Supporting Information

Mayer et al. 10.1073/pnas.1600663113

SI Appendix A: Parameterizing a Two-State Markov Chain

The parameterization that we use is based on the average frequency that the chain is in one of its states and on a characteristic timescale of state changes. Concretely, the first parameter is the fraction of generations during which a pathogen is present $\pi_{\text{env}} = \langle x \rangle = \alpha/(\alpha+\beta)$, and the second parameter is the autocorrelation time of the chain, defined from its autocorrelation function: $\langle x_t x_{t'} \rangle - \langle x \rangle^2 = \pi_{\text{env}}(1-\pi_{\text{env}})(1-\alpha-\beta)^{|t-t'|} \equiv \pi_{\text{env}}(1-\pi_{\text{env}})\exp(-|t-t'|/\tau_{\text{env}})$ or $\tau_{\text{env}} \equiv -1/\ln(1-\alpha-\beta)$. We have chosen the autocorrelation time over alternative timescales such as the persistence time of the pathogen $\tilde{\tau}_{\text{env}} = -1/\ln(1-\beta)$, because the autocorrelation time is symmetric in the switching rates: Long stretches of continuous pathogen absence play a role in the choice of a strategy as long stretches of continuous pathogen presence. Our choice of characteristic time provides a measure for the degree of predictability of the next state given the current state. For instance, for a characteristic time of zero, $1-\alpha=\beta$, no information on the next state can be gained from knowing the current state. The parameterization has the additional property that all combinations $\pi_{\text{env}} \in [0,1]$ and $\tau_{\text{env}} \in [0,\infty]$ correspond to valid values of $\alpha \in [0,1], \beta \in [0,1]$, which is not the case for all combinations of $(\pi_{\text{env}}, \tilde{\tau}_{\text{env}})$, for example.

SI Appendix B: Pattern-Search–Based Optimization for Problems with Noisy Function Evaluations

The goal of stochastic programming is to find the minimum over a bounded domain Ω of a function

$$\min_{\boldsymbol{y}\in\Omega} f(\boldsymbol{y}) = \mathbb{E}[F(\boldsymbol{y},\omega)], \qquad \textbf{[S1]}$$

which is not known explicitly, but needs to be approximated by evaluating a function $F(\boldsymbol{y},\omega)$ dependent on a random variable ω. The quality of this approximation can be increased by sampling F multiple times at the expense of higher computational cost. The numerical optimization of the long-term growth rate considered in this work falls into this class. The long-term growth rate can generally not be calculated explicitly but is approximated numerically from long but finite simulations of the population dynamics. These simulations depend on the history of pathogen presence and absence $\boldsymbol{x}_t$, which is a random variable. In the following we use the generic notation of the optimization problem of Eq. **S1**, which can be mapped onto the problem of optimizing long-term population growth rates through $f \leftrightarrow -\Lambda$, $\boldsymbol{y} \leftrightarrow (p,q,p_{\text{uptake}},c_{\text{constitutive}})$, $\omega \leftrightarrow \boldsymbol{x}_t$.

We developed an algorithm to solve a stochastic programming problem with bound constraints. The algorithm combines a compass search, a simple pattern-search algorithm that allows for easily incorporating bound constraints (28), with the idea of adapting the number of evaluations of F dynamically to control the noise in the approximation (29, 30). An advantage of the algorithm over alternative methods for noisy optimization such as stochastic approximation (31) is that it allows one to define stopping criteria in terms of parameter convergence instead of relying on more indirect stopping criteria such as decrease conditions. For the optimization problem considered in this paper, this algorithm works reliably and efficiently enough to allow for the many optimizations needed for a phase diagram such as shown in Fig. 2.

Let us define the set of search directions considered at each iteration as $\mathcal{D} = \{\pm\boldsymbol{e}_i | i = 1,\ldots,n\}$, where $\boldsymbol{e}_i$ is the ith unit vector; and let us further define $P_\Omega(\boldsymbol{y}) = \text{argmin}_{\boldsymbol{y}'\in\Omega}|\boldsymbol{y}'-\boldsymbol{y}|^2$ as the projection of a point $\boldsymbol{y}$ onto Ω (32). The projection onto box constraints considered here is particularly simple and computationally efficient as it just sets coordinate entries outside of the bounds to the bound value. Given an initial guess for the parameter vector $\boldsymbol{y}^0$, an initial step size Δ^0, and an initial number of times F should be sampled N^0, the algorithm proceeds as follows to find the optimal parameter vector to within a tolerance of Δ_{tol}:

1) **Initialize** parameter vector $\boldsymbol{y} \leftarrow \boldsymbol{y}^0$, step size $\Delta \leftarrow \Delta^0$, and number of samples $N \leftarrow N^0$.
2) **While** ($\Delta \geq \Delta_{\text{tol}}$) or ($\boldsymbol{y}$ or N updated during last iteration):
   a) **For each** step $\Delta\boldsymbol{d}$ along positive and negative coordinate directions $\boldsymbol{d} \in \mathcal{D}$:
      i) **If** $f(\boldsymbol{y}') < f(\boldsymbol{y})$ (as judged from N samples of F at both points), where $\boldsymbol{y}' = P_\Omega(\boldsymbol{y}+\Delta\boldsymbol{d})$, then update the parameter vector $\boldsymbol{y} \leftarrow \boldsymbol{y}'$.
      ii) **Else if** new point $\boldsymbol{y}+\Delta\boldsymbol{d}$ is feasible, i.e., $\boldsymbol{y}+\Delta\boldsymbol{d} \in \Omega$, and if $f(\boldsymbol{y}+\Delta\boldsymbol{d}) = f(\boldsymbol{y})$ cannot be ruled out based on N samples of F at both points (criterion below), then either one oversteps the minimum or statistical power is insufficient. Therefore, first try half-step in the same direction, and if it fails increase sampling:
         A) **If** $f(\boldsymbol{y}+(\Delta/2)\boldsymbol{d}) < f(\boldsymbol{y})$ (as judged from N samples of F at both points), then update parameter vector and reduce step size:

$$\boldsymbol{y} \leftarrow \boldsymbol{y} + \frac{\Delta}{2}\boldsymbol{d}$$

$$\Delta \leftarrow \Delta/2.$$

         B) **Else** increase sampling:

$$N \leftarrow 2N.$$

   b) **If** no updates during preceding loop, then contract pattern size:

$$\Delta \leftarrow \Delta/2.$$

For the comparisons between objective function values, we use hypothesis testing on N paired samples of F; i.e., we evaluate $F(\boldsymbol{y},\omega_i), F(\boldsymbol{y},\omega_i)$ for $\omega_i, i = 1,\ldots,N$ and calculate pairwise differences. The hypothesis testing uses a confidence level α, which indirectly controls how much the function is sampled. To correct for the multiple tests performed for different directions, we use a Bonferroni correction by using a confidence level $\alpha/(2n)$ for individual tests, where $2n$ is the number of search directions.

SI Appendix C: Analytical Insight into the Transitions Between Strategies

By analytically solving three simplified problems, we provide additional insights into the choice of strategy. For brevity of notation, we set $c_{\text{def}} = c_{\text{defense}}$, $c_{\text{con}} = c_{\text{constitutive}}$, and $c_{\text{inf}} = c_{\text{infected}}$.

When to Regulate the Response. For pathogens changing with a small characteristic timescale, there is a transition from adaptive to proto-adaptive to innate strategies for standard parameters (Fig. 2) as a function of π_{env}. For all three strategies considered here, the complete population is always protected, $q = 0$, and there is no active acquisition, $p_{\text{uptake}} = 0$. The equation for the instantaneous growth rate at generation t (Eq. **8** of the main text) thus simplifies to

$$z^t = \begin{cases} e^{-c_{\text{def}}} & \text{if } x^t = 1 \\ e^{-c_{\text{con}}} & \text{if } x^t = 0, \end{cases} \quad \text{[S2]}$$

where growth depends only on the absence or presence of pathogen during the current generation regardless of what happened at previous generations. The optimal long-term growth rate can then be calculated analytically by weighting the instantaneous growth rates in the presence and absence of pathogen by the frequency of the two environmental states

$$\Lambda = -\pi_{\text{env}} c_{\text{def}}(c_{\text{con}}) - (1-\pi_{\text{env}}) c_{\text{con}}. \quad \text{[S3]}$$

This expression for the long-term growth rate directly gives us some insight into how the frequency of the pathogen affects how much the response should be regulated. The more frequent the pathogen is, the more often the defense is actually used and thus the less it should be regulated. By maximizing Λ over $c_{\text{con}} \in [0, c_{\text{con}}^{\max}]$ for a given trade-off function $c_{\text{def}}(c_{\text{con}})$, we obtain analytical expressions for the phase boundaries. One finds the following conditions for local optimality of the three strategies:

$\pi_{\text{env}} < \pi_{\text{env}}^{(ap)}$	$\pi_{\text{env}}^{(ap)} \le \pi_{\text{env}} \le \pi_{\text{env}}^{(po)}$	$\pi_{\text{env}}^{(po)} < \pi_{\text{env}}$
$c_{\text{con}} = 0$	$0 \le c_{\text{con}} \le c_{\text{con}}^{\max}$	$c_{\text{con}} = c_{\text{con}}^{\max}$

with

$$\pi_{\text{env}}^{(ap)} = \left(1 - \left.\frac{\mathrm{d}c_{\text{def}}}{\mathrm{d}c_{\text{con}}}\right|_{c_{\text{con}}=0}\right)^{-1}, \quad \text{[S4]}$$

$$\pi_{\text{env}}^{(po)} = \left(1 - \left.\frac{\mathrm{d}c_{\text{def}}}{\mathrm{d}c_{\text{con}}}\right|_{c_{\text{con}}=c_{\text{con}}^{\max}}\right)^{-1}. \quad \text{[S5]}$$

As we assume a convex trade-off shape, we have $\pi_{\text{env}}^{(ap)} < \pi_{\text{env}}^{(pi)}$, which implies a succession of adaptive, proto-adaptive, and innate strategies for increasing π_{env} as seen in Fig. 2. If instead the trade-off function $c_{\text{def}}(c_{\text{con}})$ is concave, then the proto-adaptive phase vanishes.

When to Hedge Bets. For the very frequent pathogens, the optimal strategy is to have protection at all times, whereas for less frequent pathogens some bet hedging is often favored (Fig. 2 and Fig. S3). To understand the transition from bet-hedging innate to deterministic innate immunity, we compare the long-term growth rates of populations, using these strategies. For simplicity, we restrict the analysis to strategies with no heritability, $p = 1-q$, and no regulation, $c_{\text{def}} = c_{\text{con}}$. The fraction of protected individuals is constant across generations and the long-term growth rate can be calculated analytically as

$$\Lambda = \pi_{\text{env}} \ln[(1-p)e^{-c_{\text{inf}}} + pe^{-c_{\text{con}}}] + (1-\pi_{\text{env}}) \ln[1-p+pe^{-c_{\text{con}}}]. \quad \text{[S6]}$$

Optimizing the long-term growth rate over the fraction of protected organisms p yields

$\pi_{\text{env}} < \pi_{\text{env}}^{(0i)}$	$\pi_{\text{env}}^{(0i)} \le \pi_{\text{env}} \le \pi_{\text{env}}^{(io)}$	$\pi_{\text{env}}^{(io)} < \pi_{\text{env}}$
$p = 0$	$0 \le p \le 1$	$p = 1$

with

$$\pi_{\text{env}}^{(0i)} = \frac{e^{c_{\text{con}}} - 1}{e^{c_{\text{inf}}} - 1}, \quad \text{[S7]}$$

$$\pi_{\text{env}}^{(io)} = \frac{1 - e^{-c_{\text{con}}}}{1 - e^{-c_{\text{inf}}}}. \quad \text{[S8]}$$

This shows the existence of three regimes. For rare pathogens tolerance is optimal (as we are looking only at unregulated strategies), for frequent pathogens it is best to always protect, whereas in between bet hedging is favored. The existence of these different phases is a known result in the bet-hedging literature when both phenotypes can survive in both environmental states (33), as is the case here. The assumption $p = 1-q$ makes the derivation of this result exact when the environment itself is memoryless, $\alpha = 1-\beta$. In the presence of temporal correlations in pathogen occurrence, we expect bet-hedging strategies to be favored for a larger range of pathogen frequencies, as they can exploit the predictability of the environment.

When to Acquire Actively. For pathogens with large temporal correlations, the optimal strategy changes from an active, to a mixed, to a passive mode of acquisition (Fig. 2). To understand these transitions, we again turn to an analytical solvable limit. As these strategies are favored in the presence of temporal correlations, the limit of temporally uncorrelated strategies $p = 1-q$ considered in the previous section is not the most pertinent. We turn instead to another analytical solvable limit, in which growth rate differences are very large compared with the generation time, $c_{\text{con}} \gg 1,\ c_{\text{inf}} - c_{\text{def}} \gg 1$. In this limit, the fraction of protected individuals is Markovian as all parents of individuals in the current generation were in the favored state of the last environment (all maladapted individuals die). We note that similar results can be obtained in the limit of large environmental correlation times τ_{env} without assuming completely specialized phenotypes (34). The long-term growth rate can therefore be expressed analytically based on the probabilities Q_{ij} of observing an environmental state i followed by state j ($Q_{00} = (1-\pi_{\text{env}})(1-\alpha),\ Q_{01} = (1-\pi_{\text{env}})\alpha,\ Q_{10} = \pi_{\text{env}}\beta,\ Q_{11} = \pi_{\text{env}}(1-\beta)$) as

$$\Lambda = Q_{00} \ln(1-p) + Q_{10} \ln q + Q_{01} \ln\left[(p + p_{\text{uptake}})e^{c_{\text{def}}}\right] + Q_{11} \ln[(1-q)e^{c_{\text{def}}}] - c_{\text{uptake}}(p_{\text{uptake}}). \quad \text{[S9]}$$

By comparing the terms in which p and p_{uptake} appear in this expression, the strengths and weaknesses of the two acquisition modes become evident. Passive acquisition has a diversification cost due to unnecessary switching into state 1 in the absence of pathogen ($Q_{00} \ln(1-p)$). Active acquisition does not have this penalty, but is more difficult to implement and comes with an extra cost $c_{\text{uptake}}(p_{\text{uptake}})$ dependent on its uptake rate. As the probability Q_{00} is high for rare and temporally correlated pathogens, the relative cost of random acquisition is especially high for these pathogens, where most of the time mutations conferring gain of protection are deleterious. Optimizing the expression of the long-term growth rate over $p, p_{\text{uptake}} \in [0, 1]$, we find the following optimality conditions:

$\pi_{\text{env}} < \pi_{\text{env}}^{(cm)}$	$\pi_{\text{env}}^{(cm)} \le \pi_{\text{env}} \le \pi_{\text{env}}^{(mi)}$	$\pi_{\text{env}}^{(mi)} < \pi_{\text{env}}$
$p = 0,\ p_{\text{uptake}} > 0$	$p > 0,\ p_{\text{uptake}} > 0$	$p > 0,\ p_{\text{uptake}} = 0$

with

$$Q_{00}\left(\pi_{\text{env}}^{(cm)}\right) = \frac{Q_{01}\left(\pi_{\text{env}}^{(cm)}\right)}{g^{-1}\left[Q_{01}\left(\pi_{\text{env}}^{(cm)}\right)\right]}, \quad \text{with } g\left(p_{\text{uptake}}\right) = p_{\text{uptake}} \frac{\mathrm{d}c_{\text{uptake}}}{\mathrm{d}p_{\text{uptake}}}, \quad \text{[S10]}$$

$$\pi_{\text{env}}^{(mi)} = 1 - \left.\frac{\mathrm{d}c_{\text{uptake}}}{\mathrm{d}p_{\text{uptake}}}\right|_{p_{\text{uptake}}=0}. \quad \text{[S11]}$$

Thus, in this limit, a CRISPR-like strategy is favored for rare pathogens, an innate bet-hedging strategy for frequent pathogens, and

mixed strategies in between, in agreement with the numerical results reported in Fig. 2 of the main text.

SI Appendix D: Nonindependent Pathogen–Protection Pairs

The factorization of the recursion relation defining the population dynamics allows us to treat the problem one pathogen at a time. This makes the problem mathematically tractable and the results easily interpretable. Different protection–pathogen pairs can be treated independently, however, only if a number of assumptions are met: The costs must be additive, one protection protects against only one pathogen and vice versa, and the dynamics of different pathogens are independent. A full treatment of the general, nonfactorized problem is outside the scope of this work, but in the following we discuss how relaxing some of these assumptions affects the optimal strategy. Specifically, we consider simple cases with only two pathogen–protection pairs to build intuition of where we expect qualitative changes in optimal strategies and where and how we can relate back to the results for the factorizing case.

Nonadditive Cost of Infection. If the cost of an infection is amplified by coinfections by other pathogens, then we expect the optimal strategies to be similar to the ones emerging for a single pathogen, but with a higher effective cost of infection $c_{\text{infection}}$ (for the influence of a higher cost of infection on the phase diagram see Fig. S3*F*). The effective cost should take into account the extra cost incurred by the presence of a coinfection weighted by its probability of occurrence.

To test this intuition, we consider a simple case with two pathogens, where we impose $c_{\text{constitute}} = c_{\text{defense}}$ and $p_{\text{uptake}} = 0$. A completely unprotected organism pays a cost $c_{\text{infection}}$ if it gets infected by one pathogen and a cost $2c_{\text{infection}} + \nu$ if it gets infected by both. Solving the problem numerically shows that the optimal fraction of protection against the two pathogens increases with ν (Fig. S2) as expected. The Pearson correlation coefficient between being protected against one or the other pathogen remains small even for ν of the order of $c_{\text{infection}}$, meaning that the optimal strategy remains close to the independent case.

Nonadditive Cost of Protection. As with the case of nonadditive cost of infection, we expect nonadditive costs of protection to result in a modified effective cost of protection (for the influence of changing the cost of protection see Fig. S3*G*). However, for a nonindependent cost of protection, an optimal immune strategy might differ significantly from the factorizing case. In particular, the optimal strategy may regulate the total number of protections at a given time to either exploit the economies of scale (if protection against many pathogens is relatively cheaper) or avoid an overburdening cost (if protection against many pathogens at the same time is relatively more costly).

Some of the immune strategies that require a lot of machinery to function, such as vertebrate adaptive immunity or CRISPR-Cas immunity, might come at the expense of a large fixed investment cost, c_{system}, in addition to their state-dependent costs. This nonadditive cost can be viewed as shared equally between all pathogen–protection pairs concerned by the adaptive strategy. It does not break the independence between them, but rather adds an offset cost c_{system}/L, where L is the number of pathogen–protection pairs, which will shift the transition at which adaptive immunity becomes favorable.

Cross-Reactive Protection. In most biological defense systems, there is some degree of cross-reactivity; i.e., defense against several pathogens can be achieved with the same protection. This feature can be incorporated in our framework by introducing a more complicated form of the dependency of the number of offspring on the protection state $\boldsymbol{\sigma}$. We expect the optimal strategy to exploit cross-reactivity by having dissimilar protections that collectively tile the space of possible pathogens (18). Then, the dynamics of pathogens can be effectively reduced to the presence or absence of any of the pathogens within the scope of a given protection.

To validate this intuition, we consider a single protection that is efficient against two pathogens of frequencies $\pi_{\text{env},1}$ and $\pi_{\text{env},2}$. Assume that the cost of defense is the same whether we defend against one or both pathogens, as summarized by the costs in the table below,

$\sigma \backslash (x_1, x_2)$	$(0,0)$	$(1,0)$	$(1,0)$	$(1,1)$
0	0	c_{inf}	c_{inf}	$2c_{\text{inf}}$
1	c_{con}	c_{def}	c_{def}	c_{def},

where (x_1, x_2) indicates which one of the two pathogens is present. If the protection strategy is memoryless ($p = 1 - q$), then the long-term growth rate is

$$\Lambda = (1-\pi_{\text{env},1})(1-\pi_{\text{env},2})\ln r_{00} + \pi_{\text{env},1}(1-\pi_{\text{env},2})\ln r_{10} + (1-\pi_{\text{env},1})\pi_{\text{env},2}\ln r_{01} + \pi_{\text{env},1}\pi_{\text{env},2}\ln r_{11}, \quad \textbf{[S12]}$$

where $r_{x_1 x_2}$ is the average growth rate in environment (x_1, x_2): $r_{00} = pe^{-c_{\text{con}}} + 1 - p$, $r_{01} = r_{10} = pe^{-c_{\text{def}}} + (1-p)e^{-c_{\text{inf}}}$, $r_{11} = pe^{-c_{\text{def}}} + (1-p)e^{-2c_{\text{inf}}}$. The long-term growth rate can be alternatively expressed as

$$\Lambda = (1-\pi_{\text{env,eff}})\ln r_{00} + \pi_{\text{env,eff}}\ln r_{10} + \pi_{\text{env},1}\pi_{\text{env},2}\ln\frac{r_{11}}{r_{01}}, \quad \textbf{[S13]}$$

with $\pi_{\text{env,eff}} = \pi_{\text{env},1} + \pi_{\text{env},2} - \pi_{\text{env},1}\pi_{\text{env},2}$. The last term in this expression is small either for infrequent pathogens ($\pi_{\text{env},1}\pi_{\text{env},2} \ll \pi_{\text{env,eff}}$) or if a large fraction of the population is protected ($1 - p \ll 1$ and hence $r_{01} \approx r_{11}$). Neglecting this second-order term, we are left with the expression corresponding to a single pathogen with effective frequency $\pi_{\text{env,eff}}$, in agreement with our expectation.

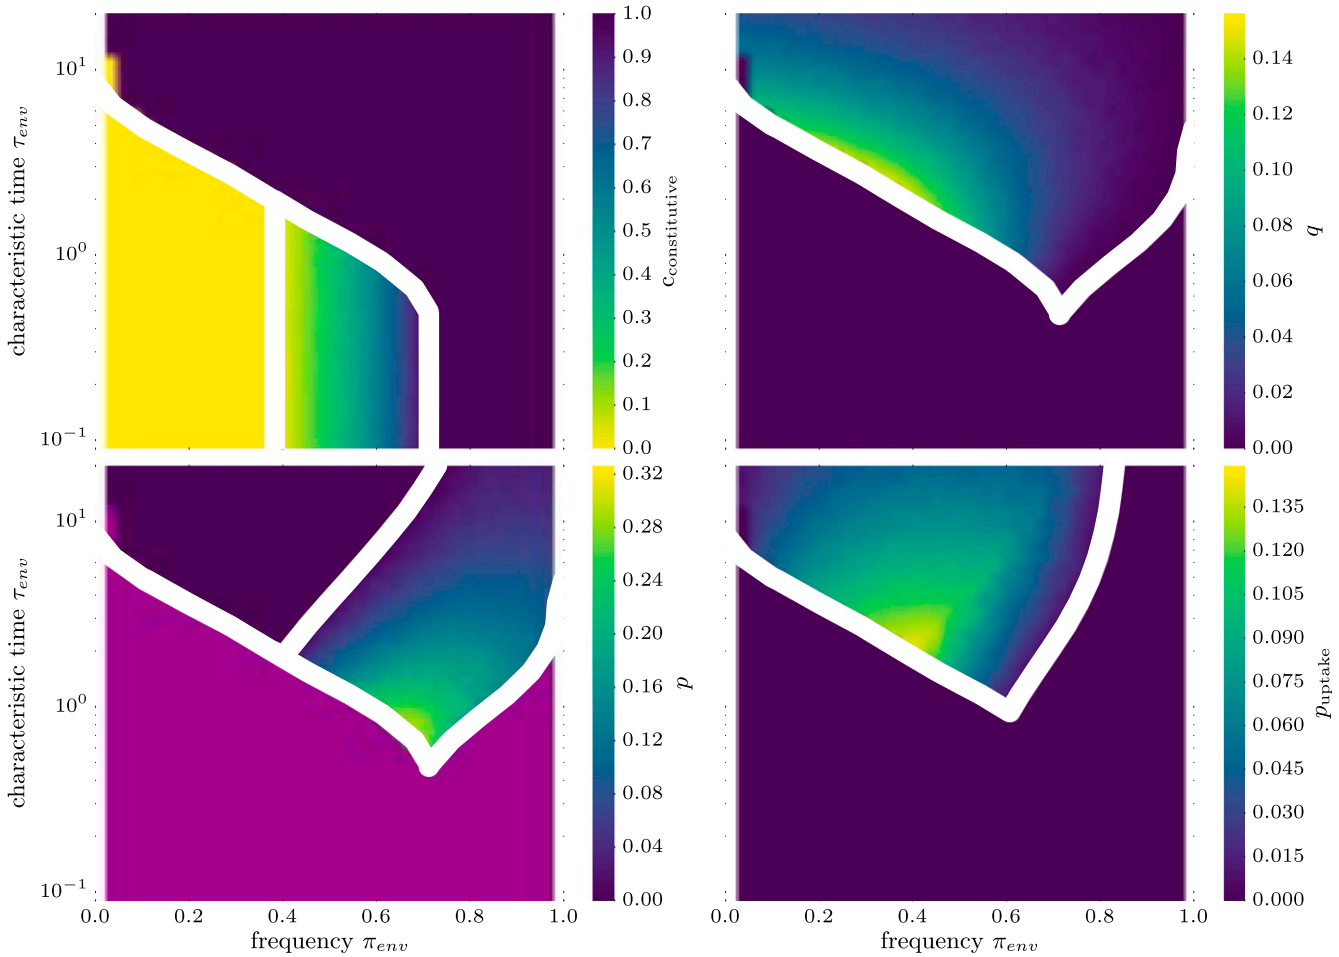

Fig. S1. Optimal parameters from a global optimization of long-term growth rate. Regions where a parameter is unconstrained at the optimum are shown in purple. Phase boundaries pertaining to the shown parameter are in white. A maximum number of 10,000 function evaluations are used for the first phase of the optimization. The second phase of the optimization is terminated at a tolerance in the parameter values of 0.005. The same model parameters as in Fig. 2 are used.

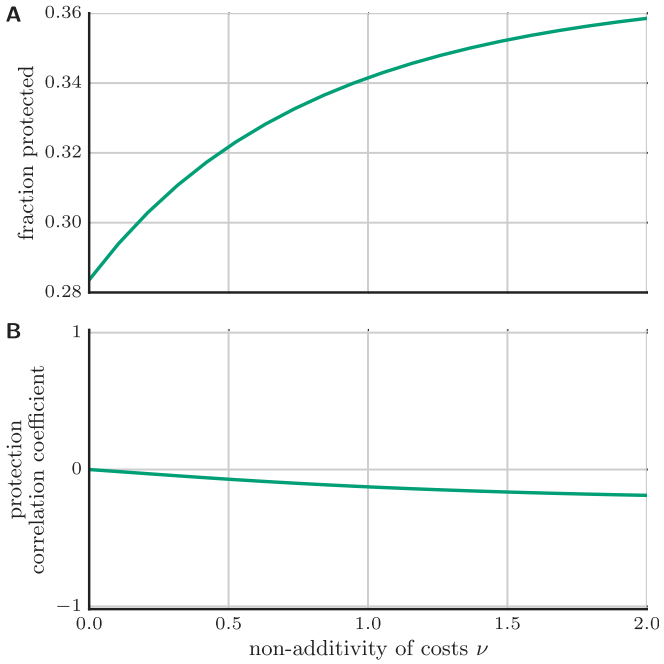

Fig. S2. Optimal protection strategy against two equally frequent pathogens $\pi_{env,1} = \pi_{env,2} = 0.4$ as a function of the degree of nonadditivity of the cost of infection ν. (*A*) Fraction of population protected against a particular pathogen. (*B*) Pearson correlation coefficient between the protection states against the two pathogens. As costs are nonadditive, the problem no longer factorizes and the optimal strategy no longer chooses protections against different pathogens independently. However, here the optimal strategy treats each pathogen almost independently, as measured by the low correlation coefficient. With an increasing cost of coinfection, more protection is needed, in agreement with our intuition that coinfection leads to higher effective costs. Parameters: $c_{infection} = 2$, $c_{defense} = c_{constitutive} = 1$, and optimization of the distribution over protection states respecting the probability simplex constraints, using an accelerated projected gradient algorithm as described in ref. 18.

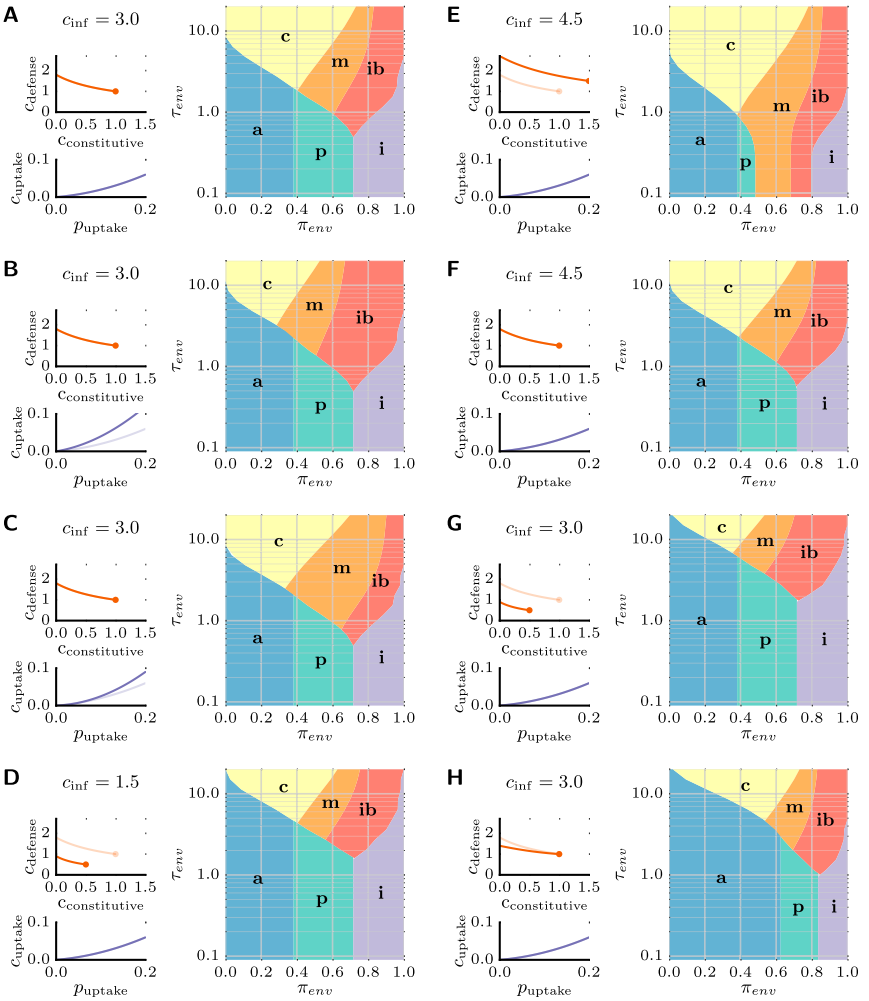

Fig. S3. Influence of parameter choice on the phase diagram presented in Fig. 2. In *A–H*, *Left* the parameter choices are shown and in *A–H*, *Right* the phase boundaries between adaptive (**a**), proto-adaptive (**p**), innate (**i**), innate bet hedging (**ib**), mixed (**m**), and CRISPR-like (**c**) strategies are shown. As a reference, lines in lighter color show trade-off and uptake cost for the parameter set used in Fig. 2. (*A*) Phase diagram for parameters used in Fig. 2. (*B*) More expensive active acquisition (c_{uptake} multiplied by a factor of 2). (*C*) Different functional form for cost of active acquisition: $c_{\mathrm{uptake}} = 0.05 \times p_{\mathrm{uptake}} + 2 \times p_{\mathrm{uptake}}^2$. (*D*) More permissive state-dependent costs (costs multiplied by a factor of 0.5). (*E*) Less permissive state-dependent costs (costs multiplied by a factor of 1.5). (*F*) Higher cost of infection. (*G*) Higher cost of immune protection. (*H*) Different functional form for cost trade-off, $c_{\mathrm{defense}} = 1.4 - 0.6 \times c_{\mathrm{constitutive}} + 0.2 \times c_{\mathrm{constitutive}}^2$.

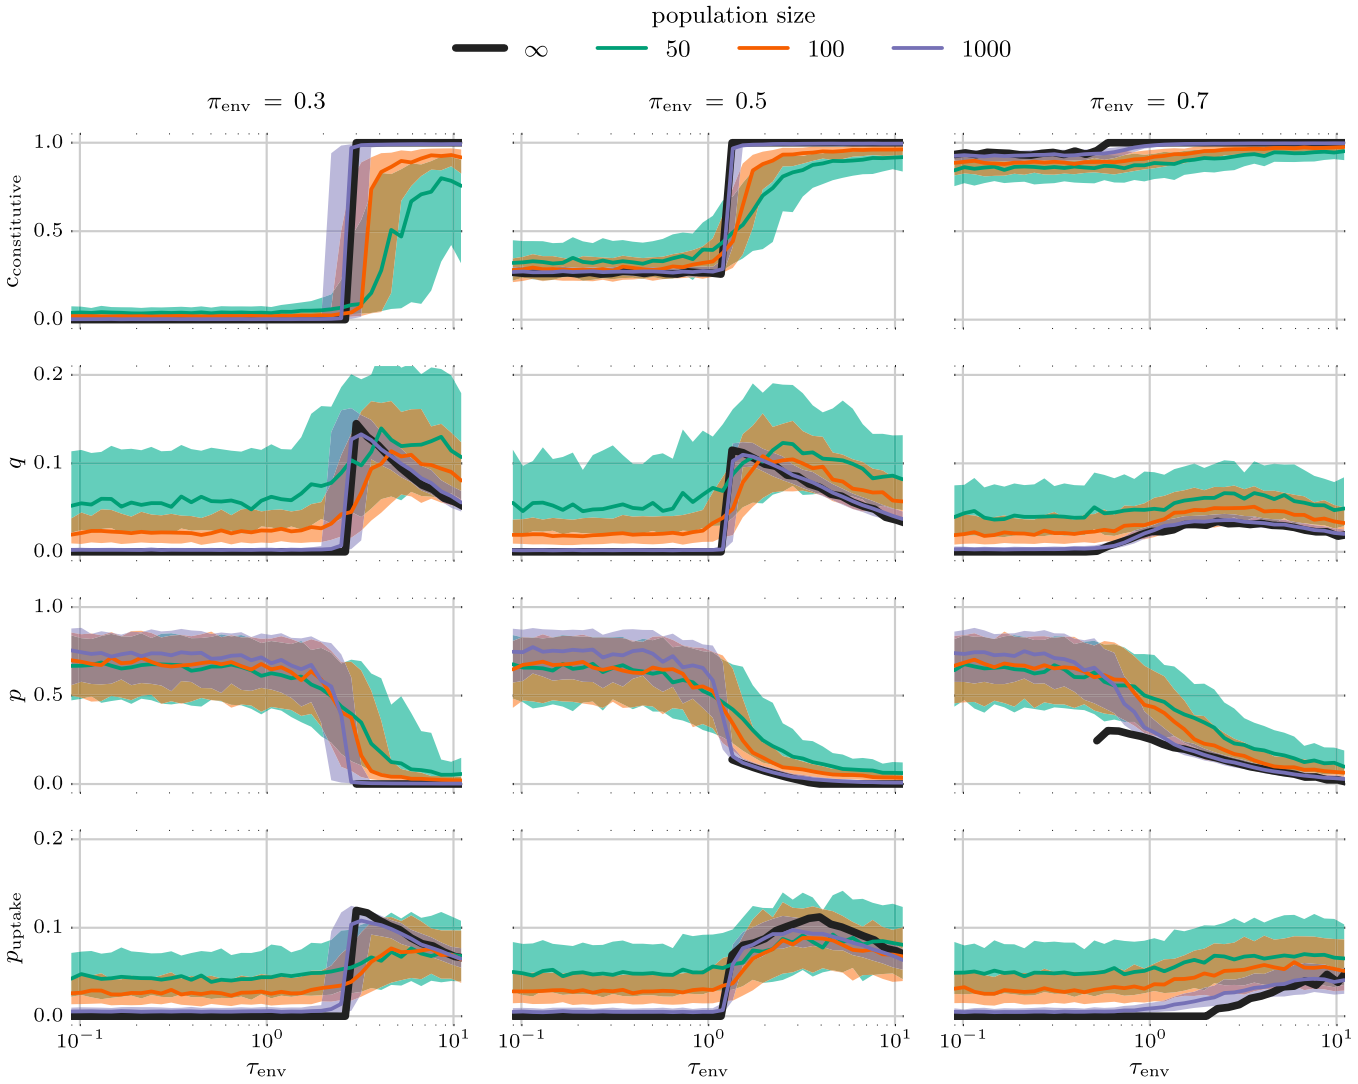

Fig. S4. Influence of finite population size on optimal immune strategies from an agent-based simulation with evolving strategy parameters (switching rates and degree of adaptability) as described in the text. For the infinite population, p is shown only for $q > 0$, because for $q = 0$ the value of p is not constrained other than being positive. Subplots show the median (solid line) and interquartile range (shaded area) of the strategy parameters at the end of a simulation of 100,000 generations length. Both are calculated from 500 independent simulations. In each simulation, the strategy parameters evolve from a random initial distribution via mutation and selection. Mutations take place with a rate $0.01 \exp(-t/10,000)$ per generation and are normally distributed with mean zero and SD $0.25 \exp(-t/10,000)$. The bound constraints on the parameters were enforced by setting the strategy parameters to the boundary value if outside after a mutation. Costs of different immune states are as in Fig. 2.