

Diversity of immune strategies explained by adaptation to pathogen statistics

Andreas Mayer^a, Thierry Mora^{b,1}, Olivier Rivoire^c, and Aleksandra M. Walczak^a

^aLaboratoire de Physique Théorique, CNRS, Université Pierre et Marie Curie, and École Normale Supérieure, 75005 Paris, France; ^bLaboratoire de Physique Statistique, CNRS, Université Pierre et Marie Curie, and École Normale Supérieure, 75005 Paris, France; and ^cLaboratoire Interdisciplinaire de Physique, CNRS and Université Grenoble Alpes, 38000 Grenoble, France

Edited by Peter G. Wolyne, Rice University, Houston, TX, and approved May 11, 2016 (received for review January 14, 2016)

Biological organisms have evolved a wide range of immune mechanisms to defend themselves against pathogens. Beyond molecular details, these mechanisms differ in how protection is acquired, processed, and passed on to subsequent generations—differences that may be essential to long-term survival. Here, we introduce a mathematical framework to compare the long-term adaptation of populations as a function of the pathogen dynamics that they experience and of the immune strategy that they adopt. We find that the two key determinants of an optimal immune strategy are the frequency and the characteristic timescale of the pathogens. Depending on these two parameters, our framework identifies distinct modes of immunity, including adaptive, innate, bet-hedging, and CRISPR-like immunities, which recapitulate the diversity of natural immune systems.

immune systems | CRISPR immunity | adaptive immunity | bet hedging | evolution of immunity

Immune systems have evolved to protect organisms against large and unpredictable pathogenic environments. However, immunity always comes at a cost (metabolic and maintenance costs, autoimmune disorders, etc. (1), and this cost must be balanced by the benefits that protection confers (2, 3). Faced with the problem of evolving a suitable defense, different organisms, from archaea to humans, have developed different strategies to identify and target pathogens, which have given rise to a diversity of mechanisms of immunity.

A large effort has been made to elucidate these mechanisms down to their molecular details in a variety of species (4–9). Beyond many differences, these studies have revealed many commonalities (10, 11), which hint at a possible general understanding of the trade-offs that shape their design (1, 2). For instance, independently of the well-known adaptive immune systems of jawed vertebrates, jawless vertebrates (e.g., lampreys) have developed an alternative adaptive system that uses a distinct molecular family of receptors, but both systems function largely in the same way, relying on the generation of a large number of diverse receptors expressed by two types of lymphocytes (B- or T-like cells). Likewise, the innate immune systems of invertebrates and vertebrates share many similarities, relying on the selected expression of germ-line Toll-like receptors upon infection. Some of the features of vertebrate immunity are even shared with bacteria, who have developed their own targeted immunity based on the CRISPR/Cas system (9, 12), which itself bears strong resemblance to genome protection through interfering RNAs in eukaryotes (13).

Independently of how they evolved and their particular molecular implementation, we may classify these diverse mechanisms into a few broad modes of immunity: heritable but not adaptable within an individual's lifetime, as innate immune systems; heritable and adaptable within a lifetime but with the benefits of adaptation being nonheritable, as adaptive immune systems; acquired from the environment and heritable, as the CRISPR/Cas system; and mixed strategies combining several of these elements. These broad distinctions call for general principles to characterize the conditions

under which one or another mode of immunity may be expected to evolve (1, 10, 11). The diversity and variability of threats from the pathogenic environment suggest that different modes of immunity may offer better protection, depending on the patterns of occurrence and virulence of pathogens or the effective population size of the protected population. Here we apply a general theoretical framework for analyzing populations in a varying environment (14) to predict the emergence of the basic forms of observed immunity.

Model

Individuals reproduce in the presence of pathogens, which randomly appear and may persist for several generations and disappear before possibly reappearing at a later time (Fig. 1A). In our framework, a given pathogen has a probability α to appear and a probability β to disappear from one generation to the next (Fig. 1B). The pathogenic dynamics are quantified both by the pathogen frequency $\pi_{\text{env}} = \alpha/(\alpha + \beta)$, which is the probability that it is present at any given generation, and by the characteristic timescale $\tau_{\text{env}} = -1/\ln(1 - \alpha - \beta)$, which sets how fast pathogens appear and disappear. Although other parameterizations may be considered, this choice for τ_{env} preserves the symmetry between the presence and absence of the pathogen (*SI Appendix A: Parameterizing a Two-State Markov Chain*).

Pathogens reduce the fitness of the individuals in the population and the immune system is designed to mitigate this effect. An individual's fecundity, defined as its expected number of descendants in the next generation ξ , depends on the pathogenic environment and its ability to protect itself against it. Each pathogen independently lowers the fecundity of an unprotected

Significance

Organisms possess many mechanisms for protecting themselves from pathogens. In addition to the well-studied innate and adaptive immune systems of vertebrates, recently discovered mechanisms such as CRISPR immunity in prokaryotes diversify the known modes of information processing used in immune protection. By classifying these different systems in terms of their rules of heritability and response to the environment, we propose a mathematical framework that recapitulates the observed natural diversity of immune systems. We show how the basic modes of immunity emerge as optimal strategies of a population adapting to a changing pathogenic environment. The proposed framework offers a unified view of immunity across species and helps rationalize the diversity of observed mechanisms.

Author contributions: A.M., T.M., O.R., and A.M.W. designed research; A.M., T.M., O.R., and A.M.W. performed research; A.M., T.M., O.R., and A.M.W. contributed new reagents/analytic tools; and A.M., T.M., O.R., and A.M.W. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: tmora@lps.ens.fr.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1600663113/-DCSupplemental.

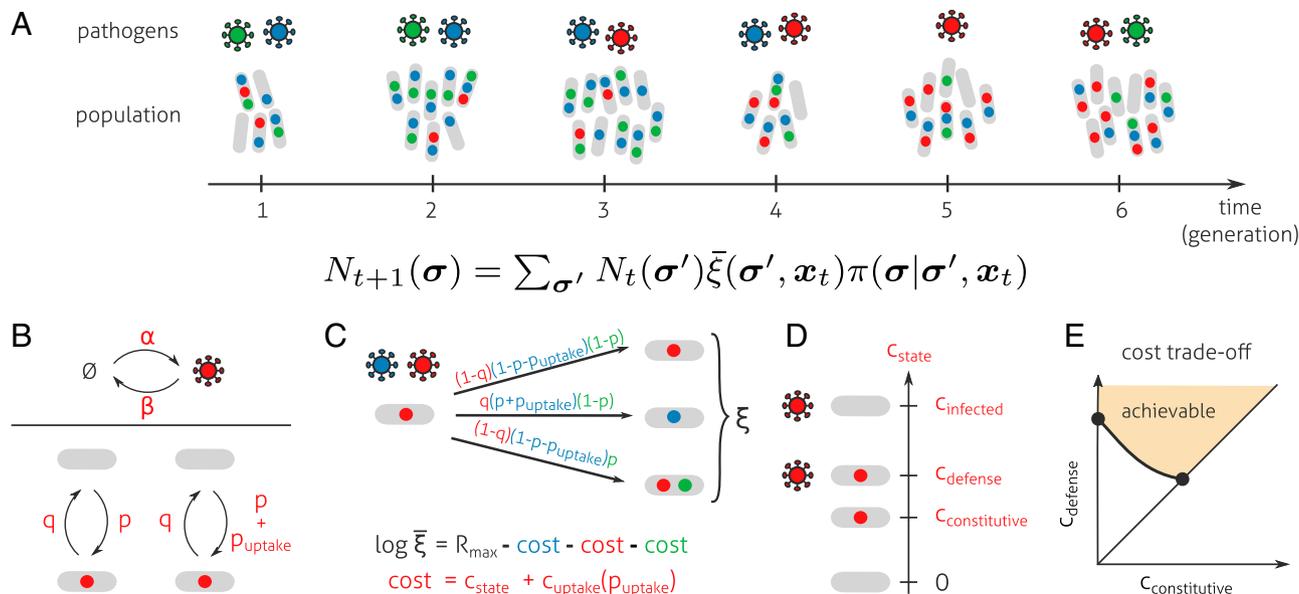


Fig. 1. A model to explore the incidence of different modes of immunity on the long-term growth of populations. (A) A population of organisms, each possibly protected against no, one, or several pathogens (no, one, or several colored circles), evolves in the presence of a pathogenic environment that varies from generation to generation. The mean number of individuals with protection σ at generation t , $N_t(\sigma)$, is given by a recursion equation involving the mean number of offspring $\bar{\xi}(\sigma', x_t)$ for individuals with protection σ' and the probability $\pi(\sigma | \sigma', x_t)$ that each of their offspring inherits a protection σ ; both of these quantities may depend on the current pathogenic environment x_t . The long-term growth rate of the population is given by $(1/t) \ln N_t$ at large t , with $N_t = \sum_{\sigma} N_t(\sigma)$ the total population size. (B) Dynamics of appearance and disappearance of pathogens x_t and immune protection σ . A pathogen appears with rate α and disappears with rate β ; these rates define the frequency $\pi_{\text{env}} = \alpha / (\alpha + \beta)$ and characteristic time $\tau_{\text{env}} = -1 / \ln(1 - \alpha - \beta)$ of the pathogen. Protection against a given pathogen is acquired spontaneously with rate p and lost from one generation to the next with rate q . Additionally, the presence of the pathogen can increase the rate of acquisition of protection by p_{uptake} , as, e.g., in the CRISPR-Cas system of prokaryotes. (C) The ξ offspring produced by an individual inherit the immune protections of their parent with rules specified in B. Each pathogen reduces the mean number of offspring $\bar{\xi}(\sigma', x_t)$ by a cost c_{state} that depends on whether the individual is in state “infected,” “defense,” or “constitutive” relative to the pathogen and by a cost $c_{\text{uptake}}(p_{\text{uptake}})$ that depends on the rate p_{uptake} at which protection is directly induced by the presence of the pathogen. (D) An unprotected organism pays a cost of infection c_{infected} if the pathogen is encountered, which is reduced to c_{defense} if it is protected. A protected organism must, however, pay a constitutive cost $c_{\text{constitutive}}$ even in the absence of the pathogen, whereas an unprotected organism pays no cost. (E) We assume a trade-off between the constitutive and defense costs: A more efficient defense (lower c_{defense}) requires more resources (higher $c_{\text{constitutive}}$).

individual by a relatively large cost factor $c_{\text{infected}} > 0$ (Fig. 1D). This cost is reduced to a lower cost $c_{\text{defense}} < c_{\text{infected}}$ when the individual is protected by its immune system; however, this protection comes at a basal cost $c_{\text{constitutive}} < c_{\text{defense}}$ of maintaining the immune defense in absence of the pathogen (Fig. 1E).

We explore the choices and trade-offs underlying various modes of immunity along three axes: adaptability, heritability, and mode of acquisition. The first one, adaptability axis, concerns how much resources are invested in the protection for the return of an efficient response. This trade-off imposes a relationship between c_{defense} and $c_{\text{constitutive}}$ (Fig. 1E): The more effective the defense is (the lower c_{defense}), the higher the maintenance cost (the higher $c_{\text{constitutive}}$). For example, having a large number of immune cells specialized against a specific pathogen allows for a quick and efficient response in the case of invasion, but this enhanced protection comes at the cost of producing and maintaining these cells in the absence of the pathogen. This strategy is adopted, for example, by much of the innate immune systems of plants and animals (4). On the contrary, the adaptive immune system keeps a very small specialized pool of lymphocytes for each potential antigen and makes them proliferate only in the case of infection (10). Having the machinery of adaptive immunity comes at some upfront investment cost, but the huge diversity of adaptive immune repertoires then allows for a response against many pathogens at essentially no marginal constitutive cost. It is this marginal constitutive cost that determines when to use a mode of defense once an organism has the machinery (*SI Appendix D: Nonindependent Pathogen-Protection Pairs*). We assume that the cost of defense grows faster

at small constitutive costs than at large ones (reflected in the convexity in Fig. 1E). The second one, heritability axis, is defined by the probability q that the protection is not transmitted to the offspring (Fig. 1B). Finally, the third one, acquisition axis, specifies how individuals may acquire the protection without inheriting it from their parent. This acquisition may occur randomly independently of the environment, with probability p , for instance by mutation or phenotypic switching, as is the case for antibiotic resistance in bacteria (15); or it can be induced by the presence of the pathogen with probability p_{uptake} , as in CRISPR-Cas immunity (Fig. 1B) (9). This mechanism comes at an extra cost $c_{\text{uptake}}(p_{\text{uptake}})$ due to maintenance and the risks of uptaking foreign genetic material (Fig. 1C), in addition to the state-dependent cost c_{state} (Fig. 1D). To account for the dangers associated with taking up foreign DNA we assume that c_{uptake} increases superlinearly with p_{uptake} .

Results

Each choice of the parameters $c_{\text{constitutive}}$, q , p , and p_{uptake} defines a specific immune strategy. This strategy is optimal if a population that adopts it outgrows in the long run any other population following a different strategy. Our goal is to characterize this optimal strategy, in particular its dependency on the two key properties of the pathogen, its frequency π_{env} and its characteristic time τ_{env} . We achieve this goal by maximizing the long-term growth rate of populations, defined by $(1/t) \ln N(t)$, where $N(t)$ is the total population size at generation t (Fig. 1A) (16). Conveniently, because the fecundity is affected independently by the different pathogens, each pathogen contributes additively to the

function of experienced infections, effectively implementing an adaptive memory within a single generation.

For slow and frequent pathogens (red phase), protection is acquired with probability $p > 0$ and lost with probability $q > 0$ independently of the presence of the pathogen. This bet-hedging strategy is implemented in bacteria that can switch on or off the expression of phage receptors (8). For slow but infrequent pathogens (yellow phase in Fig. 2A), a form of bet hedging is again present, but this time with a nonzero probability to acquire protection only in presence of the pathogen. An example of such a Lamarckian strategy is the CRISPR-Cas immune system in bacteria (9). Finally, a mixed phase (orange in Fig. 2A) is also possible where protection is randomly acquired at a rate that is increased by the presence of the pathogen.

We can gain insight into the transitions between the different phases by considering three analytically solvable simplifications of the model, detailed in *SI Appendix C: Analytical Insight into the Transitions Between Strategies*. In the first of these simplified models, we can calculate the transitions from a purely constitutive to a proto-adaptive to a purely adaptive strategy as the pathogen frequency π_{env} decreases. The second model highlights the transition from a bet-hedging to a deterministic protection, whereas the third one focuses on the transitions from a purely passive to a purely active acquisition of the protection, with a mixed phase in between.

It is instructive to examine how the parameters of immunity vary within the phases (Fig. 2C and D and Fig. S1). As one may expect, the statistical properties of the protection tend to track the pathogen statistics (18). The more frequent the pathogen is, the more prevalent the protection in the population (Fig. 2C). Likewise, the characteristic time of the protection, τ , grows with that of the pathogen, τ_{env} (Fig. 2E).

The assumptions that we made allow us to treat each pathogen–protection pair independently of each other. However, there are a number of ways in which this assumption may be questioned. We discuss several in *SI Appendix D: Nonindependent Pathogen–Protection Pairs* and we find that these generalizations do not qualitatively affect our conclusions. For example, infections could interact by inflicting more harm together than the sum of each one alone, e.g., HIV in conjunction with other diseases (Fig. S2). This case can be incorporated into our approach by considering a modified effective cost including the extra cost of coinfection. Another way in which pathogen–protection pairs could be correlated is through a nonadditive cost of protection, if the marginal cost of protection increases or decreases with the number of

protections. For example, if having protection against two pathogens is much more costly than twice the cost of having protection against just one, then the optimal strategy may be to hedge bets by keeping a subpopulation protected against one pathogen and another subpopulation protected against the other. Finally, cross-reactivity, the widespread ability of protections to recognize several pathogens, is another departure from independence, which can be partly overcome by grouping together pathogens recognized by a common protection.

Discussion

The phase portrait in Fig. 2A rationalizes the salient differences between the immune systems of prokaryotes and vertebrates. Bacterial and archaeal organisms evolve on timescales that are much closer to those of their pathogens than vertebrates. From the viewpoint of microbes, the pathogenic environment is relatively constant ($\tau_{\text{env}} > 1$), whereas for vertebrates a particular pathogenic strain is unlikely to survive a single generation ($\tau_{\text{env}} \ll 1$). Consistent with our results, vertebrates use fully heritable modes of immunity and do not rely on bet hedging. To deal with infrequent and fast-evolving pathogens such as viruses, they recourse to adaptive mechanisms by which they can up-regulate their protection in case of an invasion. The three predicted strategies—adaptive, proto-adaptive, and innate—correspond to the known modes of immunity in vertebrates (19). Prokaryotes, on the other hand, almost systematically use bet-hedging strategies. They recourse both to the CRISPR-Cas system of acquired immunity (9) and to innate immunity through, e.g., restriction endonucleases (8), which correspond to the predicted Lamarckian and innate bet-hedging strategies of the diagram, respectively. These results are robust to changes of parameters, although increasing costs can make bet hedging beneficial even at short characteristic times (Fig. S3).

Bacteria and vertebrates also have very different population sizes, which influence their overall survival probability. To evaluate this impact, we ran stochastic simulations, competing different strategies for increasing population sizes (*Materials and Methods* and Fig. S4). The phase diagram in Fig. 2A is recovered for population sizes as small as 1,000, whereas for smaller population sizes the boundaries between regimes are smeared. In small populations, adaptive strategies are generally favored over CRISPR-like strategies, and the amount of bet hedging increases. In fact, for finite populations it is always beneficial to recourse to some degree of bet hedging to react quickly to environmental changes and avoid extinction.

Table 1. Optimal strategies found in the phase diagram, their definition in terms of parameters of our framework, and biological examples

Strategy	Defining characteristics				Biological examples
	Perfect heritability, $q=0$	Acquisition mode		Adaptability, c_{defense}	
		$p > 0$	$p_{\text{uptake}} > 0$		
Innate	Yes	Yes	No	Minimal	Innate defense by recognition of pathogen-associated molecular patterns by pattern recognition receptors (4)
Protoadaptive	Yes	Yes	No	Intermediate	“Trained” innate immunity (26), especially defense by natural killer cells (17); “systemic acquired resistance” in plants (20)
Adaptive	Yes	Yes	No	Maximal	Adaptive immune systems of jawed and jawless vertebrates (10)
Innate bet hedging	No	Yes	No	Minimal	Mutation of phage receptors by bacteria (8)
CRISPR-like	No	No	Yes	Minimal	CRISPR-Cas system in bacteria and archaea (9)
Mixed	No	Yes	Yes	Minimal	Concurrent use of CRISPR-Cas system and mutations of surface molecules by bacteria defending against phages (27)

Here, we consider the case of a common environment experienced by all individuals. Having different parts of the population experience different microenvironments that are not persistent over generations does not change the population dynamics on evolutionary timescales (*Materials and Methods*). These microenvironments can result from differing infection probabilities for different subsets of the population, e.g., arising from spatial niches, or other nonpathogenic factors such as nutrient availability that influence the capability of individuals to cope with pathogens. If there are microenvironments that persist over many generations, then our results hold in each microenvironment. An optimal strategy might then exploit the additional predictability stemming from knowing the statistical properties of the microenvironments and use the microenvironment diversity as a means of bet hedging.

Our results also suggest that plants and some invertebrates, which also have long generation times compared with the variation time of pathogens, should be endowed with adaptive and proto-adaptive immune systems, in addition to innate protection mechanisms (6). Consistent with this prediction, the innate branch of the plant immune system is able to increase protection in the entire plant following a local infection through “systemic acquired resistance” (20), providing the mechanistic basis of an inducible, proto-adaptive immune system. In addition, virus-derived small interfering RNAs, which accumulate during infections, are portrayed as likely candidates of adaptive immunity in plants and invertebrates—they are induced by the virus and keep a memory of past infections (21–23). Interestingly, small RNA-based immunity has been shown to be inheritable in *Caenorhabditis elegans* (24), an invertebrate with a short generation time of around 4 days, in agreement with our result that CRISPR-like immunity is desirable in this case.

By analyzing the long-term fate of populations under minimal assumptions concerning the rules governing adaptability, heritability, and acquisition of immune protections, we have recovered the basic known modes of immunity. Remarkably our results hold even for a single pathogen. The key determinants of optimal immune strategies are found to be the statistical features of pathogen occurrence: its frequency and its characteristic timescale. As an implication, a diverse pathogenic environment, with varying statistics, will favor mixed solutions, consistent with the observation of multiple immune systems within the same organism—such as adaptive and innate immune systems in vertebrates or CRISPR and innate defense in bacteria. Naturally, the molecular implementation of these general principles differs greatly even between organisms sharing the same type of immunity. However, an evolutionary perspective that accounts for the costs and benefits of protection is enough to explain the most salient features of immunity. It will be interesting to extend our framework to account for other essential features of immunity, e.g., the acquisition of protection by horizontal transfer or the evolutionary dynamics between pathogens and their hosts. In view of our analysis, it is already less surprising that complex forms of immunity such as the adaptive immune system have evolved separately in jawed and nonjawed vertebrates, with the same general features but different molecular encodings.

Materials and Methods

Population Dynamics. The pathogenic environment is described by an L -dimensional vector \mathbf{x} (symbols in boldface type refer to vectors), where $x_i = 1$ if pathogen i is present and 0 otherwise. Protection of an organism against these pathogens is also described by an L -dimensional vector σ , where $\sigma_i = 1$ if the protection (antibody, T-cell receptor, CRISPR spacer) against pathogen i is present and 0 otherwise.

We consider the dynamics of a population of organisms reproducing at discrete times t . At each generation, each individual produces a stochastic number ξ of offspring, whose distribution depends on the state σ of that individual and the environment \mathbf{x}_t . We denote its mean by $\bar{\xi}(\sigma, \mathbf{x}_t)$. Let $N_t(\sigma|\mathbf{x}_{t<t+1})$ be the mean number of organisms in the population at time t with protection σ , for a given environment history ($\mathbf{x}_{t<t+1}$) (16). The change

in population composition from one generation to the next is governed by the reproductive success of individuals in each state σ , modified by stochastic state switching from parents to offspring

$$N_{t+1}(\sigma|\mathbf{x}_{t<t+1}) = \sum_{\sigma'} N_t(\sigma'|\mathbf{x}_{t<t}) \bar{\xi}(\sigma', \mathbf{x}_t) \pi(\sigma|\sigma', \mathbf{x}_t), \quad [1]$$

where $\pi(\sigma|\sigma', \mathbf{x}_t)$ is the switching probability from protection state σ' to state σ . Note that the protection state switching probability, which represents to what extent protection is inherited, acquired, or lost, generally depends on the state \mathbf{x}_t of the environment. For ease of notation, we omit in the following the condition on the environment ($\cdot|\mathbf{x}_{t<t}$) when referring to conditional means.

A similar recursion to Eq. 1 can be written for the fraction of the population in each state, $n_t(\sigma) = N_t(\sigma)/N_t$, with $N_t = \sum_{\sigma} N_t(\sigma)$ the total population size

$$n_{t+1}(\sigma) = \frac{1}{Z_t} \sum_{\sigma'} n_t(\sigma') \bar{\xi}(\sigma', \mathbf{x}_t) \pi(\sigma|\sigma', \mathbf{x}_t), \quad [2]$$

where Z_t is a normalization constant enforcing $\sum_{\sigma} n_t(\sigma) = 1$. The population size is given by $N_t = N_0 \prod_{t=0}^{t-1} Z_t$, so that the long-term growth rate, $\Lambda = \lim_{T \rightarrow \infty} (1/T) \ln N_T$, is given by

$$\Lambda = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \ln(Z_t). \quad [3]$$

The strategy with maximal long-term population growth rate outperforms in the long run any other strategy for almost every sequence of environments in populations of infinite size. This rate thus provides a measure of long-term fitness (16).

We assume that the mutation and inheritance probabilities of different pathogen–protection pairs are independent of each other, i.e., that $\pi(\sigma|\sigma', \mathbf{x}_t)$ factorizes over the pathogens

$$\pi(\sigma|\sigma', \mathbf{x}_t) = \prod_i \pi_i(\sigma_i|\sigma'_i, x_{it}). \quad [4]$$

The entries of $\pi_i(\sigma_i|\sigma'_i, x_{it})$ are given in Fig. 1B: $\pi_i(1|0, x) = \rho + x\rho_{\text{uptake}}$ and $\pi_i(0|1, x) = q$.

In addition, the effects of different pathogen–protection pairs on the growth rate are taken to be additive (Fig. 1C), so that

$$\ln \bar{\xi} = R_{\max} - \sum_{i=1}^L [c_{\text{infection},i} (1 - \sigma_i) x_i + c_{\text{constitutive},i} \sigma_i (1 - x_i) + c_{\text{defense},i} \sigma_i x_i + c_{\text{uptake}} (\rho_{\text{uptake},i})], \quad [5]$$

where R_{\max} is the growth rate in the absence of any immune cost. With these assumptions, the distribution $n_t(\sigma)$ also factorizes over i

$$n_t(\sigma) = \prod_{i=1}^L [r_i^t \sigma_i + (1 - r_i^t) (1 - \sigma_i)], \quad [6]$$

where r_i^t is the fraction of the population having protection i at time t . Plugging this ansatz into Eq. 2 with Eqs. 4 and 5 yields the following recursion for r_i^t :

$$r_i^{t+1} = \frac{\left[(1 - r_i^t) e^{-c_{\text{infection},i} x_i^t} (\rho_i + \rho_{\text{uptake},i} x_i^t) + r_i^t e^{-c_{\text{defense},i} x_i^t - c_{\text{constitutive},i} (1 - x_i^t)} (1 - q_i) \right]}{\left[(1 - r_i^t) e^{-c_{\text{infection},i} x_i^t} + r_i^t e^{-c_{\text{defense},i} x_i^t - c_{\text{constitutive},i} (1 - x_i^t)} \right]}. \quad [7]$$

The recursion depends on the sequence of x_i^t , which is a stochastic binary process switching from 0 to 1 with probability α and from 1 to 0 with probability β (Fig. 1B). Note that the sequence x_i^t is the same for the whole population (a quenched variable in the statistical mechanics sense). We have $Z_t = e^{R_{\max}} \prod_{i=1}^L z_i^t$, with

$$z_i^t = e^{-c_{\text{uptake}} (\rho_{\text{uptake},i})} \left[(1 - r_i^t) e^{-c_{\text{infection},i} x_i^t} + r_i^t e^{-c_{\text{defense},i} x_i^t - c_{\text{constitutive},i} (1 - x_i^t)} \right]. \quad [8]$$

From Eq. 3, it then follows that

$$\Lambda = R_{\max} + \sum_{i=1}^L \left(\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \ln z_i^t \right). \quad [9]$$

The long-term growth rate is a sum of independent terms for each pathogen-protection pair, which allows us to treat the problem of maximizing the long-term growth rate one pathogen at a time.

Numerical Solution. The cost function of the optimization, Λ , can be approximated by solving the recursion equation describing the relative frequency of organisms with different protection states in the population for a large enough number of generations (we used at least 10^6 generations). Our goal is to optimize Λ over the four parameters $p, q, \rho_{\text{uptake}}, c_{\text{constitutive}}$ (or over the subset of free parameters for a given strategy) constrained to their domain of definition. For numerical purposes, all four parameters are first mapped onto the unit interval $[0, 1]$. The noise in the evaluation of Λ arising from its approximation from finite time data makes the optimization challenging. Because the process is ergodic, averaging over very long periods is equivalent to repeating the process multiple times. The noise can therefore be reduced by both prolonged simulation or repeated sampling at the expense of a higher computational cost per function evaluation. To find the global optimum of this noisy function under the bound constraints on the parameters we use a two-phase algorithm. In the first phase the DIRECT algorithm (25) provides us with a rough, but global optimization for which we use a relatively low-quality approximation. The results of this first phase are then refined by a pattern-search algorithm with an adaptive sampling of the function (described in detail in *SI Appendix B: Pattern-Search-Based Optimization for Problems with Noisy Function Evaluations*), using the parameters $\Delta_{\text{tol}} = 0.0005$ (for Fig. S1 $\Delta_{\text{tol}} = 0.005$), $\alpha = 0.005$.

To obtain a phase diagram such as the one shown in Fig. 2A we first performed a global optimization over all four parameter values as described above, for every environment condition $(\pi_{\text{env}}, \tau_{\text{env}})$ (Fig. S1). Based on the optimal parameters found in this step, we defined the features of the emerging phases (Table 1). All phases are defined by a subset of the variables lying at a constraint boundary. To calculate precise phase boundaries, we find the frequency of pathogens π_{env} at a given characteristic time τ_{env} for which the difference in long-term growth rates between a given pair of strategies vanishes. To obtain the root of the difference function, we use a bisection algorithm. To decrease noise, the difference is calculated across

pairs of simulations using the same sequence of pathogens $\{x_t\}$ and the function is sampled adaptively to ascertain statistical significance. The bisection algorithm is run up to a tolerance of 0.025 in π_{env} and then the precise position of the root is interpolated assuming linearity of the difference function within the interval. To prevent, e.g., the mixed strategy from reducing to a CRISPR-like strategy, we impose that the parameters that are not set to a fixed value in a particular strategy are not closer than a tolerance 0.005 (0.0005 for q) of the boundary.

All source code associated with this manuscript is available online at <https://github.com/andim/evolimmune>.

Simulations with Finite Populations Sizes. To study the influence of the effects of finite population size we perform direct agent-based simulations of a population of adapting individuals with strategies evolving on a slow timescale. The population has a finite size N that remains fixed over the course of the simulation. At every generation the parents of the N individuals are drawn from the individuals making up the previous generation with probabilities proportional to the mean number of offspring $\bar{\xi}$ of these individuals. The offspring's state σ is determined from the state of its parent σ' according to the switching rates $\pi(\sigma|\sigma', x_t)$ defined previously. Along with the state σ , the switching rates themselves, $\pi(\sigma|\sigma', x)$, as well as the degree of adaptability, $c_{\text{constitutive}}$ —in other words, the parameters defining the immune strategy—are also transmitted to the offspring. They also change from parent to offspring, although at a much slower rate than the state to preserve a clear separation of timescales between short-term and long-term adaptations. In this setup, selection acts on the strategies. After an equilibration phase, we collect statistics on the strategies adopted by individuals in the population. To get rid of the effect of deleterious mutations that do not eventually fix in the populations the mutation rate and size were scaled down exponentially with time. As population size is finite deleterious mutations can fix in the population, which means that even in the limit of zero mutation rate there remains a spread in the distribution of strategies. Hence we represent not only the median as a measure of the central tendency of a parameter, but also the interquartile range as a measure of its spread.

ACKNOWLEDGMENTS. This work was supported by European Research Council Starting Grant 306312.

- Schmid-Hempel P (2005) Evolutionary ecology of insect immune defenses. *Annu Rev Entomol* 50:529–551.
- Lochmiller RL, Deerenberg C (2000) Trade-offs in evolutionary immunology: Just what is the cost of immunity? *Oikos* 88(1):87–98.
- McKean KA, Yourth CP, Lazzaro BP, Clark AG (2008) The evolutionary costs of immunological maintenance and deployment. *BMC Evol Biol* 8:76.
- Medzhitov R, Janeway CA, Jr (1997) Innate immunity: The virtues of a nonclonal system of recognition. *Cell* 91(3):295–298.
- Magnadóttir B (2006) Innate immunity of fish (overview). *Fish Shellfish Immunol* 20(2):137–151.
- Jones JD, Dangl JL (2006) The plant immune system. *Nature* 444(7117):323–329.
- Strand MR (2008) The insect cellular immune response. *Insect Sci* 15(1):1–14.
- Labrie SJ, Samson JE, Moineau S (2010) Bacteriophage resistance mechanisms. *Nat Rev Microbiol* 8(5):317–327.
- Marraffini LA, Sontheimer EJ (2010) CRISPR interference: RNA-directed adaptive immunity in bacteria and archaea. *Nat Rev Genet* 11(3):181–190.
- Boehm T (2011) Design principles of adaptive immune systems. *Nat Rev Immunol* 11(5):307–317.
- Little TJ, Hultmark D, Read AF (2005) Invertebrate immunity and the limits of mechanistic immunology. *Nat Immunol* 6(7):651–654.
- Horvath P, Barrangou R (2010) CRISPR/Cas, the immune system of bacteria and archaea. *Science* 327(5962):167–170.
- Malone CD, Hannon GJ (2009) Small RNAs as guardians of the genome. *Cell* 136(4):656–668.
- Rivoire O, Leibler S (2014) A model for the generation and transmission of variations in evolution. *Proc Natl Acad Sci USA* 111(19):E1940–E1949.
- Gniadkowski M (2008) Evolution of extended-spectrum β -lactamases by mutation. *Clin Microbiol Infect* 14(Suppl 1):11–32.
- Rivoire O, Leibler S (2011) The value of information for populations in varying environments. *J Stat Phys* 142:1124–1166.
- Vivier E, et al. (2011) Innate or adaptive immunity? The example of natural killer cells. *Science* 331(6013):44–49.
- Mayer A, Balasubramanian V, Mora T, Walczak AM (2015) How a well-adapted immune system is organized. *Proc Natl Acad Sci USA* 112(19):5950–5955.
- Murphy K, Travers P, Walport M (2001) *Janeway's Immunobiology* (Garland Science, New York), 7th Ed, Vol 2.
- Spoel SH, Dong X (2012) How do plants achieve immunity? Defence without specialized immune cells. *Nat Rev Immunol* 12(2):89–100.
- Voinnet O (2001) RNA silencing as a plant immune system against viruses. *Trends Genet* 17(8):449–459.
- Waterhouse PM, Wang MB, Lough T (2001) Gene silencing as an adaptive defence against viruses. *Nature* 411(6839):834–842.
- Aliyari R, et al. (2008) Mechanism of induction and suppression of antiviral immunity directed by virus-derived small RNAs in *Drosophila*. *Cell Host Microbe* 4(4):387–397.
- Rechavi O, Minevich G, Hobert O (2011) Transgenerational inheritance of an acquired small RNA-based antiviral response in *C. elegans*. *Cell* 147(6):1248–1256.
- Jones DR, Perttunen CD, Stuckman BE (1993) Lipschitzian optimization without the Lipschitz constant. *J Optim Theory Appl* 79(1):157–181.
- Netea MG, et al. (2016) Trained immunity: A program of innate immune memory in health and disease. *Science* 352(6284):aaf1098.
- Westra ER, et al. (2015) Parasite exposure drives selective evolution of constitutive versus inducible defense. *Curr Biol* 25(8):1043–1049.
- Kolda TG, Lewis RM, Torczon V (2003) Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Rev* 45:385–482.
- Anderson EJ, Ferris MC (2001) A direct search algorithm for optimization with noisy function evaluations. *SIAM J Optim* 11:837–857.
- Deng G, Ferris MC (2009) Variable-number sample-path optimization. *Math Program* 117(1):81–109.
- Spall JC (1998) Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Trans Aerosp Electron Syst* on 34:817–823.
- Boyd SP, Vandenberghe L (2004) *Convex Optimization* (Cambridge Univ Press, Cambridge, UK).
- Donaldson-Matasci MC, Lachmann M, Bergstrom CT (2008) Phenotypic diversity as an adaptation to environmental uncertainty. *Evol Ecol Res* 10:493–515.
- Kussell E, Leibler S (2005) Phenotypic diversity, population growth, and information in fluctuating environments. *Science* 309(5743):2075–2078.

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

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

which is not known explicitly, but needs to be approximated by evaluating a function $F(y, \omega)$ 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 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 $f \leftrightarrow -\Lambda$, $y \leftrightarrow (p, q, p_{\text{uptake}}, c_{\text{constitutive}})$, $\omega \leftrightarrow x_t$.

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 $\mathcal{D} = \{\pm e_i | i = 1, \dots, n\}$, where e_i is the i th unit vector; and let us further define $P_{\Omega}(y) = \arg\min_{y' \in \Omega} |y' - y|^2$ as the projection of a point 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 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 $y \leftarrow y^0$, step size $\Delta \leftarrow \Delta^0$, and number of samples $N \leftarrow N^0$.
- 2) **While** ($\Delta \geq \Delta_{\text{tol}}$) or (y or N updated during last iteration):
 - a) **For each** step Δd along positive and negative coordinate directions $d \in \mathcal{D}$:
 - i) **If** $f(y') < f(y)$ (as judged from N samples of F at both points), where $y' = P_{\Omega}(y + \Delta d)$, then update the parameter vector $y \leftarrow y'$.
 - ii) **Else if** new point $y + \Delta d$ is feasible, i.e., $y + \Delta d \in \Omega$, and if $f(y + \Delta d) = f(y)$ 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** $f(y + (\Delta/2)d) < f(y)$ (as judged from N samples of F at both points), then update parameter vector and reduce step size:
$$y \leftarrow y + \frac{\Delta}{2}d$$

$$\Delta \leftarrow \Delta/2.$$
 - B) **Else** increase sampling:
$$N \leftarrow 2N.$$
 - b) **If** no updates during preceding loop, then contract pattern size:
$$\Delta \leftarrow \Delta/2.$$

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

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

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 $F(y, \omega_i), F(y', \omega_i)$ for $\omega_i, i = 1, \dots, N$ 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 $c_{\text{con}} \in [0, c_{\text{con}}^{\text{max}}]$ 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:

$$\frac{\pi_{\text{env}} < \pi_{\text{env}}^{(ap)}}{c_{\text{con}} = 0} \quad \left| \quad \frac{\pi_{\text{env}}^{(ap)} \leq \pi_{\text{env}} \leq \pi_{\text{env}}^{(po)}}{0 \leq c_{\text{con}} \leq c_{\text{con}}^{\text{max}}} \quad \left| \quad \frac{\pi_{\text{env}}^{(po)} < \pi_{\text{env}}}{c_{\text{con}} = c_{\text{con}}^{\text{max}}}$$

with

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

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

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

$$\frac{\pi_{\text{env}} < \pi_{\text{env}}^{(oi)}}{p = 0} \quad \left| \quad \frac{\pi_{\text{env}}^{(oi)} \leq \pi_{\text{env}} \leq \pi_{\text{env}}^{(io)}}{0 \leq p \leq 1} \quad \left| \quad \frac{\pi_{\text{env}}^{(io)} < \pi_{\text{env}}}{p = 1}$$

with

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

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

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

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

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:

$$\frac{\pi_{\text{env}} < \pi_{\text{env}}^{(cm)}}{p = 0, p_{\text{uptake}} > 0} \quad \left| \quad \frac{\pi_{\text{env}}^{(cm)} \leq \pi_{\text{env}} \leq \pi_{\text{env}}^{(mi)}}{p > 0, p_{\text{uptake}} > 0} \quad \left| \quad \frac{\pi_{\text{env}}^{(mi)} < \pi_{\text{env}}}{p > 0, p_{\text{uptake}} = 0}$$

with

$$Q_{00}(\pi_{\text{env}}^{(cm)}) = \frac{Q_{01}(\pi_{\text{env}}^{(cm)})}{g^{-1} [Q_{01}(\pi_{\text{env}}^{(cm)})]}, \quad \text{with } g(p_{\text{uptake}}) = p_{\text{uptake}} \frac{dc_{\text{uptake}}}{dp_{\text{uptake}}}, \quad [\text{S10}]$$

$$\pi_{\text{env}}^{(mi)} = 1 - \frac{dc_{\text{uptake}}}{dp_{\text{uptake}}} \Big|_{p_{\text{uptake}}=0}. \quad [\text{S11}]$$

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. S3F). 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. S3G). 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 non-additive 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 σ . 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,

$\sigma \setminus (x_1, x_2)$	(0, 0)	(1, 0)	(0, 1)	(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

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

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

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

with $\pi_{\text{env},\text{eff}} = \pi_{\text{env},1} + \pi_{\text{env},2} - \pi_{\text{env},1}\pi_{\text{env},2}$. The last term in this expression is small either for infrequent pathogens ($\pi_{\text{env},1}\pi_{\text{env},2} \ll \pi_{\text{env},\text{eff}}$) 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 $\pi_{\text{env},\text{eff}}$, in agreement with our expectation.

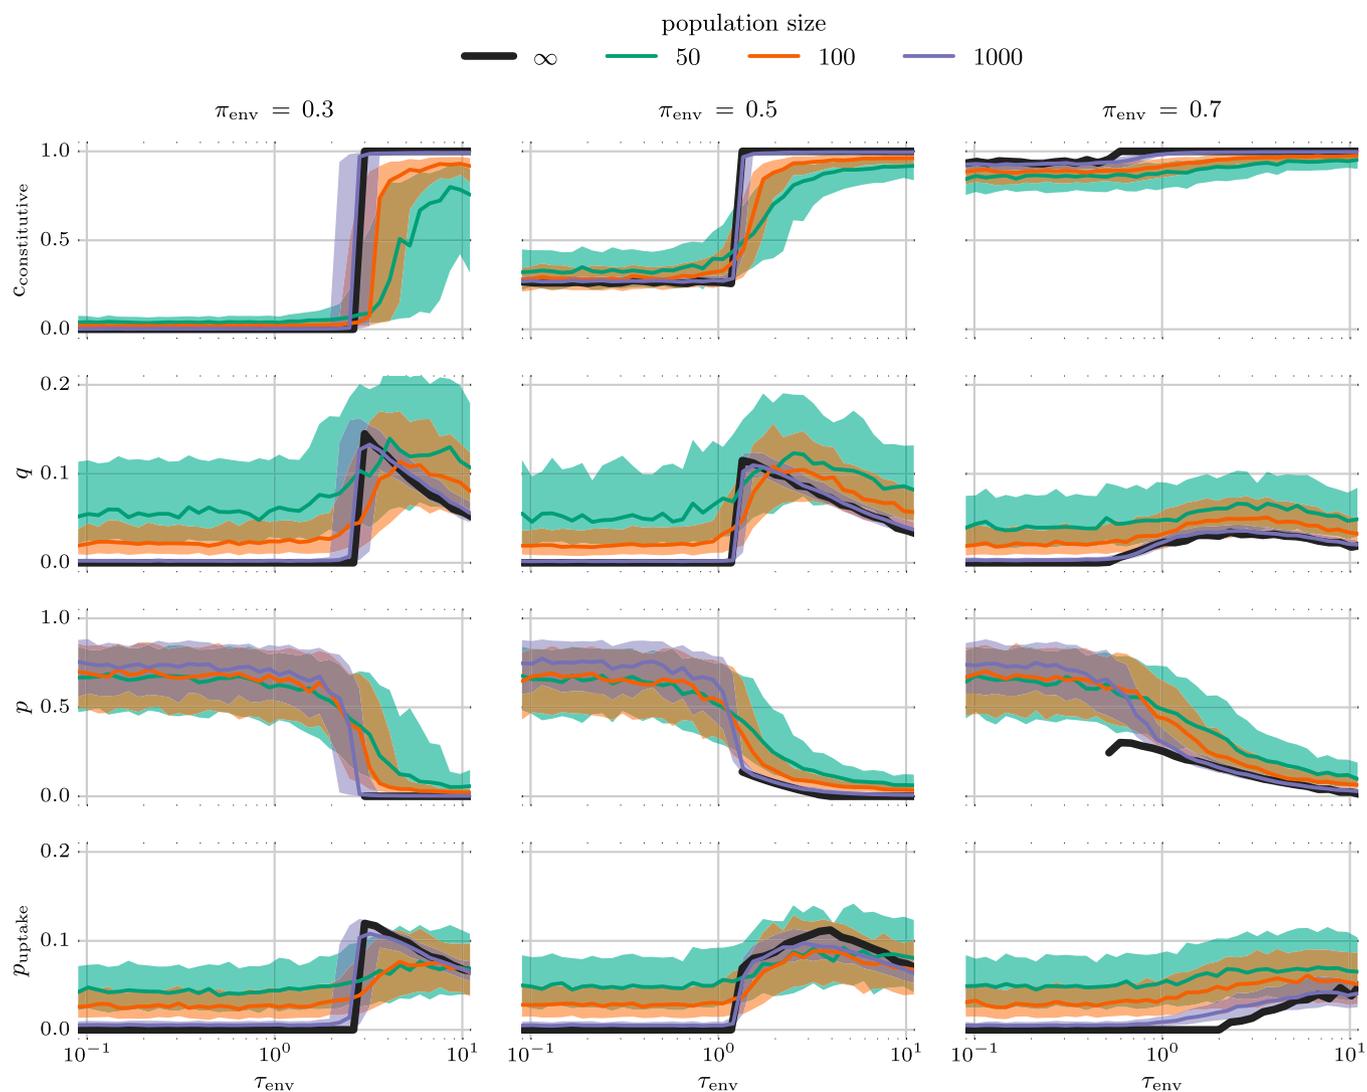


Fig. 54. Influence of finite population size on optimal immune strategies from an agent-based simulation with evolving strategy parameters (switching rates and degree of adaptability) as described in the text. For the infinite population, p is shown only for $q > 0$, because for $q = 0$ the value of 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.