, we suppose $L_c$ sites of the $L$ total sites are fixed across $\eta$, while the others are unique to each antigen. This constraint results in the possibility that some antibody sequences have a positive fitness for all antigens. Such antibodies are called generalists. The number of generalists is a function of the threshold $T$ and the lengths of the conserved $L_c$ and variable $L_v = L - L_c$ regions. In particular, if $T_{\text{sites}} > L_c + \frac{1}{2}L_v$, there are no generalists. A simple equation to compute the fraction of antibodies with positive fitness for an antigen that are generalists is given as follows:

$$\frac{\Omega_g}{\Omega} = \frac{\sum_{j=0}^{L}\binom{L_c}{j}\binom{L_v}{T_{\text{sites}}-j}}{\sum_{k=T_{\text{sites}}}^{L}\binom{L}{k}} \tag{3}$$

where, $\binom{N}{m}$ is the combinatorial function, $\Omega_g$ is the number of generalists, and $\Omega$ is the number of antibodies with positive fitness. By rule, this function is zero if $m > N$ or $m < 0$. Here, $j$ indices the number of sites matched in the conserved portion of the antibody and $k$ indices the number of overall sites matched along the string.

Here, we considered $L = 19$, $L_c = 12$, and $T = 11$. The proportion of antibodies with positive fitness that are generalists for these parameter choices is $\approx 1.3\%$. This choice is qualitatively similar to the analysis developed in [1] based on experiments there. The analysis can be repeated for longer sequences and the results are qualitatively unchanged; the primary effect of changing L is explained by the change in entropy, as predicted by SI Equation 3.

www.pnas.org/cgi/doi/10.1073/pnas.1914586117

### C. Finite population simulation

Affinity maturation is an evolutionary process for antibodies with complex population dynamics[2]. Here, we first model this process using a simplified canonical birth-death-mutation model - a 'Yule' process [3] - commonly used to study evolutionary dynamics. The 'Yule' process ignores many of the molecular details of affinity maturation while still enabling us to develop a minimal model of evolutionary dynamics in time-varying environments. We then verify our results with an independent simulation that accounts for population dynamics and complexities inherent to affinity maturation.

#### *1. Yule Process*

The key ingredients in a Yule process are:

- **Mutation** with rate $\mu$ per individual, in which a single site on the genome is mutated (i.e. a single bit-flip of $\mathbf{x}$).
- **Birth-death** with rate $\lambda$ per individual.
- The population size is maintained at a carrying capacity $K$ by modulating the probability of replication by a factor $(1 - N/K)$. With the particular fitness function of this model, we have $\Pr(\mathbf{x} \text{ reproduces}) = \Theta\left(E(\mathbf{x}, h^\eta) - T\right)(1 - N/K)$. If the individual does not reproduce, it is removed from the population. Here, $N$ is the current population size, $K = 500$ is a carrying capacity that prevents the population from growing indefinitely, and $F^{(\eta)}$ is the fitness of that individual in the (current) landscape $\eta$.

At each event, time is advanced by the usual exponentially distributed amount. The environment $\eta$ is taken to alternate between $\eta = 1$ and $\eta = 2$ every $\tau_{\text{epoch}}$.

Choosing units of time by setting the birth-death rate $\lambda = 1$, we set the mutation rate to $\mu = 0.05$. Finally, we set the carrying capacity $K = 500$. We can infer from these choices that $s = 1, \epsilon = 2 - \frac{2N}{K}$.

We evolved an initially monoclonal population of size $N = 10$ (initialised with $\mathbf{x} = \mathbf{h}^\eta$ for all individuals). Simulations were run for either a fixed time $t = 100$ (in units of $\lambda$) or for $t = 10\tau_{\text{epoch}}$, whichever is longer, and we performed 25 replicates for each value of $\tau_{\text{epoch}}$. For each run, we saved the number of generalists and the overall population size at the end of the simulation. In Fig. 2b of the main text, we plotted the proportion of trials that had more than 10 generalist antibodies at the end of the simulation, finding that the proportion was high for an intermediate rate of cycling.

#### *2. Affinity Maturation Inspired Model*

To more directly model affinity maturation, we also simulate a model that mirrors the known dynamics of B-cells in germinal centers [1, 4, 5]. The steps are as follows:

- B-cell clones expand without significant mutation in the first week after vaccination. We model such a formation of germinal centers by taking a B-cell with an antibody that meets the binding affinity threshold for one antigen and replicate it to a size of 1500 B-cells.
- We model the reproduction and somatic hypermutation phase of affinity maturation in the dark zone of the germinal center by allowing each B-cell to duplicate twice with a mutation rate of 0.00625 per replication per base pair.
- We then model the selection phase in the light zone by determining if each B-cell in our B-cell population can internalize antigens it encounters on a follicular dendritic cell (FDC). We say that a B-cell can internalize an antigen if its antibody's binding affinity for that antigen, as given by Equation 1 meets a threshold, $T$.
- B-cells receive T-cell help to avoid an apoptosis signal as a function of whether or not they internalized antigen. As in earlier work[1, 4, 5], we assume the probability B-cells do not receive help increases with binding energy is proportional to $\exp(\alpha(F^{(\eta)}(\mathbf{x}) - F_{\text{threshold}}))$.
- The surviving B-cells are recycled into the dark zone. We repeat the steps above until the B-cell population grows to be larger than 2000 or the process has cycled 100 times. These choices model antigen depletion on the follicular dendritic cells (FDCs) in the germinal center and antigen decay.

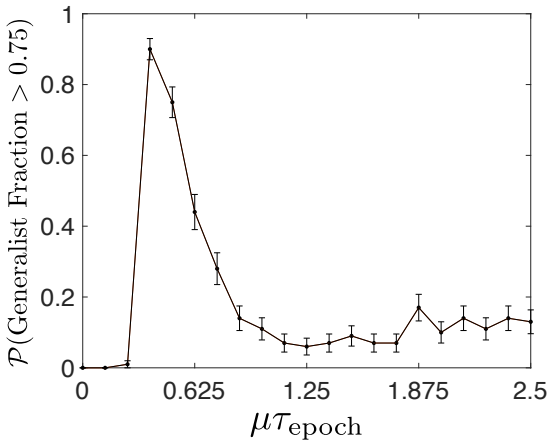


FIG. 1. Evolving generalists using a detailed model of affinity maturation. We simulated cycling antigens in a model with known details of the population dynamics of B-cells in germinal centers. Here, we plot the probability that the population at the end of affinity maturation has at least 75% of its population in a generalist genotype. We observe a resonant peak in this probability, similar to results presented in the main paper for the simpler population dynamics model based on a Yule process.

We present the results of this simulation in Fig. 1.

We find that the results from using this affinity maturation-specific evolutionary scheme are qualitatively similar to the minimal Yule process model. Note that this affinity maturation model incorporates numerous ingredients, particular to affinity maturation, that are not captured by the Yule process used in the earlier section. E.g., the specifics of birth and death, carrying capacity and details of how the affinity maturation process terminates differ. And yet we obtain qualitatively similar results, showing that our results are primarily tied to the broad topology of the fitnesss landscape and the ratio of broadly relevant timescales and not to particular details of evolutionary population dynamics.

### D. Evolution between Specialists and Generalists: $\chi_{\mathbf{s}\to\mathbf{g}}$ and $\chi_{\mathbf{g}\to\mathbf{g}}$

As discussed in the main text, there are two possible failure modes: (1) at cycling rates too fast, the population does not have the time to evolve a generalist before adverse selection result in population extinction, (2) at cycling rates too slow, the population loses its ability to maintain generalists. In order to illustrate the tension between cycling too fast and cycling too slow, we compute two quantities:

- $\chi_{\text{g}\to\text{g}}$, the fraction of trials starting from an initially generalist population maintaining at least 20% generalists after one epoch.
- $\chi_{\text{s}\to\text{g}}$, the fraction of trials in which a monoclonal population of specialists, initialized at $\mathbf{x} = \mathbf{h}^{(1)}$, evolve a single generalist within an epoch.

Both $\chi_{\text{g}\to\text{g}}$, $\chi_{\text{s}\to\text{g}}$ are computed from 50 replicates for each $\tau_{\text{epoch}}$. As plotted in Fig. 2c, we observe that $\chi_{\text{s}\to\text{g}}$ starts initially at 0 and rises with $\tau_{\text{epoch}}$, while $\chi_{\text{g}\to\text{g}}$ starts at 1 and falls with $\tau_{\text{epoch}}$.

**Population traces:** Fig. 2d shows population traces in single runs. As in Section I C, we initialized a population at $\mathbf{x} = \mathbf{h}^{(1)}$ with parameters as above. Generalist fraction is defined as $\frac{\text{number of generalists}}{\text{population size}}$.

We use the following values of $\tau_{\text{epoch}}$ for Fig.2(d): Fast cycling $\tau_{\text{epoch}} = 1$ (Fig. 2d(i)), Intermediate cycling $\tau_{\text{epoch}} = 60$ (Fig.2d(ii)), Slow cycling $\tau_{\text{epoch}} = 400$ (Fig.2d(iii))

### E. Timescale Analysis

Our numerical study found that an intermediate timescale of environmental cycling, $\tau_{\min} < \tau_{\text{epoch}} < \tau_{\max}$, was most effective at obtaining generalists. Here we estimate the bounds, $\tau_{\min}$and $\tau_{\max}$, in terms of the mutation rate $\mu$, the population size $N$, the length of the genotype $L$, and the distance from the initial ancestral genotype to the generalist, $d_{i\to g}$.

*Finding the generalist: $\tau_{min}$*

Consider a population initialised as a specialist for antigen 1. For sufficiently strong selection pressure ($s > \mu \log N$), purifying selection drives the population to extinction if a generalist has not been discovered before the environment switches to antigen 2. Thus we demand that $\tau_{\text{epoch}}$ is long enough for the population to evolve a generalist in a single epoch.

As fitness is uniform across the specialist region, the population must discover the generalist by diffusion. An initially monoclonal population of size $N$ diffuses out from the initial genotype. If the population size is much smaller than the number of possible genotypes ($N \ll 2^L$), there are two possible regimes of the diffusive search:

1. The initial genotype is far from the generalist: more precisely, the population size $N$ is smaller than the set of sequences between the initial genotype and the generalist. In terms of the Hamming distance between the initial genotype to the generalist, $d_{i\to g}$: $\sum_{d=0}^{d_{i\to g}} \binom{L}{d} > N$

   In this regime, finding the generalist is a rare event, requiring time $\mu\tau_{\min} \sim 2^L$, i.e. the time taken to explore all of genotype space. It is therefore extremely improbable that the generalist will be found, and population extinction is likely.

2. The initial genotype is close to the generalist: that is, the generalist is sufficiently close to the initial condition that it may be found by the diffusing population of antibodies:$\sum_{d=0}^{d_{i\to g}} \binom{L}{d} < N$

   For $L = 19$ and $N = 500$, as used in the simulation, this suggests that for $d_{i\to g} \leq 4$ the generalist may be reasonably found by diffusion.

Assuming that we are in the latter regime, we may estimate $\tau_{\min}$ from $\langle d(t)\rangle$, the average distance away from the initial condition that an individual has diffused in time $t$, by solving:

$$\langle d(\tau_{\min})\rangle \sim d_{i\to g} \tag{4}$$

We compute $\langle d(t)\rangle$ as follows: the probability that a diffusing individual may be found at (Hamming) distance $d$ from its initial genotype is:

$$P(d,t) = \binom{L}{d} \mathrm{e}^{-\mu t} \sinh^d\left(\frac{\mu t}{L}\right) \cosh^{L-d}\left(\frac{\mu t}{L}\right) \tag{5}$$

Thus, $\langle d(t)\rangle = \sum_{d=0}^{L} d\, P(d,t) = L\mathrm{e}^{-\frac{\mu t}{L}} \sinh\left(\frac{\mu t}{L}\right)$. Inserting into Eq. 4 and solving for $\tau_{\min}$:

$$\tau_{\min} = \frac{L}{2\mu} \log\left(\frac{L}{L - 2d_{i\to g}}\right) \approx \frac{d_{i\to g}}{\mu} \tag{6}$$

where the approximation is valid for $d_{i\to g}/L \ll 1$. For values used in the simulation ($L = 19$, $d_{i\to g} = 3$), we obtain $\mu\tau_{\min} \approx 3.5$, which is consistent with our numerical results.

*Maintaining a generalist: $\tau_{max}$*

We may similarly estimate the upper bound for effective cycling, $\tau_{\max}$. Supposing that an evolving population has found the generalist, it must now remain localised there. For this to happen, the environment must switch rapidly enough to prune, by purifying selection, those individuals that diffuse away from the generalist.

Consider a monoclonal population of size $N$ at a generalist at time $t = 0$. For simplicity, let us suppose that the generalist has to accrue $d_{g\to s}$ mutations to become a specialist. The number of generalists in sequence space, $\Omega_g$, is then approximated by the volume of the Hamming ball of radius $d_{g\to s}$

$$\Omega_g = \sum_{k=0}^{d_{g\to s}} \binom{L}{k} \approx \binom{L}{d_{g\to s}} \tag{7}$$

where we have replaced the sum by its dominant term, valid when the genome length $L$ is large and $d_{g\to s} \ll L$.

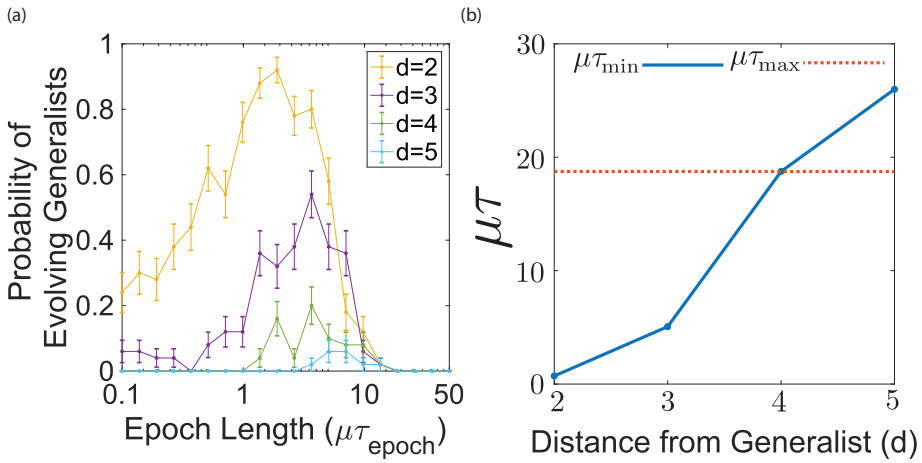


FIG. 2. (a) We compute the probability of evolving generalists for varying Hamming distances from the generalist, given by $d$, using the method described in Section I C. We note that as $d$ increases, the time at which generalists begins to be evolved increases, though the time at which they probability decays remains fixed. (b) We compute $\chi_{s\to g}$ for each initial condition and take $\tau_{\min}$ to be the smallest time where $\chi_{s\to g} > 0.6$. We find that this rises, eventually rising above $\tau_{\max}$ for large enough $d$. By construction, $\tau_{\max}$ is independent of $d$.

From Eq. 5, the number of generalists remaining at time $t$ is:

$$P_{\text{generalists}}(t) = \mathrm{e}^{-\mu t} \sum_{k=0}^{d_{g\to s}} \sinh^k\left(\frac{\mu t}{L}\right) \cosh^{L-k}\left(\frac{\mu t}{L}\right) \binom{L}{k}$$
$$\approx \mathrm{e}^{-\mu t} \left(\frac{\mu t}{L}\right)^{d_{g\to s}} \Omega_g \tag{8}$$

where we have once again replaced the sum by its dominant term, assumed that $\mu t/L$ is small (valid when genome length $L$ is large), and used Eq. 7 to write $\binom{L}{d_{g\to s}} \approx \Omega_g$.

Then, $\tau_{\max}$ is defined as the time taken for the occupancy of the generalist region to fall below 1 individual, $P_{\text{generalists}}(\tau_{\max}) = 1/N$, i.e. the solution of

$$\mathrm{e}^{-\mu t} \left(\frac{\mu t}{L}\right)^{d_{g\to s}} = \frac{1}{N\Omega_g} \tag{9}$$

For $t > 1\mu$, the expression on the left hand side is dominated by the exponential; we thereby solve for $t$ to obtain:

$$\tau_{\max} \sim \frac{1}{\mu} \log \Omega_g N \tag{10}$$

For the parameters used in the simulations ($N = 500$, $\Omega_g \approx 7 \times 10^5$, as computed from Eq. 3), we have $\mu\tau_{\max} \approx 20$, which is consistent with the numerical results (see $\chi_{g\to g}$ in Fig. 4c in the main text).

*Existence of an intermediate timescale*

The existence of an intermediate timescale $\tau_{\text{epoch}}$ that produces generalists requires $\tau_{\min} < \tau_{\max}$. However, the above expressions make it clear that $\tau_{\min} < \tau_{\max}$ only if (a) the initial condition is close enough to generalists (small $d_{i\to g}$), (b) the fraction of generalists relative to specialists is large enough (large $\Omega_g$).

Fig SI 2a shows how the yield of generalists at intermediate timescales disappears as the initial conditions are made less favorable. This panel was constructed using the simulation method described in Section I C with variations on the initial condition for $\mathbf{x}$ such that the Hamming distance between $\mathbf{x}$ and the generalist varied by some distance d.

Fig SI 2b shows that $\tau_{\min}$ rises in this limit and exceeds $\tau_{\max}$. We approximated $\tau_{\min}$ by computing the smallest time for which $\chi_{s\to g} > 0.6$ and $\tau_{\max}$ by computing the largest time for which $\chi_{g\to g} > 0.6$.

When $\tau_{\min} < \tau_{\max}$ is not satisfied, there is no intermediate timescale. The time needed for generalists to specialize is shorter than the time needed to evolve generalists from specialists. In this case, the chirp protocol described below is still successful at producing generalists.

### F. Simultaneous presentation of antigens

The cycling strategy explored in this paper may not be practical in the fast limit in the context of B-cell affinity maturation. A common practical alternative is vaccination with a cocktail of antigens, i.e., simultaneous exposure to multiple antigens.

Such simultaneous exposure to multiple antigens is mathematically equivalent to fast cycling of those antigens if specific microscopic assumptions about antibody-antigen interactions in germinal centers hold[1]. During the affinity maturation process, folicular dendritic cells (FDCs) host antigens on their surface for B-cells to interact with, and if antibodies expressed on B-cells bind the antigens with high enough affinity, B-cells internalize those antigens. This process enables those B-cells to avoid apoptosis, and thus proliferate and continue affinity maturation.

There are two currently experimentally unresolved hypotheses about antigen presentation by FDCs:

1. Antibodies are fit only if they can bind *ALL* presented antigens: In this hypothesis, FDCs only present a single antigen or present antigens in a spatially heterogeneous manner. Consequently, each B-cell is randomly exposed to a single antigen at the selection stage of the affinity maturation process. A B-cell must be able to bind *ALL* presented antigens to survive selection.

Such simultaneous presentation is qualitatively similar to the fast cycling limit studied in this paper. When presented with such a cocktail vaccine, antibodies starting from a naive repertoire are expected to go extinct since such antibodies typically cannot bind all antigens with high affinity, as seen in the experiments of [1].

2. Antibodies are fit if they can bind *ANY* one of the presented antigens: If the FDCs present the antigens in a homogeneous manner, each B-cell only needs to bind *ANY* single one of the presented antigens to avoid apoptosis.

In this case, specialists are fit enough to survive early rounds of selection and evolve generalists. But generalists cannot be maintained in preference over specialists unless selection pressures are fine tuned (e.g., specialists are strongly out-competed by generalists once generalists evolve, despite specialists having significant fitness to begin with).

### G. Death and fast cycling

In the fast cycling limit, $\tau_{\text{epoch}} \to 0$, the fitness of specialist antibodies is the average of their fitness in different environments; as seen Eqn. 2, this fitness is $s(\epsilon - 2)$.

As discussed above for simultaneous presentation, we only consider the case where fast cycling corresponds to hypothesis (1), where antibodies need to bind all antigens to survive. Hence, we need $s(\epsilon - 2)$ to be sufficiently negative, so that a specialist population of size $N$ typically dies out before reaching the generalists in this fast cycling limit. Since the latter process takes time $\tau_{\min} \sim d_{\text{i}\to\text{g}}/\mu$ as derived earlier and the initial population size is $N$, the condition for a specialist population to go extinct in the fast limit is $N\exp(s(\epsilon-2)\tau_{\min}) \sim N\exp(s(\epsilon-2)/\mu) < 1$. Assuming that $1 < \epsilon < 2$, we find $s > \mu \log N$ as a conservative criterion independent of $\epsilon$.

Thus, cycling is necessitated because of population extinction in the fast cycling limit (or equivalently, in the averaged environment). Population extinction does appear to be relevant in affinity maturation [1]. However, given this reliance on population extinction, one can ask whether the cycling strategies proposed here are relevant to other problems. For example, in other evolutionary contexts, can one evolve generalists easily in the fast cycling or averaged environment limit by violating the $1 < \epsilon < 2$ condition?

In fact, reducing death in this manner reduces purifying selection and hence does not always make it easier to evolve generalists. To see this, note that without death, the average fitness of specialists is positive. Consequently, the purifying selection needed to proliferate generalists over specialists is much weaker. Such reduced purifying selection is especially relevant in evolutionary processes that terminate at finite population sizes.

As a concrete example, in Sec I of the SI, we have run simulations of an affinity maturation process that terminates once the population exceeds a threshold, a realistic termination criterion[1]. There, we find that if population death is removed (i.e., $\epsilon > 2$), we still fail to find generalists in the fast limit because the germinal centers are filled with a large number of specialists which typically terminates the process. Thus, the principles developed here have larger relevance to any context of evolving generalists where there is sufficient purifying selection.

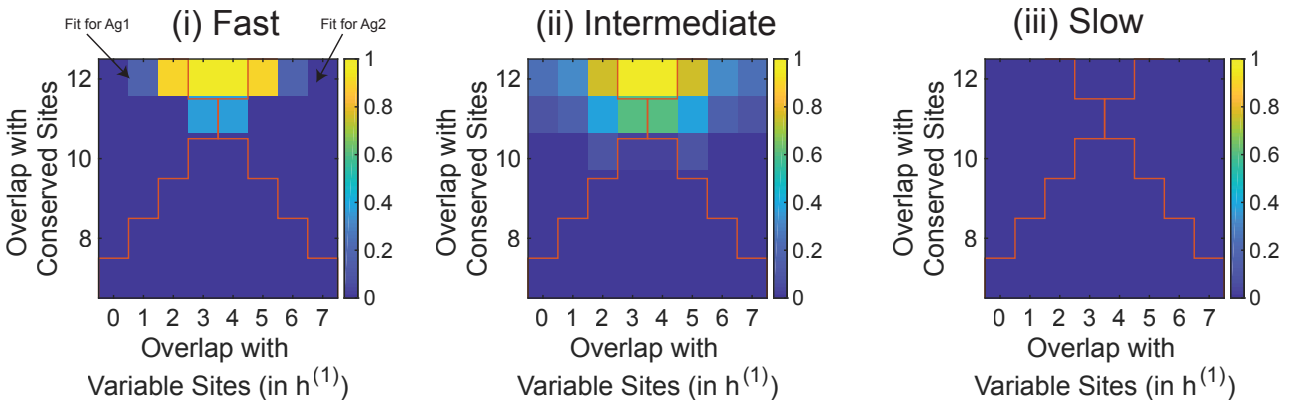


FIG. 3. Intermediate cycling increases the effective attractor size of generalists in models of entropically disfavored generalists. Here, the color in the color map represents the proportion of trials leading to a generalist, where yellow corresponds to all trials, while blue corresponds to no trials leading to a generalist. We initialize our population at size $N = 10$ with a carrying capacity $K = 500$. We maintain $L = 19$, $L_c = 12$, and $T = 11$. We ran our simulations for a total of ten epochs. We ran each ordered pair of conserved and variable matches for 50 replicates in environment $\eta = 1$ and assumed symmetry over environments. (i) Fast cycling ($\mu\tau_{\text{epoch}} = 0.05$) can only evolve generalists when initial conditions are already very close to generalists. (ii) Intermediate cycling ($\mu\tau_{\text{epoch}} = 3$) increases the number of viable initial conditions. (iii) Slow cycling ($\mu\tau_{\text{epoch}} = 15$) does not lead to generalists for any initial condition since this limit is unable to maintain generalists.

### H. Dependence on initial repertoire

In the main paper, we initialize our population from a genotype that exactly binds the antigen characterized by genotype $\mathbf{h}^{(1)}$. In SI Fig 3, we show that not only does intermediate cycling increase the likelihood of evolving generalists from a given initial condition, it also increases the number of initial conditions (e.g., initial B-cell repertoire) that can lead to generalists. Hence, intermediate cycling increases the effective attractor size of the fitness peak associated with a generalist. In SI Fig 3, we consider cycling fast ($\mu\tau_{\text{epoch}} = 0.05 < \mu\tau_{\text{min}}$), cycling at an intermediate rate ($\mu\tau_{\text{epoch}} = 3$), and cycling slowly ($\mu\tau_{\text{epoch}} = 15 > \tau_{\text{max}}$). Blue regions in the heatmap correspond to initial conditions that led to few surviving populations after ten epochs. Yellow regions in the heatmap correspond to initial conditions that led to many surviving populations after ten epochs. We ran each ordered pair of conserved matches and variable matches for 50 replicates in environment $\eta = 1$ and symmetrized over environments.

### I. Chirp protocol

**Trade-off in fixed frequency cycling:** The anticorrelated behavior of $\chi_{\text{s}\to\text{g}}$ and $\chi_{\text{g}\to\text{g}}$ is indicative of a trade-off between evolving generalists and maintaining them in the population.

We first assess this by only considering simulation runs used in Fig 2b that did not result in extinction and computed the number of generalists at the end of such simulations. This number is plotted in Fig 3a. We note that as epoch length increases, the number of generalists remaining in the population decreases, but the probability of a population evolving a single generalist increases.

To illustrate this point we compute two quantities for each simulation for a given $\tau_{\text{epoch}}$:

- The number of generalists (if non-zero): The number of generalists is simply the average number of generalists in the population for simulations where the population survived an evolutionary run. This is plotted on the y-axis of Fig 3c.

- The probability of a surviving population: The proportion of trials for a given $\tau_{\text{epoch}}$ that a population does not go extinct during the evolutionary run. This is plotted on the x-axis of Fig 3c.

By plotting these two quantities against each other, as in Fig 3c in black dots, we observe a tradeoff front.

**Chirp Cycling Breaks the Trade-off:** This trade-off leads us to proposing a 'chirp' protocol. In a 'chirp' cycling protocol, we decrease the length of the epoch using a multiplicative factor after each cycle, enabling us to take advantage of high probability of population survival (favored by slow cycling) and still obtain high yield (favored by

fast cycling). We update $\tau_{\text{epoch}}$ according to the following rule:

$$\tau_{\text{epoch}} \leftarrow k\tau_{\text{epoch}} \tag{11}$$

where $k$ is some number smaller than 1. We continue evolving the population until $\tau_{\text{epoch}} << \lambda$. Plotted in Fig 3a is a time trace of the population size and the fraction of the population that is of a generalist genotype. Generalist fraction is $\frac{\text{Number of Generalists}}{\text{Population Size}}$. We note that as the length of each epoch decreases, the generalist fraction decreases less in time, until eventually, it remains stabilized at $\approx 1$. Additionally, fluctuations in population size are suppressed. We ran the chirp protocol for 25 replicates and computed the number of generalists at the end of each run, if the population survived the run, and the probability that the population survived a chirp protocol. Plotting this in Fig 3d. demonstrates that the tradeoff boundary has been broken. Finally, we compared the chirp protocol to fixed frequency cycling for $L_c$ ranging from 11 to 14. We computed the mean number of generalists observed over 50 replicates in fixed frequency cycling and plot the results in Fig 3b. We compare this to the mean number of generalists discovered under chirp cycling. For this particular chirp, we set $\kappa = \frac{1}{6}$ and initial $\mu\tau_{\text{epoch}} = 5$. We demonstrate that even in regimes where entropic cost is high, chirped cycling can yield generalists robustly.

We note that a similar tradeoff can be observed in Figure 4c of the main text. This indicates that the chirp protocol will work for models with other fitness landscape topologies, so long as there exists a tension between $\chi_{\text{s}\to\text{g}}$ and $\chi_{\text{g}\to\text{g}}$.

**Chirped Cycling Evolves Generalists when $\tau_{\max} < \tau_{\min}$** By implementing the same chirped strategy for initial conditions where $\tau_{\max} < \tau_{\min}$, we find that we have high generalist yield at rates match are near the maximum probability of evolving generalists of the fixed frequency evolutionary runs, as demonstrated in Fig 3b.

## II. GENERALISTS SEPARATED BY VALLEYS

### A. Model

We model the fitness landscape of an antibody binding to an antigen with multiple epitopes through a phenomenological construction, inspired by Hopfield's spin glass landscape.

Consider an antigen $\eta$ with $P_\eta$ epitopes (i.e., sets of residues on the antigen that form binding locations for antibodies). Suppose that for each epitope, an antibody with sequence $\mathbf{h}_\alpha^{(\eta)}$ of length $L$ with entries $\pm 1$ binds with high affinity (with $\alpha \in \{1, ..., P_\eta\}$ indexing epitopes). Then, the overall binding affinity to antigen $\eta$ of any antibody with sequence $\mathbf{x}$, of length $L$ with entries $\pm 1$ is taken to be:

$$F^{(\eta)}(\mathbf{x}) = s \sum_\alpha \kappa_\alpha^{(\eta)} \left( \frac{\mathbf{h}_\alpha^{(\eta)} \cdot \mathbf{x}}{L} \right)^p \tag{12}$$

This construction naturally produces islands of high fitness around the epitope-binding antibodies $\mathbf{h}^{(\eta)}$, separated by regions of low fitness. Our results are tied to the topology of fitness islands and not to details of the functions used to achieve them, provided genotypic space is of sufficiently high dimension. The specific mathematical choice of $p$ has a limitation set by capacity; in a sequence space of dimension $L$, this method only allows us to program fewer than $\alpha L^{(p-1)}$ fitness islands. Beyond this 'capacity', there is a spin glass transition and the mathematical function above actually models a glassy landscape with many other fitness peaks. We can increase this capacity by increasing the non-linearity $p$ of the model. In what follows, we choose $p$ large enough to stay under this spin glass transition. Note that in the mathematical construction of this landscape , $-\mathbf{x}$ is as equally fit as $\mathbf{x}$. Such a degeneracy can be lifted by adding a linear term, $\sum_{i=1}^{L} h_i x_i$ to the fitness function. However, we found we did not need such a term since the degenerate pairs of fitness peaks, $\mathbf{x}$ and$-\mathbf{x}\}$, are far from each other in sequence space. Here, $\kappa_\alpha^{(\eta)}$ is the binding affinity of the ideal antibody $\mathbf{h}^{(\eta)}$ to its cognate epitope. We shall take epitope $\alpha = 1$ to be present for all antigens, thus defining the generalist. In keeping with the assumption that the generalist is less fit in any landscape, we take its binding affinity $\kappa_1^{(1)} = \kappa_1^{(2)} = 0.8$ and all other $\kappa_\alpha^{(\eta)} = 1$.

**Fitness penalty for generalists:** We impose that the height of the fitness peak associated with the generalist is lower than the peaks of the specialist ($\frac{\kappa_1^{(\eta)}}{\kappa_{\alpha\neq 1}^{(\eta)}} = 0.8$), reflecting fitness costs associated with being a generalist relative to specialists in a fixed environment.

### B. Population Simulations

We simulate a population of antibodies evolving in these landscapes by implementing a Moran process[6] with three events:

- **Environment shifts** with a deterministic rate, $\frac{1}{\tau_{\text{epoch}}}$
- **Mutation** with a rate, $\mu$ per individual, where a single site on $\mathbf{x}$ is bit-flipped
- **Reproduction** with a rate, $\lambda$ per individual, where an individual is selected from the population with probability proportional to $\exp(F^{(\eta)})$, with $F^{(\eta)}$ defined above with $p = 2$ (Hopfield model), to be duplicated and another individual from the population to be removed with uniform probability

and a population size of $N$.

In our population simulations, we impose the following values for our simulation parameters. We fix the total number of epitopes for each landscape, $P_\eta = 11$ across all $\eta$, keeping just one generalist. We impose sequence length to be $L = 100$, generating each optimal epitope-binding antibody randomly. We initialize our simulations from a monoclonal specialist initial condition of size $N = 100$ unless otherwise specified. We impose a per individual mutation rate of $\mu = 0.25$ and a reproduction rate of $\lambda = 1$. The overall selection strength is set to $s = 0.1$.

### C. Fig 4b: Resonance Peak in Generalist Discovery as a function of $\tau_{\text{epoch}}$

We ran our simulation for 50 replicates from random monoclonal initial conditions of size $N = 100$ with sequences of length $L = 100$. We initialize the landscapes with 10 specialist antibodies and 1 generalist antibody (i.e. $P_\eta = 11$ for all $\eta$). Our simulations were run for a total of 30 $\tau_{\text{epoch}}$s. We swept over $\tau_{\text{epoch}}$s. We note that during the simulations, regardless of the frequency of environmental shifts, the population remains tightly clustered as it evolves in time. This behavior corresponds to evolution in the strong selection and weak mutation limit. We consider a generalist to have been discovered if, after a run, there exists at least one individual whose overlap with the generalist antibody is 90%. Using these simulations, we demonstrate that an intermediate regime of switching enhances the discovery of generalists. We plot the results of these simulations in Fig. 4b.

**Time traces of the population from its initial condition:** To identify the reason for this resonant peak, we run the simulation for a 100 cycles and plot in SI Fig 4 how far the population is from its initial condition in hamming distance after a fixed number of epochs We initialize a population at a specialist antibody and a generalist antibody, comparing the behavior for fast cycling and slow cycling. We show that for generalist initial conditions, regardless of cycling rate, the population remains in the generalist. There is some fluctuation out but the population returns to the generalists often. Given enough time, the population will escape though this is a slow process. However, for specialist initial condition, slow cycling enables the population to escape its initial condition, while fast cycling does not allow such escape.

### D. Timescale Analysis

**Computing the minimum epoch length, $\tau_{\text{min}}$:**

Cycling benefits evolution of generalists in fitness landscapes with valleys by enabling the population to escape specialist peaks in one landscape by evolving subject to a different fitness landscape. To escape from a specialist peak in environment 1, the population must accrue enough mutations when subject to environment 2, such that when environment 1 returns again, the population is not likely to return to the original specialist peak.

We first consider the strong selection limit $sN \gg 1$. Consider a monoclonal population at a specialist peak $\mathbf{h^{(1)}}_\alpha$ in fitness landscape $F^{(1)}$. When such a population is now subject to landscape $F^{(2)}$ for a time $\tau_{\text{epoch}}$, $\mathbf{h^{(1)}}_\alpha$ serves as an initial condition of typical low fitness and will evolve towards a fitness peak $\mathbf{h^{(2)}}_\beta$ in $F^{(2)}$. If we switch back to $F^{(1)}$ after a sufficiently long time, the population genotype $\mathbf{x}$ will be sufficiently mutated compared to $\mathbf{h^{(1)}}_\alpha$ that the population will likely fix to an alternative fitness peak $\mathbf{h^{(1)}}_\beta$ in $F^{(1)}$. Let us assume that the number of such mutations needed is $d_{12}$.

Since the population in genotype $\mathbf{h^{(1)}}_\alpha$ is typically of low fitness in landscape $F^{(2)}$, most mutations are beneficial. Then, in the strong selection limit, the time needed to acquire $d_{12}$ beneficial mutations is by the mutation rate,

$$\tau_{\text{min}} \approx \frac{d_{12}}{\mu}.$$

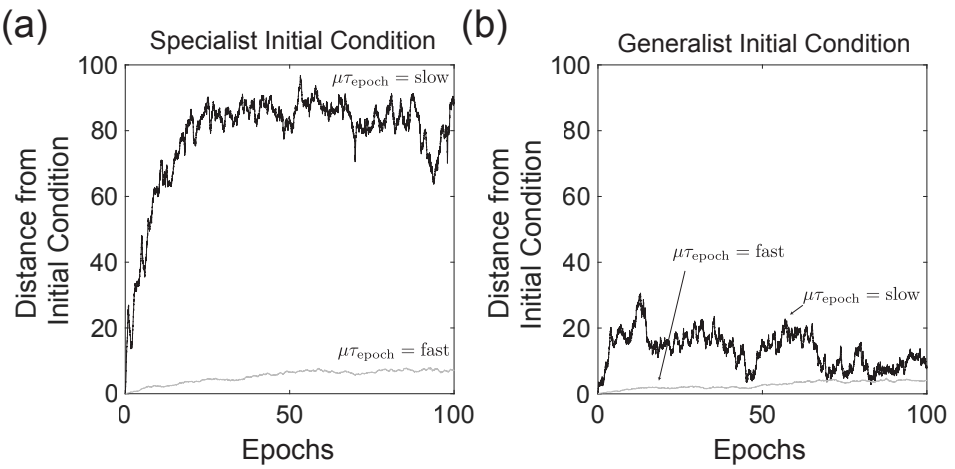


FIG. 4. Specialist populations evolve significantly through sequence space for intermediate timescale cycling but not fast cycling; generalists do not evolve significantly for any timescale cycling. (a) An initially specialist population does not evolve away from the initial genotype for fast cycling. However, with slower cycling, the population evolves to a significantly different genotype(s). (b) An initially-generalist population does not significantly evolve away from the initial genotype for fast or slow cycling.

This minimal number of mutations $d_{12}$ to escape the 'attractor basin' of a fitness peak $\mathbf{h^{(1)}}_\alpha$ is model dependent. - $d_{12}$ depends on the size of the attractor basin around $\mathbf{h^{(1)}}_\alpha$ and the correlations between $F^{(1)}$ and $F^{(2)}$. In our Hopfield-inspired model of fitness landscapes $F^{(i)}$, if the fitness peaks are randomly distributed in a sequence space of length $L$, then the empirical value of $d_{12} \sim \frac{1}{4}L$. In real fitness landscapes, this distance $d_{12}$ can vary widely for different specialists which can have attractor regions of different size.

**Computing the maximum epoch length, $\tau_{\max}$:**

Unnecessarily long times $\tau_{\text{epoch}} > \tau_{\min}$ spent in each environment is counter-productive. To see this, note that specialists are most likely to evolve to generalists in a short duration of time after an environmental switch. Any extra time spent $\tau_{\text{epoch}} > \tau_{\min}$ in the same environment is simply 'dead time' that does not increase the yield of generalists further. Hence the effective rate of evolving generalists from specialists falls as $1/\tau_{\text{epoch}}$ for $\tau_{\text{epoch}} > \tau_{\min}$.

Meanwhile, existing generalists can specialize again. Let the rate of this process be $r_{g\to s}$. The yield of generalists is reduced when this escape rate $r_{g\to s}$ from generalists to specialists is larger than the switching-induced rate from specialists to generalists $1/\tau_{\text{epoch}}$. Hence, $\tau_{\max} \sim 1/r_{g\to s}$.

The rate $r_{g\to s}$ at which generalists specialize is easily estimated since the fitness landscape does not change in time for sequences near the generalist. Hence this process is the well-studied process of an asexual population crossing a fitness valley by picking up a sequence of deleterious mutations. This process has been studied in numerous regimes with different assumptions about population sizes, selection pressure [7–9]. Here, we assume strong selection and weak mutation, allowing us to use the simple result $r_{gs} \sim \mu\, e^{-N\Delta F_g}$ result obtained from the analogy of statistical physics and population dynamics; population size $N$ plays the role of temperature and $\Delta F_g$, the fitness difference between the generalist peak and the fitness valley, plays the role of an energy barrier. Hence,

$$\tau_{\max} \approx \frac{\exp(N\Delta F_g)}{\mu}.$$

Real populations can often violate these assumptions; in that case, any other relevant result[7–9] for valley crossing rates can be used in place of $r_{g\to s}$.

### E. Fig 4c: Transitions Amongst Specialists and Generalists: $\chi_{\mathbf{s}\to\mathbf{g}}$ and $\chi_{\mathbf{g}\to\mathbf{g}}$

The presence of a resonant peak in Fig. 4b is suggestive an underlying tension between discovering the generalist and escaping the generalist, similar to that in the earlier model of entropically disfavored generalists. As such, we re-introduce the quantities $\chi_{\text{s}\to\text{g}}$ and $\chi_{\text{g}\to\text{g}}$:

- $\chi_{\text{s}\to\text{g}}$ is the proportion of trials initialized from a monoclonal specialist initial condition that evolve a generalist

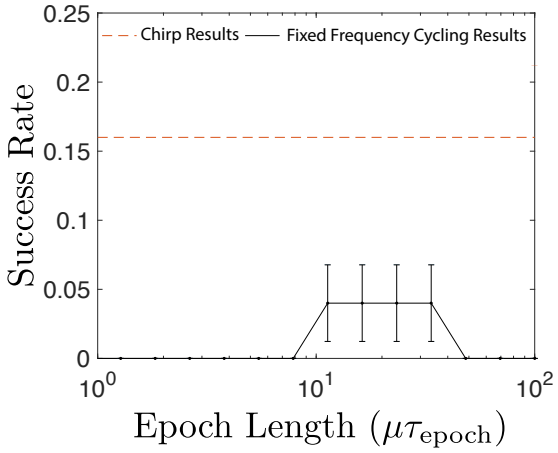


FIG. 5. We compute the success rate of finding a generalist using fixed frequency cycling and chirped strategies when the fitness of binding a generalist is 70% of binding a specialist site. We see, in black, that fixed frequency cycling finds generalists with very low probability. However, in orange, we see that chirped protocols, where $\tau_{\text{epoch}}^{(n+1)} \leftarrow \frac{15}{16}\tau_{\text{epoch}}^{(n)}$ and $\mu\tau_{\text{epoch}}^{(0)} = 100$, find the generalists at a probability $0.16 \pm 0.05$, which is higher than the best fixed frequency cycling strategy. This demonstrates the success of chirped strategies.

(i.e. a single member of the population matches 90% of the generalist) antibody within 30 epochs of an evolutionary run for a given $\tau_{\text{epoch}}$

- $\chi_{\text{g}\to\text{g}}$ is the number of trials in which a population, initialized from the generalist initial condition, maintains 20% of its population in the generalist (i.e. a given antibody maintains 90% overlap with $\mathbf{h}_1^{(\eta)}$) after 30 epochs for a given $\tau_{\text{epoch}}$.

We chose 30 epochs in the definitions above as 30 epochs are needed to give the population enough time to accrue enough cycling-induced stochasticity to explore genotype space. This extension was not necessary for entropically disfavored generalists because multiple epochs are not needed to induce cycling induced stochasticity in such landscapes. We discuss cycling induced stochasticity in more detail in SI Section II G. Plots for $\chi_{\text{g}\to\text{g}}$ and $\chi_{\text{s}\to g}$ are shown in Fig. 4c and illustrate the same behavior as in the entropically disfavored models.

### F. Chirping in Rugged Landscapes

Here, we demonstrate that chirp cycling provides benefits over fixed frequency cycling in fitness landscapes where peaks are separated by fitness valleys. We begin by changing the binding affinity of the antibody to the generalist site to $\kappa_1^{(1)} = \kappa_1^{(2)} = 0.7$. This results in the performance of fixed frequency cycling decreasing dramatically. Chirping, however, continues to provide generalists in a robust manner, as demonstrated in SI Figure 5.

### G. Cycling-induced variance and correlations between environments

To illustrate how cycling enables the discovery of generalists, we consider population trajectories during cycling. We consider the impact of the initial condition of the population on these trajectories and the impact of the correlation structure between the different environments. To this end, we introduce a measure of correlation and introduce a new simulation to capture the effective behavior of the population.

#### *1. Definition of Correlations between Environments*

We measure the correlation between landscapes, denoted $\langle F^{(1)}|F^{(2)}\rangle_s$ using the following equation:

$$\langle F^{(1)}|F^{(2)}\rangle_s \equiv \frac{c(\mathbf{h}^{(1)},\mathbf{h}^{(1)})}{c(\mathbf{h}^{(1)},\mathbf{h}^{(2)})c(\mathbf{h}^{(2)},\mathbf{h}^{(2)})} \tag{13}$$

where we have defined the function $c(\mathbf{h}^{(1)}, \mathbf{h}^{(2)}) = \sum_{\alpha_1,\alpha_2=2}^{P_1,P_2} \mathbf{h}^{(1)}_{\alpha_1} \cdot \mathbf{h}^{(2)}_{\alpha_2} / (L\sqrt{P_1 P_2})$.

This measure of correlation is high if the specialist genotypes for different antigens are highly similar in pairs; e.g., if each specialist for antigen 1 is similar to a specialist for antigen 2. As seen below, a high correlation by this measure implies that specialist antibodies do not evolve significantly due to cycling and thus generalists are not easily evolved.

Note that this measure is normalized so the measure is unaffected by the diversity $c(\mathbf{h}^{(\eta)}, \mathbf{h}^{(\eta)})$ of specialist genotypes for a single antigen $\eta$.

#### 2. *Modeling Population Trajectories with Single Walkers*

To measure the role of cycling between landscapes and the correlation structure of the landscape, we studied the dynamics of single walkers. This is justified as the population is shown to be roughly monoclonal in its evolutionary trajectories. Single walkers were simulated via the well-known Metropolis-Hasting algorithm[10]. We preserve definitions of $\mathbf{x}$, $F^{(\eta)}(\mathbf{x})$, and all related quantities from before. The process is as follows:

- Randomly select a single site to mutate to create new variant $\mathbf{x}'$ from original $\mathbf{x}$
- Compute fitness of new variant
- Accept new variant with probability $\exp(\beta(F^{(\eta)}(\mathbf{x}') - F^{(\eta)}(\mathbf{x})))$ and repeat.

Because of the differences between single walker dynamics and population dynamics, we include an overall scale for the landscapes, $\beta$. $\beta$ is chosen to be $\beta = 4$.

#### 3. *Fig 4f: Cycling-Induced stochasticity*

We begin by considering antigens with uncorrelated specialists (ie, $\langle F^{(1)}|F^{(2)}\rangle_s \approx 0$). Starting from two initial conditions, a generalist antibody and a specialist antibody for antigen $\eta = 1$, we evolve the walker for $k$ proposals in the presence of antigen 2, and then allowed the walker enough proposals to relax to a stable solution in the presence of antigen 1. By computing the final state for 20 different walkers in a given landscape, and averaging over 20 random landscapes, we can compute the variance in the final positions of the walkers. This is accomplished by computing the average pairwise distance between walkers in the same landscape, and then averaging over landscapes. To demonstrate the importance of the number of proposals, $k$, which is serving as a proxy for $\tau_{epoch}$, we swept over $k$. The result is plotted in Fig. 4f.

We see that when starting from a specialist initial condition, cycling-induced variance rises when $\tau_{epoch}$ is sufficiently large. Generalists, as predicted, are unaffected by cycling, as those genotypes are fit in both environments.

#### 4. *Fig 4g: Impact of Correlation Structures Between Cycled Environments*

We then repeated the same simulations as above with increasing correlation structure between the landscapes. We enforced that $k = 250$ was large enough to ensure high stochasticity in uncorrelated environments. We see that as correlation between the landscapes rises, cycling-induced stochasticity decreases. This indicates that generalist discovery hinges on the landscapes being sufficiently uncorrelated. The results are plotted in Fig. 4g.

### H. More than 2 Antigens

Throughout the main text, we only consider the evolution of generalists in the presence of two antigens. Here, we demonstrate that when we increase the number of antigens, discovery of the generalist becomes easier. We demonstrate this in the Moran simulations by increasing the number of randomly constructed landscapes with a single generalist peak shared across the landscapes. In order to probe a dynamic range of antigen number, we weight the generalist to be smaller than in previous trials, setting $\kappa_1^{(\eta)} = 0.07$, rather than 0.08. We ensure each antigen is presented at least once, cycling for at least 30 epochs. We maintain the parameters used before.

We find that increasing the number of antigens increases the discovery-likelihood of generalists, as plotted in SI Fig 6. We interpret these results as due to effectively reduced correlation between landscapes since correlations shared across a subset of the antigens may not be shared across all antigens. For example, the population can settle into a limit cycle between specialists with only 3 antigens but evolution in the presence of other antigens allows escape

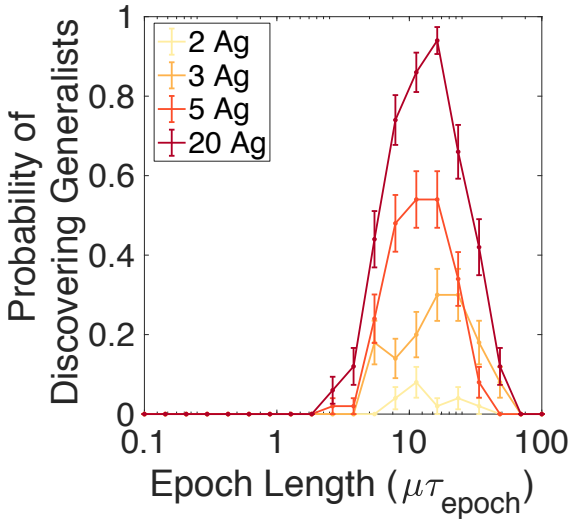


FIG. 6. An increased number of distinct antigens makes it easier to evolve generalists through cycling. We increased the number of distinct antigens in stepsfrom 2 through 20, each with 10 randomly chosen specialist epitopes, and one generalist epitope common to all of them. We cycle between these landscapes using the Moran simulation described in Section II B for 30 $\tau_{\text{epoch}}$. We consider the evolutionary run to have evolved a generalist if at least one antibody has an overlap of 90% with the generalist. We find that increasing the number presented increases the likelihood that generalist genotypes are discovered.

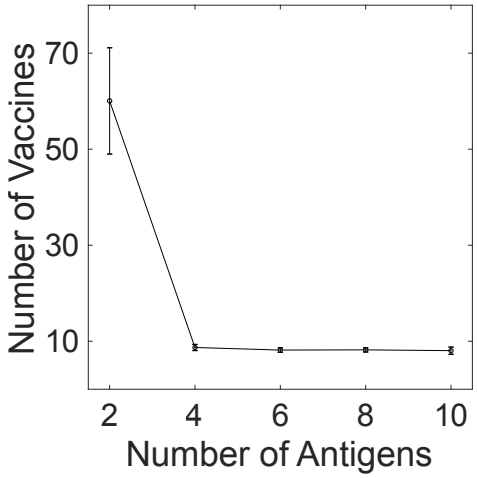


FIG. 7. The number of vaccine shots required to evolve generalists is reduced if distinct antigens are used. We repeat the simulation described in SI Figure 6, except we run until a generalist is discovered for a fixed $\mu\tau_{\text{epoch}} = 40$. We compute the average first arrival time across 50 replicates and the standard error. We find that increasing the number of distinct antigens used decreases the number of doses needed up.

from such cycles. As a result, we increase the rate of evolution from specialists to generalists without enhancing the reverse process.

We further consider the number of cycles needed for the first generalist to appear as a function of the number of antigens. This can be interpreted as the number of vaccine doses needed when using a particular number of strains in the vaccine course. We probe this by running Moran dynamics at a fixed $\mu\tau_{\text{epoch}} = 40$ with $\kappa_1^{(\eta)} = 0.08$, as in other simulations, for as many epochs as needed to discover a generalist. We run 50 replicates of this simulation, reporting the average time at which a generalist first appeared across those replicates. Our choice of epoch length is the epoch length at probability of evolving generalists appears to maximize. We find that the number of vaccine doses decreases as the number of antigen strains increases, as plotted in Fig. 7.

## I. Molecular specificity

A critical requirement of antibodies, including broadly neutralizing antibodies, is molecular specificity. This is, antibodies must show higher binding affinity for their particular target and low binding affinities for all other antigens. We quantify molecular specificity of antibodies in our models by comparing the binding energies of antibodies to antigens featuring the conserved epitope to the binding energies of antibodies to antigens without the conserved epitope. We use these comparisons to identify parameter regimes in which antibodies show molecular specificity; all analyses in the paper are carried out in such regimes.

### 1. *Molecular Specificity in the Entropic Model*

We consider antibodies of length $L = 20$. We impose that the first 15 sites of this antibody bind to a conserved region on some set of antigens. We further impose that for the antibody to be considered to bind to an antigen, its binding energy, as given by Equation 1 to be below $\frac{T}{L} = -\frac{1}{2}$, which is to say that $-\frac{1}{L}\sum_i^L h_i x_i + \frac{T}{L}$ must be negative. We randomly generate 1000 antigens featuring the conserved eptiope and compute the binding energies of the antibody to the antigens. We then compare these binding energies to the binding energy of the same antibody against 1000 randomly generated antigens that are not obligated to feature the conserved epitope. We present the results of this in Fig. 8a, with the antigens featuring the conserved epitope in blue and the antigens without in black. We observe that while the antibody strongly binds all antigens featuring the conserved epitope, it only binds a small fraction of random antigens, showing molecular specificity in this parameter regime.

To ensure molecular specificity is achieved, we must ensure that the number of antigens to which an antibody binds must be small compared to the space of all antigens. We determine the choice of binding energy thresholds $T$ that enforces molecular specificity by first stating that the fraction of antigens bound by a particular antibody is given by $\frac{\sum_{i=T_{\text{sites}}}^{L}\binom{L}{i}}{2^L}$, where the numerator represents the volume of the Hamming ball associated with the antigens that the antibody binds and the denominator represents the space of all antigens. We note that for $T \sim O(1)$, the volume of the Hamming ball is similar to that of the whole space, and for this choice of binding energy threshold, molecular specificity is not achieved. For $T \sim L$, the volume of the Hamming ball can be upper bounded by $\frac{L^L}{T^T}\left(\frac{1}{(L-T)^{L-T}}\right)$. This results in a vanishingly small fraction antigens being bound by our antibody when $T \sim L$. We work only in this regime.

### 2. *Molecular Specificity in Landscapes with Barriers*

We begin by initializing an antibody of length $L = 100$ to bind to a conserved epitope. We then randomly construct 1000 antigens using the prescription described in Section II A with 11 epitopes, 1 fixed across all antigens, and the remaining 10 random. We compute the binding energy of the antibody against these antigens. We then compare these binding energies to the binding energy of the antibody to 1000 antigens, each with 11 randomly generated epitopes. The results are plotted in Fig. 8b, with the binding energies associated with the conserved epitope in blue and the binding energies with random epitopes in black. We observe a large separation between the binding energies of the the antigens featuring the conserved epitope and the antigens without. In general, we expect that for conditions where the number of epitopes is below the Hopfield capacity[11], the probability that a random antigen is bound by an antibody that does not bind one of its epitopes to be vanishingly small.

## J. HIV Antibody Data

The success of our proposed cycling strategy depends on specific assumptions about correlations between different antigens. In particular, antigens need to be sufficiently correlated in that the same generalist antibodies can bind them (e.g., the antigens share a epitope). And yet antigens need to be sufficiently *uncorrelated*: i.e., specialist antibodies that bind different antigens must be sufficiently distinct as measured by $\langle F^{(1)}|F^{(2)}\rangle_s$ (e.g., the specialist epitopes on antigens must be sufficiently distinct).

We sought to test whether these correlation conditions are met by antibodies evolved in response to real HIV strains.

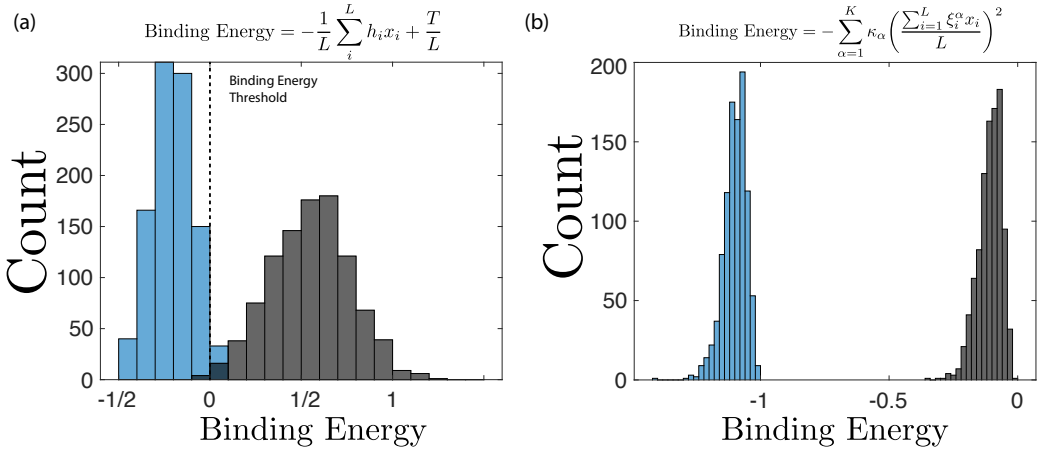


FIG. 8. (a) We fix an antibody of length $L = 20$. We construct 1000 antigens, each of which with a conserved portion of length 15. We compute the binding energies of the antibody with each of these antigens and plot them in blue. We note that the binding energy for each of these falls below the binding energy threshold. We compare this to the binding energy of the same antibody to 1000 random antigens, plotted in black, and find it to be unlikely that the binding energy to fall below the threshold, indicating the antibody is unlikely to bind random antigens. (b) We fix an antibody of length $L = 100$ to bind to some conserved epitope. We generate 1000 antigens, each with 11 epitopes, 1 of which is conserved. We compute the binding energy of the antibody to these antigens and plot them in blue. We compare these binding energies to the binding energy of the same antibody to 1000 antigens, each with their own random 11 epitopes, plotted in black. We observe that a large separation in the binding energies, suggesting it is unlikely that an antibody will spontaneously bind a random antigen.

### *1. Antibody sequences and Binding Affinity Data*

Several works have studied observed antibodies from individuals afflicted with different strains of HIV[12–14]. These works sequenced the observed antibodies, studied their binding affinities to different strains, and proposed intermediate antibodies in between the germline antibody and the discovered broadly neutralizing antibody. They evaluated the binding affinities of each of these antibodies using the ELISA assay. The binding affinity data is presented in SI Table I. The mutational distance between each antibody is given in SI Table II.

**Antibody Sequence Data:** Two classes of antibodies are presented here: mature antibodies observed in patients during their course with HIV and antibody sequences inferred to be [12–14] intermediate between the germline and the mature broadly neutralizing antibody. The natural antibodies appear with the prefix 'CH', and the inferred antibodies, which were synthesized, appear with the prefix 'IA'.

**Antibody Binding Data:** The binding affinity of each antibody to two different strains of HIV, 31D8gp120/293F and 11D8gp120, is evaluated using the ELISA assay. Particular values for binding are presented in table I. We impose a cutoff of 10 log(AUC) to indicate when an antibody has bound an HIV strain. By this rule,

- 31D8gp120/293F is bound by antibodies IA2, IA3, CH105, and CH103.
- 11D8gp120 is bound by antibodies CH186, CH187, CH200, and CH103.

### *2. Constructing landscapes $F^{(1)}$ and $F^{(2)}$*

Let $F^{(1)}$ and $F^{(2)}$ define the fitness landscape of antibody space corresponding to 31D8gp120/293F and 11D8gp120 respectively. In each landscape, the experimentally discovered and synthetically produced antibodies will define disconnected neighborhoods of antibodies that are fit for that landscape.

We begin constructing these landscapes by converting the sequence of each antibody into a binary vector with entries $\pm 1$, noting that each antibody is length $L = 121$. We accomplish this randomly generating a binary vector with entries $\pm 1$ of length $L$ to represent the unmutated common ancestor. Then, using the sequence data given by [12–14], we determine where each antibody differs from the unmutated common ancestor and introduce a binary spin vector for each that preserve the differences from the unmutated common ancestor as presented in the real data. We define the sequences associated with $F^{(\eta)}$ as $\mathbf{h}_\alpha^{(\eta)}$. We set $\mathbf{h}_1^{(1)} = \mathbf{h}_1^{(2)}$ to the sequence of the generalist antibody CH103.

| HIV strain, Ab | CH105 | CH186 | CH187 | CH200 | IA2 | IA3 | CH103 |
|---|---|---|---|---|---|---|---|
| 31D8gp120/293F | 13.52 | 1.13 | 0.00 | 5.80 | 13.34 | 13.01 | 13.63 |
| 11D8gp120 | 8.97 | 13.59 | 10.21 | 10.92 | 9.12 | 6.82 | 10.92 |

TABLE I. Binding affinity of different antibodies (columns) to two different HIV strains (rows), measured via the ELISA assay (units of the logarithm of the area under the curve (logAUC) of the absorbance of the sample)[12–14]. Higher values reflect stronger affinity. We consider an antibody to be a specialist for a strain using a cutoff of 10 logAUC. Note that only CH103 is a generalist in this dataset.

| Antibody | UCA | CH105 | CH186 | CH187 | CH200 | IA2 | IA3 | CH103 |
|---|---|---|---|---|---|---|---|---|
| UCA | 0 | 27 | 8 | 16 | 20 | 27 | 19 | 28 |
| CH105 | 27 | 0 | 25 | 24 | 38 | 21 | 9 | 24 |
| CH186 | 8 | 25 | 0 | 11 | 20 | 25 | 25 | 26 |
| CH187 | 16 | 24 | 11 | 0 | 27 | 23 | 19 | 24 |
| CH200 | 20 | 38 | 20 | 27 | 0 | 37 | 31 | 38 |
| IA2 | 27 | 21 | 25 | 23 | 37 | 0 | 15 | 4 |
| IA3 | 19 | 9 | 25 | 19 | 31 | 15 | 0 | 19 |
| CH103 | 28 | 24 | 26 | 24 | 38 | 4 | 19 | 0 |

TABLE II. Mutational distances (Hamming Distance) between antibody sequences for antibodies observed in an HIV patient who eventually developed bnAbs. Sequences for these antibodies are found in [12–14]. Using the raw sequence data and the binding energy presented in I, we can construct fitness landscapes $F^{(1)}$ and $F^{(2)}$ with fitness peaks that reflect these mutational distances.

We take the fitness of each antibody, represented by $\mathbf{x}$ with entries $\pm 1$ and length $L$, to be:

$$F^{(\eta)}(\mathbf{x}) = s \sum_{\alpha} \left( \frac{\mathbf{h}_{\alpha}^{(\eta)} \cdot \mathbf{x}}{L} \right)^{p}. \tag{14}$$

In the main text, we take $p = 10$ to stay below the spin glass transition for the sequences under consideration. $s$ is a scalar that controls overall magnitude of fitness, which we take to be $s = 200$.

### *3. Simulations*

We simulated evolution using the technique described in Section II G 2. Given that mutation rates in B-cells undergoing somatic hypermutation are taken to be $10^{-3}$ per base pair per division[1], we choose our epoch length to be long enough that the population accumulates 100 mutations. This corresponds to an epoch length that allows 800 total divisions. The initial condition for these simulations was set to be the unmutated common ancestor (UCA). We note that if the population is not started from the UCA, the simulation fails to find successful antibodies.

### *4. Shuffled assignment*

We earlier demonstrated that the correlation structure of the landscape impacted the ability of the landscape to effectively cycle its way to a generalist. Here, we find that

$$\langle F^{(1)} | F^{(2)} \rangle_s = 0.43 \tag{15}$$

which reflects the distance between specialist sequences for the two strains in the data of [12–14] (see Table II). With such low correlations between specialists, the simulations discover generalists around 60% of the time when cycling in this landscape, as shown in Fig. 5b.

To understand how increasing the correlation structure can impact generalist discovery in real data, we artificially shuffled the antibody binding data. In particular, we treated CH186 as a specialist antibody for 11D8gp120 and

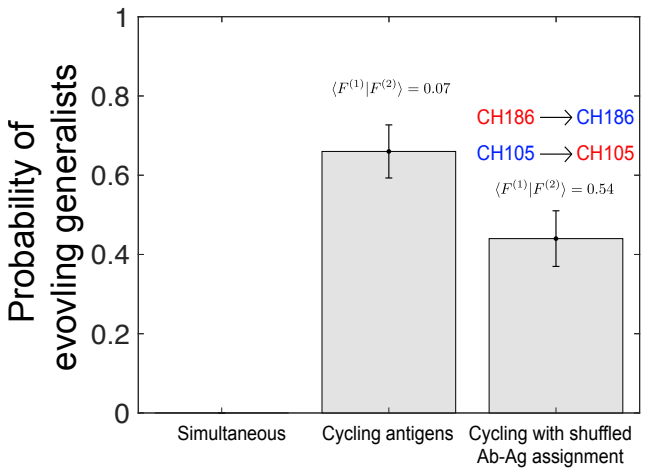


FIG. 9. We construct a landscape only using sites along the antibodies which feature genotypic diversity, reducing the overall length of the antibodies from $L = 121$ to $L = 47$. We set $p = 10$ and $s = 200$ and repeat the simulation described in Section II J 3 and observe qualitatively similar results to those observed in Figure 5 of the main text.

CH105 as a specialist antibody for 31D8gp120/293F and then followed the same construction of landscapes described in II J 2. In the new constrution, we find that the two things are substantially more correlated.

$$\langle F^{(1)}|F^{(2)}\rangle_s = 0.78 \tag{16}$$

Then, after running simulations with changing environments, we find that recovery rates of the generalist drops significantly, as shown in Fig. 5b. This demonstrates that our results are relevant when the fitness landscapes are sufficiently uncorrelated.

### 5. *Sequences Restricted to Variable Regions*

In Fig. 9, we repeat the simulation described above, but construct the patterns using only the antibody sites that are variable across the antibodies considered. This restriction reduces the length of the antibody sequences from $L = 121$ to $L = 47$. As a result, the correlation measure for the unshuffled landscape drops from 0.43 to 0.07, while the correlation measure for the shuffled landscape drops from 0.78 to 0.54. We set the epoch length to be longer than before, as a result of the differences in the magnitudes of correlations and the sequence length. The results of the simulation are shown in Fig.9.

Despite the resulting quantitative differences, qualitatively, these results are similar to those displayed in Figure 5 of the main text for the full $L = 121$ length sequences.

Thus, our conclusions primarily depend on the correlations between the fitness landscape and not the mathematical details of how we construct the fitness landscape. In particular, the dimensionality of sequence space affects our results to the extent that the dimensionality changes correlation structure across environments. We expect the effects of cycling to be weaker in lower dimensions, such as the case explored in [15] where there are fewer paths from specialists to generalists. For example, in 1 dimension, cycling can be entirely unproductive if cycling-induced evolution repeatedly traps the population at specialist peaks adjacent to the generalist genotype.

[1] S Wang, et al., Manipulating the selection forces during affinity maturation to generate cross-reactive hiv antibodies. *Cell* **160**, 785 – 797 (2015).
[2] K Rajewsky, Clonal selection and learning in the antibody system. *Nature* **381**, 751–758 (1996).
[3] YE Maruvka, DA Kessler, NM Shnerb, The birth-death-mutation process: A new paradigm for fat tailed distributions. *PLOS ONE* **6**, 1–7 (2011).
[4] Cobey Sarah, Wilson Patrick, Matsen Frederick A., The evolution within us. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **370**, 20140235 (2015).

[5] Childs Lauren M., Baskerville Edward B., Cobey Sarah, Trade-offs in antibody repertoires to complex antigens. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **370**, 20140245 (2015).
[6] PAP Moran, Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society* **54**, 60–71 (1958).
[7] DB Weissman, MM Desai, DS Fisher, MW Feldman, The rate at which asexual populations cross fitness valleys. *Theor. Popul. Biol.* **75**, 286–300 (2009).
[8] E van Nimwegen, JP Crutchfield, Metastable evolutionary dynamics: crossing fitness barriers or escaping via neutral paths? *Bull. Math. Biol.* **62**, 799–848 (2000).
[9] K Jain, J Krug, Deterministic and stochastic regimes of asexual evolution on rugged fitness landscapes. *Genetics* **175**, 1275–1288 (2007).
[10] N Metropolis, AW Rosenbluth, MN Rosenbluth, AH Teller, E Teller, Equation of state calculations by fast computing machines. *The Journal of Chemical Physics* **21**, 1087–1092 (1953).
[11] DJ Amit, H Gutfreund, H Sompolinsky, Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.* **55**, 1530–1533 (1985).
[12] M Bonsignori, et al., Maturation pathway from germline to broad hiv-1 neutralizer of a cd4-mimic antibody. *Cell* **165**, 449 – 463 (2016).
[13] F Gao, et al., Cooperation of b cell lineages in induction of hiv-1-broadly neutralizing antibodies. *Cell* **158**, 481 – 491 (2014).
[14] HX Liao, et al., Co-evolution of a broadly neutralizing hiv-1 antibody and founder virus. *Nature* **496**, 469 EP – (2013).
[15] E Kussell, S Leibler, A Grosberg, Polymer-population mapping and localization in the space of phenotypes. *Phys. Rev. Lett.* **97**, 068101 (2006).