# Eco-evolutionary control of pathogens

**Michael Lässig[a,1,2] and Ville Mustonen[b,1,2]**

[a]Institut für Biologische Physik, Universität zu Köln, 50937 Köln, Germany; and [b]Organismal and Evolutionary Biology Research Programme, Department of Computer Science, Institute of Biotechnology, Helsinki Institute for Information Technology, University of Helsinki, 00014 Helsinki, Finland

Control can alter the eco-evolutionary dynamics of a target pathogen in two ways, by changing its population size and by directed evolution of new functions. Here, we develop a payoff model of eco-evolutionary control based on strategies of evolution, regulation, and computational forecasting. We apply this model to pathogen control by molecular antibody–antigen binding with a tunable dosage of antibodies. By analytical solution, we obtain optimal dosage protocols and establish a phase diagram with an error threshold delineating parameter regimes of successful and compromised control. The solution identifies few independently measurable fitness parameters that predict the outcome of control. Our analysis shows how optimal control strategies depend on mutation rate and population size of the pathogen, and how monitoring and computational forecasting affect protocols and efficiency of control. We argue that these results carry over to more general systems and are elements of an emerging eco-evolutionary control theory.

biophysics | population genetics | immune systems | control theory

Control of human pathogens is a central goal of medicine. Important examples are antimicrobial and antiviral therapies and vaccinations; similarly, cancer therapies aim to control tumor cell populations. Biological hosts, notably the human immune system, face related issues of pathogen control. In most cases, control targets pathogen populations with fast-paced replication and evolution. Its goal is to alter these dynamics: to prevent or elicit an evolutionary process of the pathogen or to curb the pathogen population by reducing its ecological niche. Pathogen control has seen spectacular successes (e.g., in the eradication of smallpox and in HIV combination therapies) (1). However, control is often compromised by escape evolution of the pathogen, highlighting the importance to factor pathogen evolution into control protocols (2, 3). Promising evolutionary avenues include adaptive pathogen control and cancer therapy (4–6), vaccination, drug development and immunotherapy strategies based on evolutionary predictions (7–10), and controlled evolution of immune antibodies (11–13). However, we need quantitative relations between leverage and cost of control in order to generate optimization criteria and protocols that are comparable across systems. These are central elements of an eco-evolutionary control theory.

Because population dynamics and evolution are stochastic processes, any eco-evolutionary control operates on the likelihood of future states. Successful control turns a likely process into an unlikely one (e.g., the evolution of antibiotic resistance) or vice versa (e.g., the evolution of a broadly neutralizing antibody). In a broader scientific context, directing a stochastic process toward a future objective is a classic subject of control theory (14, 15). There is a well-established conceptual and computational framework to optimize control protocols, given complete knowledge of the dynamical rules and the ability to forecast likely future outcomes. However, the swords of eco-evolutionary control are blunter, and establishing an appropriate control theory faces new challenges. First, the control of an evolving population is based, at best, on limited dynamical information and forecasting capabilities. Here, we compare three modes of control update dynamics: by Darwinian evolution of a biotic host system, by reg-

ulation (which requires sensing of the current pathogen state as input), and by computation (which requires sensing and forecasting). For human interventions, optimizing control is inextricably linked to predictive evolutionary analysis, which is a topic of high current interest but far from a comprehensive understanding (16). Second, control theory has to factor in the underlying biological mechanism of control. Host–pathogen interactions are often based on biomolecular interactions, such as drug–target or antibody–antigen binding (17). The form of these interactions imposes specific constraints on control forces and their leverage on the pathogen system, which are discussed below. Third, developing an appropriate dynamical model of control update and pathogen response calls for a merger of control theory with ecological dynamics and population genetics. These questions are the topic of the present paper.

In the first part, we develop dynamical principles of eco-evolutionary control. The evolution and population dynamics of the pathogen are governed by intrinsic forces, including fitness and entropy of pathogen traits, and by the additional selective force imposed by control. We derive general minimum-leverage relations that specify the strength of control needed to alter the evolution of the pathogen toward the host's control objective. The control force has a payoff function in the host system, which includes the pathogen load and the cost of control, and is updated in response to the pathogen. For control by evolution or regulation, host and pathogen follow similar dynamical rules.

## Significance

Vaccinations and therapies targeting evolving pathogens aim to curb the pathogen and to steer it toward a controlled evolutionary state. Control is leveraged against the pathogen's intrinsic evolutionary forces, which in turn, can drive an escape from control. Here, we analyze a simple model of control, in which a host produces antibodies that bind the pathogen. We show that the leverages of host (or external intervention) and pathogen are often highly imbalanced: an error threshold separates parameter regions of efficient control from regions of compromised control, where the pathogen retains the upper hand. Because control efficiency can be predicted from few measurable fitness parameters, our results establish a proof of principle how control theory can guide interventions against evolving pathogens.

The fixed points of the coupled dynamics, which are relevant for long-term control, are coevolutionary (Nash) equilibrium points. Through control by computation, however, a host can globally maximize its integrated payoff over an extended control period, often trading an initial dip for a later gain. We show that computational control defines a class of dynamical fixed points, called computational equilibria, which differ from Nash equilibria and reflect the added value of computation.

The second part focuses on applications to biomolecular control of pathogens. We analyze a minimal model of control, in which the host produces antibodies that bind to the pathogen. This model captures two complementary control modes, which are associated with different host–pathogen interactions. For ecological control, bound antibodies impede pathogen growth. The control objective is to reduce the pathogen's carrying capacity; a deleterious collateral is the evolution of resistance. For evolutionary control, bound antibodies reduce pathogenicity. The control objective is the adaptive evolution of an antibody binding site (epitope) in the pathogen population; a collateral is the concurrent increase of the carrying capacity. We develop an analytical solution of the minimal model for ecological and evolutionary control. The solution includes maximum-payoff stationary and time-dependent protocols for antibody dosage, and it maps efficiency phase diagrams that delineate parameter regimes of efficient and compromised control.

The control theory of the minimal model has a number of key characteristics we argue to be general features of biomolecular control. First, control phase diagrams contain an error threshold, marking a switch in molecular antibody–antigen recognition and a rapid change of control efficiency. Second, control by computation shows striking differences to evolutionary and regulatory host protocols: computational equilibrium points can reach higher payoff than Nash equilibria and in turn, be outcompeted by protocols with intermittent time dependence. Third, protocols and success of control depend on the pathogen's mutation rate and population size, highlighting how control depends on the underlying population genetics and ecology. We discuss implications of these results for biomedical applications: how information processing impacts mode and efficiency of control and how measurements of core pathogen and host data can be used to predict control outcomes.

## Eco-evolutionary Control Theory

**Eco-evolutionary Dynamics of Pathogens.** Consider a population of pathogens with a quantitative, heritable trait $G$ that is a target of host–pathogen interactions, including control. Selection on the trait is described by a fitness landscape $f_p(G, \zeta)$; this function depends on a time-dependent control protocol $\zeta(t)$ generated by a host system. The eco-evolutionary process of the pathogen population has three components, which are detailed in *Methods* and *SI Appendix*.

First, common mutations with individually small trait effects $\Delta G$ produce a heritable trait distribution $\rho(G, t)$ that is acted upon by selection. As described by [17] in *Methods*, this Darwinian process generates a trait distribution peaked at its mean $\Gamma(t)$, which evolves by diffusion in a free fitness landscape $\psi_p(\Gamma, \zeta)$. This landscape generates an evolutionary force given by its gradient,

$$\partial_\Gamma \psi_p(\Gamma, \zeta) = 2 N_p^e(\Gamma) \, \partial_\Gamma \bar{f}_p(\Gamma, \zeta) + \partial_\Gamma S_p(\Gamma). \qquad \textbf{[1]}$$

The selective force component is generated by the mean population fitness $\bar{f}_p(\Gamma, \zeta)$; below, we will approximate $\bar{f}_p(\Gamma, \zeta) \simeq f_p(\Gamma, \zeta)$, which is appropriate for a peaked trait distribution. The mutational component is generated by the entropy $S_p(\Gamma)$, defined as the log density of sequence states with trait value $\Gamma$. This force measures changes in mutational target, which impede evolution by gain of function ($\partial_\Gamma S_p \, d\Gamma < 0$) but facilitate loss of function ($\partial_\Gamma S_p \, d\Gamma > 0$). The effective population size $N_p^e(\Gamma)$,

which equals the coalescence time of the evolutionary process, enters the trait's heritable variation and its Darwinian response to selection and control (18, 19) (*Methods*).

Second, mutations of large effect $\Delta G$ shift the trait from an initial value $\Gamma_1$ to a value $\Gamma_2 = \Gamma_1 + \Delta G$; the substitution process involves a mean trait $\Gamma(t) = (1 - x(t)) \, \Gamma_1 + x(t) \, \Gamma_2$ and a free fitness $\psi_p(t) = (1 - x(t)) \, \psi_p(\Gamma_1, \zeta(t)) + x(t) \, \psi_p(\Gamma_2, \zeta(t))$ that depend on the frequency $x(t)$ of the mutant allele. While small-effect mutations generate evolutionary trajectories toward a local fitness maximum, large-effect mutations can cross fitness valleys and bridge between different local maxima and their basins of attraction. Large-effect trait shifts can also be generated by pairs of a deleterious and a compensatory mutation. The relative weight of diffusive and large-effect trait evolution will prove to be an important determinant of eco-evolutionary control.

Third, the pathogen population dynamics follows a minimal ecological model, which is given by [18] in *Methods*. This model describes a stochastic birth–death process with basic reproductive rate $\bar{f}_p(\Gamma, \zeta)$ in an ecological niche of constraint $c$. Given a static or slowly varying reproductive rate, these dynamics generate population size fluctuations around a carrying capacity $\bar{N}(\Gamma, \zeta) \simeq \bar{f}_p(\Gamma, \zeta)/c$.

In an individual pathogen population, the controlled evolutionary dynamics defines a control path $(\Gamma, \zeta)$ that tracks the evolving trait $\Gamma(t)$ and the underlying control protocol $\zeta(t)$. We define the evolutionary flux over a time interval $(t_1, t_2)$,

$$\Theta_p(\boldsymbol{\Gamma}, \boldsymbol{\zeta}) = \int_{t_1}^{t_2} \partial_\Gamma \psi_p(\Gamma(t), \zeta(t)) \, \dot{\Gamma}(t) \, dt, \qquad \textbf{[2]}$$

which sums the entropy increments and the scaled fitness increments generated by trait changes along the path $(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$. The flux $\Theta_p(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$ is an important building block of eco-evolutionary control theory: it characterizes the likelihood of the evolutionary path $\boldsymbol{\Gamma}$ given the control protocol $\boldsymbol{\zeta}$. This is shown in *SI Appendix*, building on previous results in stochastic thermodynamics and evolutionary statistics (20, 21). Pathogen evolution often involves large flux amplitudes ($\Theta_p \gg 1$), which can be generated by strong selection or sufficiently complex evolutionary traits. In this case, we obtain a deterministic criterion: observable evolutionary paths have a nonnegative flux,

$$\Theta_p(\boldsymbol{\Gamma}, \boldsymbol{\zeta}) \geq 0; \qquad \textbf{[3]}$$

that is, the evolution of the pathogen acts to increase its instantaneous free fitness $\psi_p$. Importantly, this relation holds not only for the full process but for any subperiod $(t_1, t_2)$, as long as flux amplitudes remain large. Pathogen evolution is often counteracted by control acting to decrease $\psi_p$.

**Minimum-Leverage Relations for Pathogen Evolution.** Eco-evolutionary control is exerted by altering selection: the controlled system is governed by a free fitness landscape $\psi_p(\Gamma, \zeta) = \psi_b(\Gamma) + f_c(\Gamma, \zeta)$, which contains a background component $\psi_b(\Gamma)$ and a host-dependent control component $f_c(\Gamma, \zeta)$. We now use the flux inequality [3] to determine lower bounds on $f_c(\Gamma, \zeta)$ (i.e., a minimum leverage of the controlling onto the controlled system required for successful control).

A control protocol $\zeta$ can elicit a trait value $\Gamma_e$ from an initial value $\Gamma_0$, if there is a trait path $\boldsymbol{\Gamma}$ from $\Gamma_0$ to $\Gamma_e$ that fulfils the flux condition [3] for any segment covered in an arbitrary subperiod $(t_1, t_2)$. That is, the trait $\Gamma(t)$ moves uphill in the local free fitness gradient given by [1] (20). This amounts to the local minimum-leverage condition

$$2 N_p^e(\Gamma) \, \partial_\Gamma f_c(\Gamma, \zeta) > -2 N_p^e(\Gamma) \, \partial_\Gamma f_b(\Gamma) - \partial_\Gamma S_p(\Gamma), \qquad \textbf{[4]}$$

which compares the control force with the evolutionary forces in the absence of control. In particular, the control force must

bridge fitness valleys of the background landscape $f_b(\Gamma)$, which often requires high transient control amplitudes. Control must also overcome bottlenecks of the pathogen entropy landscape $S_p(\Gamma)$, which arise for evolution by gain of function. The deterministic minimum-leverage condition [4] is valid up to small fitness troughs that can be crossed by common mutations at low effective population size.

If the population states $\Gamma_0$ and $\Gamma_e$ are linked by large-effect mutations and if the control amplitude $\zeta$ changes slowly over typical substitution periods, the minimum-leverage relation [4] takes the simpler form

$$\psi_p(\Gamma_e, \zeta) > \psi_p(\Gamma_0, \zeta), \qquad [5]$$

which relates the free fitness values at the end points of the evolutionary path $\boldsymbol{\Gamma}$ (*SI Appendix*) (20, 22). In the case of a constant $N_p^e$, this inequality further reduces to $\Theta_c + \Theta_b > 0$; that is, the control leverage $\Theta_c = 2N_p^e[f_c(\Gamma_e, \zeta) - f_c(\Gamma_0, \zeta)]$ has to exceed the drop in free fitness of the uncontrolled system, $-\Theta_b = \psi_b(\Gamma_0) - \psi_b(\Gamma_e)$. The end point minimum-leverage condition applies whenever large-effect substitutions are frequent on timescales relevant for the control task (i.e., for sufficiently large populations and sufficiently long control periods). A case in point is the maintenance of a controlled pathogen state $\Gamma_e$ by stationary control. The condition [5] says that escape mutations from $\Gamma_e$ to $\Gamma_0$ are suppressed by negative selection. In large populations, this condition is often necessary for successful control. However, as shown by an example given below, it can be undercut in small populations when escape mutations are rare.

Fig. 1 illustrates the minimum-leverage relation for a minimal model of control with a time-independent amplitude $\zeta$. The pathogen free fitness landscape $\psi_p(\Gamma, \zeta)$ has two local maxima, the wild type $\Gamma_{wt}$ and the optimal evolved state $\Gamma_e^*(\zeta)$, which are assumed to be linked by large-effect mutations. We consider two complementary control scenarios: ecological control aimed at reducing the pathogen load $\bar{N}_p$ while maintaining the wild-type trait $\Gamma_{wt}$ against the evolution of a resistant state $\Gamma_e^*$ (Fig. 1A) and evolutionary control aimed at maintaining an evolved equilibrium $\Gamma_e^*$ against reversal to $\Gamma_{wt}$ (Fig. 1B). In both cases, the minimum-leverage relation [5] is seen to delineate a strong control (SC) regime, where the control objective is achieved, and a weak control (WC) regime, where the objective is missed. These regimes and their dependence on the control amplitude $\zeta$ will be further discussed below.

The inequalities [4 and 5] specify the minimum leverage that a controlling host system must exert on the controlled pathogen system in order to elicit a feature that would not evolve spontaneously and to maintain this feature against reverse evolution toward the wild type. These relations are formally related to the maximum-work theorem of thermodynamics, which specifies the minimum-work uptake (or maximum-work release) associated with a given free energy change of a thermodynamic system. Unlike in thermodynamics, however, the minimum-leverage relations say nothing about cost and benefit of control for the controlling system. This requires explicit modeling of the host's control dynamics and of the host–pathogen interactions, to which we now turn.

**Control by Darwinian Evolution or Regulation.** How can pathogen evolution under control be optimized for the controlling host system? To address this question, we have to specify a host payoff function and the resulting dynamics of the control amplitude $\zeta(t)$. Eq. 19 in *Methods* specifies a stochastic update rule: $\zeta(t)$ changes by small increments (i.e., by diffusion) in the payoff landscape $\psi_h(\Gamma, \zeta)$. This local rule, which is analogous to the evolution of the pathogen trait by common mutations, depends only on the payoff gradient $\partial_\zeta \psi_h(\Gamma(t), \zeta(t))$ at the instantaneous point of the control path $(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$. It can be realized by Darwinian
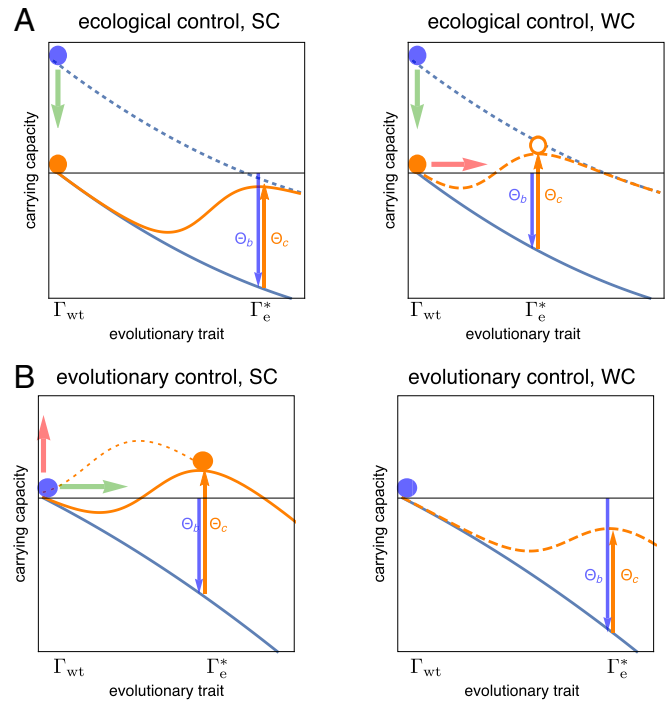


**Fig. 1.** Modes and leverage of eco-evolutionary control. A pathogen population with mean trait $\Gamma$ under control with amplitude $\zeta$ lives in a free fitness landscape $\psi_p(\Gamma, \zeta)$ (orange lines), which is the sum of a background component $\psi_b(\Gamma)$ (blue lines) and a control landscape $f_c(\Gamma, \zeta)$. An evolutionary path from the wild type $\Gamma_{wt}$ to an optimal evolved state $\Gamma_e^*$ involves the control leverage $\Theta_c = 2N_p^e[f_c(\Gamma_e^*, \zeta) - f_c(\Gamma_{wt}, \zeta)]$ (orange arrows) and a change in background free fitness, $\Theta_b = \psi_b(\Gamma_e^*) - \psi_b(\Gamma_{wt})$ (blue arrows). (A) Ecological control, starting from an uncontrolled wild-type pathogen (blue dots), has the objective of reducing the pathogen's carrying capacity (green arrows)—here by antibody binding—and the collateral effect of resistance evolution (red arrows). SC ($\Theta_c + \Theta_b < 0$) suppresses the evolution of resistance and generates a stable wild type (orange dot; i.e., the reverse path from $\Gamma_e^*$ to $\Gamma_{wt}$ fulfils the minimum-leverage condition [5]). WC ($\Theta_c + \Theta_b > 0$) triggers the evolution of resistance (orange circle). (B) Evolutionary control has the objective of eliciting a new pathogen state (green arrow) and the collateral of increasing its carrying capacity (red arrow). Dynamical control elicits the evolved trait along a path of positively selected trait increments, which requires elevated transient control amplitudes (dotted orange line). SC ($\Theta_c + \Theta_b > 0$) generates a stable evolved trait (orange dot; i.e., the path from $\Gamma_{wt}$ to $\Gamma_e^*$ fulfils the minimum-leverage condition [5]). WC ($\Theta_c + \Theta_b < 0$) cannot elicit an evolved state (or triggers a reversal to the wild type).

evolution in a host population, where the control amplitude is a quantitative trait peaked at its population mean value $\zeta(t)$ (12). In this case, the payoff function takes again the form of a free fitness; i.e., $\partial_\zeta \psi_h(\Gamma, \zeta) = 2\, N_h^e(\zeta)\, \partial_\zeta \bar{f}_h(\Gamma, \zeta) + \partial_\zeta S_h(\zeta)$, where $\bar{f}_h(\Gamma, \zeta)$ is the mean population fitness, $S_h(\zeta)$ is an entropy generated by the molecular encoding of the control amplitude, and $N_h^e(\zeta)$ is an effective population size (*Methods*).

Control by local update rules can readily lead to a local maximum of the host's payoff $\psi_h(\Gamma, \zeta)$. In contrast, so-called greedy protocols globally maximize the instantaneous payoff, as described by [20] in *Methods*. In a complex payoff landscape, this involves bridging payoff valleys and requires large host populations that harbor large-effect control amplitude changes in their standing variation. Otherwise, evolutionary update of the control amplitude becomes mutation limited and often prohibitively slow. In appropriate host systems, as in the example given below, global payoff maximization can be implemented more rapidly by regulation. Given a recurrent pathogen, a trained regulatory network can sense the instantaneous pathogen load and generate

an approximately optimal control amplitude as output. Because regulation is faster than the eco-evolutionary pathogen dynamics, it produces a greedy maximization of the instantaneous host payoff. These dynamics can again be modeled as a stochastic or deterministic process in the payoff landscape $\psi_h(\Gamma, \zeta)$ (*Methods*).

Given a control path $(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$, we define the host flux

$$\Theta_h(\boldsymbol{\Gamma}, \boldsymbol{\zeta}) = \int_{t_1}^{t_2} \partial_\zeta \, \psi_h(\Gamma(t), \zeta(t)) \, \dot{\zeta}(t) \, dt, \qquad \textbf{[6]}$$

in analogy to the pathogen flux given by [2]. Under a joint stochastic process of pathogen evolution and (gradient or greedy) instantaneous control, these fluxes take symmetric roles. The total flux $\Theta_p + \Theta_h$ is related to the probability of the path $(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$ (*SI Appendix*). In the deterministic limit, instantaneous-update protocols have a nonnegative flux,

$$\Theta_h(\boldsymbol{\Gamma}, \boldsymbol{\zeta}) \geq 0; \qquad \textbf{[7]}$$

that is, control acts to increase the instantaneous host payoff $\psi_h$. These dynamics are often counteracted by pathogen evolution acting to decrease $\psi_h$.

A stable fixed point $(\Gamma^\dagger, \zeta^\dagger)$ of deterministic pathogen evolution and instantaneous-update control satisfies the conditions

$$(\Gamma^\dagger, \zeta^\dagger) = \arg \max_\Gamma \psi_p(\Gamma, \zeta^\dagger) = \arg \max_\zeta \psi_h(\Gamma^\dagger, \zeta), \qquad \textbf{[8]}$$

where the maxima are to be taken over all trait values and control amplitudes accessible over the relevant period. These conditions define a Nash equilibrium of an evolutionary game with payoff functions $\psi_p$ and $\psi_h$ (23). In the strong-selection limit $(N_p^e, N_h^e \gg 1)$, they reduce to conditional maximization of the fitness functions $f_p$ and $f_h$. In other words, control by Darwinian evolution or regulation leads to optimization problems familiar from other coevolutionary and ecological systems.

**Control by Computation.** A more ambitious goal, and the subject of control theory, is to optimize protocols toward an objective defined over an extended period of the dynamics, including the future. Here, we use a scoring function of the form

$$\Omega(\boldsymbol{\Gamma}, \boldsymbol{\zeta}) = \Psi(\boldsymbol{\Gamma}, \boldsymbol{\zeta}) - \lambda \, T_\delta(\boldsymbol{\Gamma}, \boldsymbol{\zeta}). \qquad \textbf{[9]}$$

The first term is the time integral of the host payoff function, $\Psi(\boldsymbol{\Gamma}, \boldsymbol{\zeta}) = \int \psi_h(t) \, dt$. For computational control of human pathogens, this function may include survival, life quality, and public health components. The second term penalizes the duration of controlled adaptation, $T_\delta(\boldsymbol{\Gamma}, \boldsymbol{\zeta})$, defined as the time needed to reach a payoff maximum $\psi_h^*$ by a margin $\delta$, with a coefficient $\lambda \geq 0$. This term weighs in the speed of the control dynamics toward a given objective, which is often crucial in biomedical applications but comes at the price of a reduced payoff score $\Psi$. Maximizing $\Omega$ requires computation preempting the future evolution of the pathogen. This problem can be cast in the form of a so-called Hamilton–Jacobi–Bellman (HJB) equation for the optimal score with given boundary conditions. In *Methods*, we solve the HJB equation analytically for deterministic control with scoring functions of the form [9]. Computational protocols often steer through payoff valleys, trading transient periods with $\Theta_h < 0$ and a concurrent decline of $\psi_h$ for a later gain. Such protocols violate the flux condition [7]; hence, they cannot be realized by any instantaneous-update rule in the payoff landscape $\psi_h(\Gamma, \zeta)$. As discussed below, this does not exclude more complex evolutionary or regulatory circuits programming investments for future payoff gains.

A stationary state $(\Gamma^*, \zeta^*)$ that maximizes $\Omega$ satisfies the conditions

$$(\Gamma^*, \zeta^*) = \arg \max_\Gamma \psi_p(\Gamma, \zeta^*) = \arg \max_\zeta \psi_h(\Gamma^*(\zeta), \zeta); \qquad \textbf{[10]}$$

that is, it is evolutionarily stable against pathogen escape mutations and maximizes $\psi_h$ by preempting pathogen response before it happens. We refer to $(\Gamma^*, \zeta^*)$ as a computational equilibrium point. As we will show below, the optimality condition [10] can lead to different fixed points than the Nash equilibrium condition [8]. Specifically, computational control can reach higher stationary payoff $\psi_h^*$, but the control amplitude $\zeta^*$ has to be maintained by computation against a payoff gradient. Explicitly time-dependent computational control, including the protocols of adaptive trait formation and metastable control discussed below, can reach high payoff faster than evolutionary or regulatory protocols and maintain a higher average payoff than any stationary protocol.

### Pathogen Control by Antibodies

**Antibody–Antigen Interactions.** We now focus on a specific control scenario, in which a host exerts control by producing antibodies that bind to a pathogen (also referred to as antigen in this context). The probability that the pathogen is bound,

$$P_{\text{bind}}(G, \zeta) = \frac{1}{1 + \zeta^{-1} \exp(\epsilon G)}, \qquad \textbf{[11]}$$

depends on the pathogen trait $G$ (with wild-type value $G_{\text{wt}} = 0$) and the antibody density or dosage $\zeta$ (measured in units of the dissociation constant or half-maximal inhibitory concentration [IC50] of the wild-type pathogen). We consider two cases, which cover the control modes illustrated in Fig. 1. For ecological control ($\epsilon = +1$), $G$ is a resistance trait (the log of the dissociation constant; i.e., a population with evolved resistance $[\Gamma_e^* > 0]$ has reduced binding). For evolutionary control ($\epsilon = -1$), $G$ is the epitope affinity (the log of the association constant); this mode is to elicit an evolved population ($\Gamma_e^* > 0$) with increased binding.

**Pathogen Fitness and Host Payoff Landscapes.** We assume that host and pathogen live in coupled landscapes of the form

$$f_p(\Gamma, \zeta) = f_{p,0} - c_p \Gamma - \epsilon q_{ph} \, P_{\text{bind}}(\Gamma, \zeta),$$
$$\psi_h(\Gamma, \zeta) = \psi_{h,0} - c_h \zeta - q_{hp} \, L_p(\Gamma, \zeta). \qquad \textbf{[12]}$$

The pathogen has a background fitness $f_{p,0}$ in the uncontrolled wild type, a background cost $c_p > 0$ per unit of the trait (the linear form is taken for simplicity), and a binding-dependent control term $f_c = -\epsilon q_{ph} \, P_{\text{bind}}$ with a selection coefficient $q_{ph} > 0$. The host has a background payoff $\psi_{h,0}$ in the absence of control and pathogens, a production cost $c_h$ per unit of antibody, and an interaction term depending on the pathogen load $L_p(\Gamma, \zeta)$ with a selection coefficient $q_{hp} > 0$. In the case of ecological control, the load is generated by the full pathogen population, $L_p(\Gamma, \zeta) = f_p(\Gamma, \zeta)$, where the pathogen is assumed to be at its carrying capacity $\bar{N} = f_p/c$ given by [18] and measured in units of $c$. In the case of evolutionary control, we use a load function $L_p(\Gamma, \zeta) = (1 - P_{\text{bind}}(\Gamma, \zeta)) f_p(\Gamma, \zeta)$, assuming that bound pathogens lose their pathogenicity. Fig. 2 shows the landscapes of [12] for ecological and evolutionary control. In both cases, the pathogen has two local fitness maxima (solid and dashed lines) with a rank order depending on the antibody dosage.

**Optimal Stationary Control.** We now focus on pathogens that can be contained by sustained treatment but cannot be eradicated by a short-time protocol (this is, currently, an appropriate assumption for HIV and some cancers). First, we compute the maximum-payoff stationary control protocol, as given by the
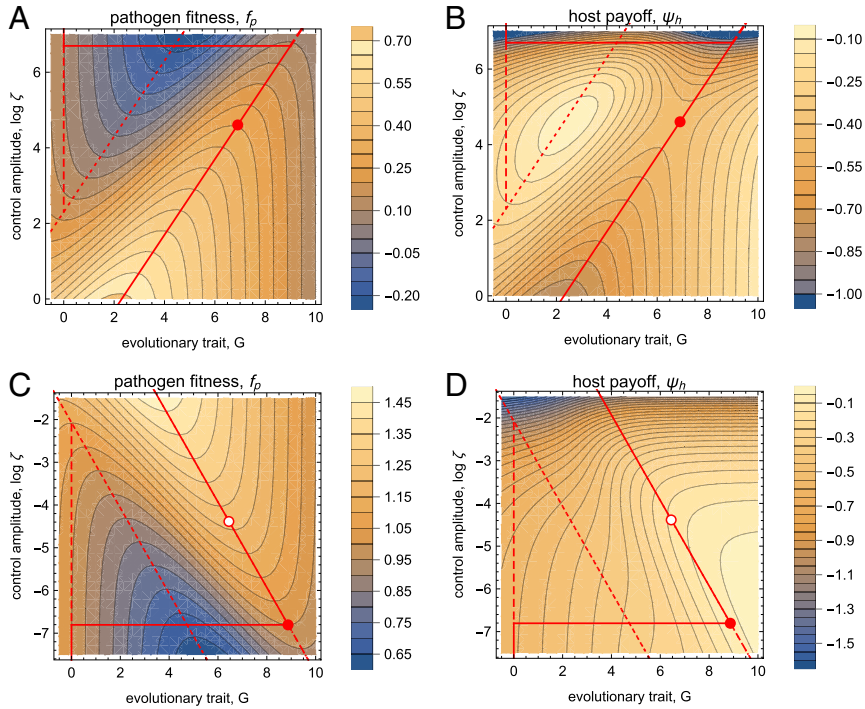
**Fig. 2.** Fitness and payoff landscapes of pathogen control. (*A* and *B*) Ecological control and (*C* and *D*) evolutionary control. Pathogen fitness $f_p$ (*A* and *C*) and host payoff $\psi_h$ (*B* and *D*) are shown as functions of the log control amplitude (antibody dosage), $\log \zeta$, and the pathogen trait, $G$. Specific control loci: pathogen fitness minimum (dotted), pathogen fitness maxima (stable: solid; metastable: dashed; stability switch: horizontal lines), computational equilibrium control ($\log \zeta^*$, $G^*$) (dots), Nash equilibrium of evolutionary control, ($\log \zeta^\dagger$, $G^\dagger$) (circles). Model parameters: $f_{p,0} = 1$, $q_{ph} = 1/q_{ph} = 0.9$, $c_p = 0.09$, $c_h = 0.001$ (ecological control); $f_{p,0} = 1$, $q_{ph} = 0.9$, $q_{hp} = 5/9$, $c_p = 0.09$, $c_h = 5$ (evolutionary control).

computational equilibrium condition [**10**], which is relevant for long-term treatment. Given a large pathogen population under stationary control, the trait evolves to a dosage-dependent fitness maximum, $\Gamma^*(\zeta) = \arg\max_\Gamma \psi_p(\Gamma, \zeta) \simeq \arg\max_\Gamma f_p(\Gamma, \zeta)$, with a negligible contribution of the entropy. In the landscapes of [**12**], the resulting optimal control ($\Gamma^*$, $\zeta^*$) can be computed analytically ([**27–30**] in *Methods*). In the case of ecological control, we find two control regimes, SC and WC, as shown schematically in Fig. 1. These regimes are separated by a transition at

$$c_h^\bullet = q_{ph} q_{hp} \exp\left(1 - \frac{q_{ph}}{c_p}\right). \qquad [\mathbf{13}]$$

In the SC regime ($c_h < c_h^\bullet$), the optimal protocol keeps the pathogen in its wild type, maintaining antibody–antigen binding and suppressing the evolution of resistance. In the WC regime ($c_h > c_h^\bullet$), control is compromised by evolved pathogen resistance (this case is shown in Fig. 2*A*). At the transition, the resistance trait switches from $\Gamma^* = 0$ (SC) to $\Gamma^* = q_{ph}/c_p - 1$ (WC), akin to the order parameter of a first-order phase transition, while $\psi_h$, $f_p$, and $\zeta^*$ remain continuous (*SI Appendix*, Fig. S1). Antibody–antigen binding switches from $P_{\text{bind}} \sim 1$ (SC) to $P_{\text{bind}} \ll 1$ (WC); that is, the transition can be interpreted as an error threshold of molecular recognition (24). To map the phase diagram of optimal stationary control, we define the control efficiency $\eta^* = (\psi_h(\Gamma^*, \zeta^*) - \psi_{h,0})/\delta\psi_{\max}$ as the payoff gain relative to its maximum, $\delta\psi_{\max} = q_{ph} q_{hp}$, which is reached for perfect control at no cost. Fig. 3*A* shows the efficiency as a function of the cost parameters ($c_p$, $c_h$). The error threshold $c_h^\bullet(c_p)$ (yellow line) marks a rapid decline from $\eta^* \sim 1$ in the SC regime to $\eta^* \ll 1$ in most of the WC regime.

For stationary evolutionary control, we find a similar emergence of two control regimes with an error threshold

$$c_h^\bullet = q_{hp} \frac{c_p}{q_{ph}} \left(1 - \frac{c_p}{q_{ph}}\right) \exp\left(\frac{q_{ph}}{c_p} - 1\right). \qquad [\mathbf{14}]$$

In the SC regime ($c_h < c_h^\bullet$), control maintains an evolved trait $\Gamma^* > \Gamma_{\text{wt}}$ with antibody–antigen binding beneficial to host and pathogen. In the WC regime ($c_h > c_h^\bullet$), control is too weak to maintain the evolved trait, and the pathogen reverts to the unbound wild type. At the transition, $\Gamma^*$ and $\zeta^*$ switch, while $f_p$ and $\psi_h$ remain continuous (*SI Appendix*, Fig. S1). Fig. 3*B* shows the control efficiency $\eta^* = (\psi_h(\Gamma^*, \zeta^*) - \psi_{h,0})/(q_{hp} f_{p,0})$ as a function of the cost parameters ($c_p$, $c_h$). In this case, $\eta^*$ is a decreasing function of both parameters (i.e., host and pathogen cost impede evolutionary control). The error threshold $c_h^\bullet(c_p)$ marks the transition between $\eta^* > 0$ (SC) and $\eta^* = 0$ (WC).

**Instantaneous-Update Control.** Next, we consider the dynamical accessibility of the computational equilibrium point. In the case of ecological control, ($\Gamma^*$, $\zeta^*$) is a stable fixed point of deterministic pathogen evolution and instantaneous dosage update, which satisfies the Nash equilibrium condition [**8**] (*Methods*). Hence, this point can be reached by a control dynamics based on host evolution or regulation; examples of such control paths are shown in Fig. 4*A*. In the SC regime, the pathogen wild type is stable ($\Gamma^* = \Gamma_{\text{wt}}$); the conditional host payoff landscape $\psi_h(\Gamma_{\text{wt}}, \zeta)$ has a single peak $\zeta^*$ that can be reached by local or greedy instantaneous-update protocols. In the WC regime, control paths start with a large-effect escape mutation of the pathogen (orange arrow), which is followed by a coupled dynamics of common pathogen mutations and antibody dosage changes. A control path with local stochastic update gets localized to ($\Gamma^*$, $\zeta^*$);
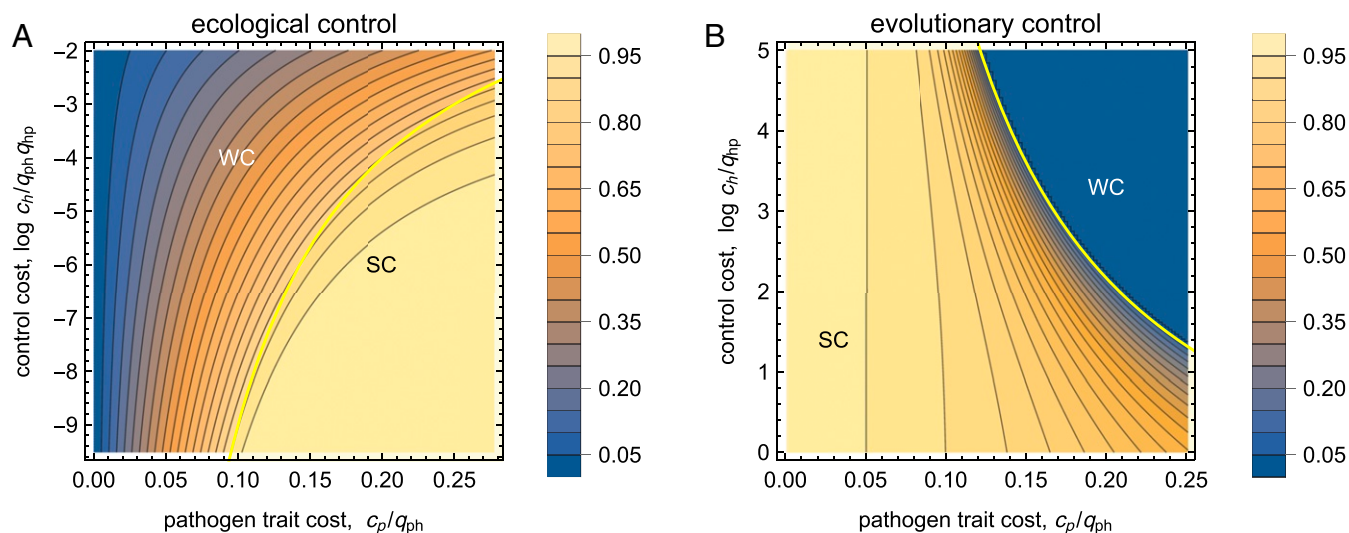
**Fig. 3.** Phase diagrams of stationary control. (*A*) Ecological control and (*B*) evolutionary control. The control efficiency $\eta^*$ is shown as a function of the (scaled) cost parameters $c_p$ and $c_h$. A yellow line marks the error threshold $c_h^{\bullet}(c_p)$ between WC and SC.

deterministic paths with local or greedy update converge to the same point. Host payoff and pathogen fitness evolve in a non-monotonic way along these control paths (Fig. 4*B*), reflecting the mutually deleterious effects of host and pathogen in the landscapes of [**12**]. The host and pathogen fluxes, however, are positive throughout in accordance with the flux conditions [**3**, **4**, and **7**] (Fig. 4*C*). In this case, computational update generates control paths with a similar convergence to the fixed point $(\Gamma^*, \zeta^*)$; evolutionary, regulatory, and computational protocols have only transient differences in their time-dependent efficiency (*SI Appendix*, Fig. S2). We conclude that the ecological control equilibrium can be dynamically reached and maintained in a robust way.

**Computational Control of Adaptive Trait Formation.** In contrast, the optimal stationary (computational equilibrium) protocol $(\Gamma^*, \zeta^*)$ for evolutionary control in the SC regime is not a Nash equilibrium of deterministic pathogen evolution and instantaneous dosage update. Hence, it cannot be reached or maintained by an evolutionary or regulatory host system. Time-dependent computational protocols maximizing the score $\Omega(\Gamma, \zeta)$ in the fitness and payoff landscapes [**12**] can be evaluated analytically (Fig. 4*D* and *Methods*). These protocols have two phases. Starting from a wild-type pathogen, we apply a high initial dosage $\zeta_{in}$ to trigger an intermediate-effect gain-of-function mutation of the pathogen that establishes antibody binding at a trait value $G_{in}$ (orange arrow). This initial-phase protocol flattens the pathogen fitness valley for trait formation but generates a drop in host payoff (Fig. 4*E*). In a second "breeding" phase, small-effect trait and dosage changes increase host payoff and steer the control path toward the stationary point $(\Gamma^*, \zeta^*)$. The pathogen evolutionary flux is positive throughout this process, following the conditions [**3** and **4**]. The host flux has negative increments in the initial phase and final phase, which violate the condition [**3**] and mark strong deviations of the computational protocol from evolutionary or regulatory protocols (Fig. 4*F*). These instantaneous-update protocols converge to a Nash equilibrium point $(\Gamma^{\dagger}, \zeta^{\dagger})$ or to the WC fixed point $(\Gamma_{wt}, 0)$, both of which have smaller payoff than the computational equilibrium protocol $(\Gamma^*, \zeta^*)$ (*SI Appendix*, Fig. S2).

Computational protocols of adaptive trait formation maximize the score $\Omega$ by jointly tuning payoff and duration of the initial and the breeding phase. The optimal protocol depends

on the speed scoring parameter $\lambda$; protocols with large $\lambda$ steer the pathogen along paths of near-maximal speed of trait evolution (*SI Appendix*, Fig. S3). Importantly, the optimal protocol also depends on the pathogen population size, which affects the mutational supply in the adaptive process (*SI Appendix*, Fig. S3). To map these effects, we evaluate the score of a time-dependent path relative to the optimal stationary protocol, $\Delta\Omega = \Delta\Psi - \lambda T_{\delta}$ with $\Delta\Psi = \int [\psi_h(\Gamma(t), \zeta(t)) - \psi_h^*] \, dt$; we call $(-\Delta\Psi)$ the cost of adaptation. In *Methods*, we solve an extended HJB equation for $\Delta\Omega$ with a maximum-likelihood approximation for the initial phase, which uniquely relates the gain-of-function amplitude $G_{in}$ to the initial dosage $\zeta_{in}$. These protocols have a scoring function of the form

$$\Delta\Omega = \frac{1}{\theta} \omega(\theta, G_{in}/\epsilon_0), \qquad [15]$$

which depends on the sequence diversity of the pathogen trait, $\theta = 2\mu N_e$, and on the gain-of function trait measured in units of the mutational scale $\epsilon_0$ (*Methods*). In small populations ($\theta \lesssim 1$), the maximum-score protocol has $G_{in}/\epsilon_0 \sim 1$; larger values of $G_{in}$ are suppressed by an exponentially longer waiting time for a gain-of-function mutation. In large populations ($\theta \gtrsim 1$), the maximum-score protocol has $G_{in}/\epsilon_0 \sim \log \theta$. This amplitude is set by a simple rule: the optimal $G_{in}$ is the largest adaptive gain-of-function amplitude likely to be seeded from standing variation of the initial pathogen population (*Methods*). That is, the optimal protocol of adaptive trait formation eliminates the bottleneck of waiting for de novo gain-of-function mutations by circumnavigating the pathogen fitness valley (*SI Appendix*, Fig. S4). Larger values of $G_{in}$ are still suppressed by waiting time; smaller ones require an excess control cost $c_h\zeta_{in}$ and additional breeding time. In both regimes, the cost of adaptation decreases, and $G_{in}$ monotonically increases with increasing $\theta$ (Fig. 4*G*). We conclude that in small pathogen populations, adaptive trait formation requires higher control effort and is more costly.

**Metastable Computational Control.** Strikingly, optimal computational protocols can be time-dependent even for long-term control with the stationary objective of maximizing the host payoff integral $\Omega = \Psi$ (we set $\lambda = 0$ here). As an example, consider ecological control in the WC regime. The control protocols discussed above, which lead to the equilibrium point $(\Gamma^*, \zeta^*)$, apply

**Fig. 4.** Control dynamics. (*A–C*) Ecological control by instantaneous-update protocols. (*A*) Control paths $(\Gamma, \zeta)$ generated by local deterministic (dashed orange) and greedy (solid orange) control dynamics. These paths start with a pathogen escape mutation (gray square) and converge to the optimal stationary protocol $(\Gamma^*, \zeta^*)$ (red dot), which is a computational and a Nash equilibrium. (*B*) Time-dependent pathogen fitness, $f_p(t)$, and host payoff, $\psi_h(t)$, of these control paths. (*C*) Pathogen flux, $\Theta_p(t)$, and host flux, $\Theta_h(t)$. Both fluxes are monotonically increasing functions of the path coordinate $\Gamma(t)$. (*D–G*) Evolutionary control for adaptive trait formation. (*D*) Deterministic computational control path in the SC regime (orange line). This path maximizes the score $\Omega(\Gamma, \zeta)$ for given values of the speed parameter and of the pathogen sequence diversity (here $\lambda = 0$, $\theta = 1$) (*SI Appendix*, Fig. S3). The initiation phase $(\Gamma \leq G_{in})$ with a gain-of-function mutation (gray square) is followed by a breeding phase $(\Gamma > G_{in})$; the path converges to the optimal stationary protocol (red dot), which is a computational but not a Nash equilibrium. (*E*) Pathogen fitness, $f_p(t)$, and host payoff, $\psi_h(t)$, along the computational control path. (*F*) Fluxes $\Theta_p(t)$ and $\Theta_h(t)$ as functions of the path coordinate $\Gamma(t)$. This path has negative increments of the host flux $(\partial\Theta_h/\partial\Gamma < 0)$ at the start and in the final segment. (*G*) Cost of adaptation, $(-\Delta\Omega)$ (orange), and initial gain-of-function trait, $G_{in}$ (blue), of the maximum-score computational protocol as a function of the pathogen diversity, $\theta$ for $\lambda = 0$. (*H–K*) Metastable ecological control. (*H*) Time-dependent protocol $\zeta(t)$ with baseline dosage $\zeta_{base}$, boost dosage $\zeta_{res}$, and boost duration $T_{res}$ (orange); pathogen escape mutant frequency $x(t)$ (blue) using a detection threshold $x(t) > 0.05$ to initiate rescue boost (dashed line). (*I*) Pathogen mean fitness, $\bar{f}_p(t)$, and host payoff, $\psi_h(t)$. (*J*) Fluxes, $\Theta_p(t)$ and $\Theta_h(t)$. (*K*) Efficiency of maximum-score metastable protocols, $\eta_{ms}$ (orange), as a function of the pathogen diversity, $\theta$, together with the efficiency of the optimal stationary protocol, $\eta^*$ (red dashed). For $\theta < \theta_c$, metastable control outperforms stationary control. Parameters: $\epsilon_0 = 1$, others as in Fig. 2.

to sufficiently large pathogen populations that evolve resistance by frequent escape mutations (Fig. 4*A*). In smaller populations, such mutations become rare, and the pathogen will stay in a metastable state (here, the wild type) over extended periods. Computational control can reinforce metastability by deepening the pathogen fitness valley for escape mutations; we refer to such protocols as metastable control. Provided we can detect escape mutants at low frequency, we can keep the bulk pathogen population permanently in the metastable state by a two-state protocol (Fig. 4*H*). As long as no escape mutant is detected, we apply a baseline antibody dosage $\zeta_{base}$ to jointly curb pathogen load and escape rates. After an escape mutant is detected above a threshold frequency $x_0$, we apply a rescue boost of higher dosage $\zeta_{res}$ over a period $T_{res}$ to drive that mutant to loss. Metastable control shows intermittent drops in host payoff

(Fig. 4*I*) and flux (Fig. 4*J*), which mark strong deviations from instantaneous-update protocols.

For metastable control over a long period $T$, the average score relative to the optimal stationary state per unit of time, $\Delta\omega = \Delta\Omega/T$, can be written in the form

$$\Delta\omega = \Delta\omega_{base} - \theta\, c_{res}, \qquad \text{[16]}$$

where $\Delta\omega_{base}$ refers to the baseline protocol and $c_{res} > 0$ is proportional to the average payoff cost per rescue. This expression governs the joint optimization of baseline and boost protocols (Eq. **37** in *Methods* and *SI Appendix*, Fig. S4). For $\Delta\omega_{base} > 0$ and $\theta < \theta_c = \Delta\omega_{base}/c_{res}$ (i.e., at sufficiently low mutation rates or in sufficiently small populations), metastable control outperforms stationary control. This is shown in Fig. 4*K*, which

compares the efficiency of maximum-score metastable and stationary protocols, $\eta_{ms}$ and $\eta^*$, as functions of $\theta$.

For evolutionary control in the WC regime, a similar protocol can stabilize a metastable evolved trait $\Gamma_e$ against reversal to the wild type. Alternative protocols exploiting metastable pathogen states include temporarily suspending control (25) or switching to a secondary defense against escape mutants. All of these time-dependent protocols depend on $\theta$ in a similar way, indicating that metastable control becomes more effective in small pathogen populations.

## Discussion

In this paper, we solve the optimization problem for a minimal model of pathogen control operating on realistic biomolecular host–pathogen interactions. Specific features of host and pathogen biology enter the eco-evolutionary control theory of this system at several stages. In particular, control mechanisms are based on physiological effects of host–pathogen interactions—here, antibody–antigen binding—and the objective functions of control, Eq. **12**, account for the fitness and payoff effects of these interactions. As we have shown, these biological features strongly impact the eco-evolutionary control dynamics and the efficiency of optimized protocols.

A salient feature of eco-evolutionary control is the emergence of high- and low-efficiency parameter regimes separated by an error threshold of molecular recognition (Fig. 3). This behavior can be traced to the nonlinearities in the Hill function of antibody–antigen binding, Eq. **11**. First, binding-mediated control has a diminishing-return leverage, which is bounded by the pathogen selection coefficient $q_{ph}$. Second, the control cost depends exponentially on the chemical potential ($\log \zeta$), which determines the evolved pathogen state $\Gamma_e$ (Fig. 2). Hence, the minimum-leverage control [**5**] has a moderate cost in the SC regime but is too expensive or mechanistically impossible in the WC regime. Because such nonlinearities are a generic feature of biomolecular interactions, we expect error thresholds to emerge also in more complex models of pathogen control.

In biomedical applications of pathogen control, it is crucial to predict the likely efficiency of control prior to any intervention. Our results show that few independently measurable cost parameters can inform such estimates. The pathogen cost parameters ($c_p$ and $q_{ph}$ in the minimal model) are routinely measured in dosage–response assays (26). The control cost parameter $c_h$ can, for example, be estimated as the metabolic production cost per unit of antibody (27, 28). The arguably most case-specific parameter is the pathogen cost to the host. This parameter can be measured by fitness assays in microbial systems but is more complex to assess for human hosts.

Optimal control protocols and their efficiency also depend on mutation rate and population size of the pathogen. Eqs. **15** and **16** describe opposing effects: adaptive formation of pathogen traits becomes less costly in large populations; maintenance of traits by metastable control is more efficient in small populations. These effects reflect differences between control strategies. Evolutionary control aimed at eliciting a pathogen trait has to circumnavigate valleys of the pathogen's fitness and entropy landscape in order to catalyze its adaptive dynamics. In contrast, metastable control has to broaden fitness valleys surrounding a metastable state in order to suppress pathogen adaptation by escape mutations. In both cases, we can compute optimized control paths by combining probabilistic, discrete jumps across fitness valleys with continuous dynamics on smooth flanks in between. These navigation principles and their computational implementation are expected to extend to strong-selection control in more complex landscapes.

An equally important issue for control by humans and by natural hosts is optimizing the control dynamics in tune with the monitoring and computation capabilities of the host system (29). Instantaneous-update dynamics following local payoff gradients can often be realized efficiently by Darwinian evolution in the host (i.e., by variation of and selection on antibody levels). Protocols based on regulation can, in principle, circumnavigate payoff valleys and implement an approximate maximization of the instantaneous host payoff. However, signaling and regulatory networks require prior training by learning or by evolution, and they generate an additional cost to the host. Optimization of control toward future objectives depends, to various degree, on computation. The minimal model displays the relative efficiency of these control modes. For ecological control, the computational equilibrium point (i.e., the stationary state of maximal host payoff) is also a Nash equilibrium and can be reached by instantaneous-update protocols; for evolutionary control, this point can only be reached by computation. In more complex systems, the success of human computational control is limited by the ability to predict pathogen evolution. An example is vaccine selection for human influenza based on predictive analysis, where current methods have a prediction and control horizon of about one year (8, 16). For biotic, nonhuman host systems, it remains a fascinating question how far evolutionary and regulatory mechanisms can emulate control by computation.

A case in point is metastable control, which realizes the time-independent objective of maintaining a controlled pathogen state by a time-dependent protocol responding to recurrent pathogen attacks. In the control theory literature, this class of protocols is known as closed loop control (14). Here, we have solved a minimal model of metastable control, which maintains a metastable pathogen state against escape mutations by a two-state protocol of baseline control and rescue boosts. We find that metastable control can outperform the computational equilibrium protocol ($\Gamma^*, \zeta^*$) under two conditions: there is a metastable point of high host payoff (undercutting the quasistatic minimum-leverage relation [**5**]), and escape mutants have a substantial intrinsic cost (generating positive selection for the metastable state during rescue boosts). In the minimal model, an optimized resource allocation between baseline and rescue protocols can only be achieved by computation. We can compare this model with the adaptive immune system of vertebrates, which controls multiple pathogens by a complex pattern of resource allocation (30). From the perspective of control theory, the immune system mounts a bilayer baseline protocol of naive and memory B cells; control boosts involving the evolution of target-specific high-affinity antibodies (affinity maturation) respond to acute infections by specific antigens (akin to escape mutations in the minimal model). Remarkably, affinity maturation continues for several weeks after an infection, which is an investment toward future infections by similar antigens. Together, adaptive immunity may be regarded as an instance of metastable control by nested circuits of evolution (of antibody–antigen affinities) and regulation (of antibody levels).

In summary, this paper establishes a conceptual framework and infers navigation principles for eco-evolutionary control based on a minimal model of host–pathogen interactions with stationary control objectives. Several complexities of biomedical control problems are not captured by the minimal model. For example, the human immune system has a complex antibody repertoire, and pathogens often have multiple antigenic binding sites (microbial and viral epitopes, cancer neoantigens). This generates multiple antibody–antigen interactions (i.e., multiple potential control channels with independent control amplitudes, leverage, and cost parameters). Conversely, pathogens with a high mutation rate often have multiple channels of escape mutations from a given antibody (31). Consequently, the fitness and payoff

landscapes of host and pathogen live in a multidimensional parameter space; an appropriate control dynamics is to navigate this space toward high-efficiency protocols. An important example is the competitive selective dynamics of broadly neutralizing and specific antibodies for HIV (11, 13). Another layer of complexity arises for pathogens embedded in microbiota, which are multispecies systems with a tightly connected ecology. The control of a given pathogen can perturb the entire microbiota, generating cross-resistance of multiple pathogens and complex collateral effects for the host (32). Gearing up eco-evolutionary control theory to these systems is an important avenue for future work.

## Methods

**Stochastic Pathogen Dynamics.** We consider a population of pathogens with a quantitative trait $G$, which has a peaked trait distribution $\rho(G)$ characterized by its mean $\Gamma$ and variance $\Delta$. We describe the evolution of the trait and the concurrent population dynamics in an ecological niche by coupled stochastic equations for the mean $\Gamma$ and the population size $N$,

$$\dot{\Gamma} = D \frac{\partial \psi_p(\Gamma, t)}{\partial \Gamma} + \chi_\Gamma(t), \qquad [17]$$

$$\dot{N} = f(\Gamma, t)N - cN^2 + \chi_N(t), \qquad [18]$$

with white noise $\chi_\Gamma(t), \chi_N(t)$ of mean 0 and variance $\langle \chi_\Gamma(t)\chi_\Gamma(t')\rangle = D\,\delta(t-t')$, $\langle \chi_N(t)\chi_N(t')\rangle = N\,\delta(t-t')$. By [1], the trait dynamics of Eq. 17 depend on the entropy landscape $S(\Gamma)$, which is defined as the log density of states with trait value $\Gamma$, and the fitness seascape $f(\Gamma, t)$, which is often explicitly time dependent. The trait diffusion constant, $D = U\epsilon_0^2$, is set by the total rate $U$ and the mean square trait effect $\epsilon_0^2$ of mutations at genomic loci encoding the trait (19, 33). The trait response to selection, $\Delta = 2DN_e$, is also proportional to the effective population size $N_e$, which equals the coalescence time of the evolutionary process (18, 19). This quantity depends on the population size dynamics in a model-dependent way, often generating an inhomogeneous response to selection, $\Delta(\Gamma) = 2DN_e(\Gamma)$. The population dynamics of [18] depend on the fitness $f(\Gamma, t)$, which sets the basic reproductive rate, and the constraint parameter $c$ of the ecological niche. Given a static or slowly varying fitness function, these dynamics generate population size fluctuations around a carrying capacity $\bar{N}(\Gamma, t) \simeq f(\Gamma, t)/c$. Details of this stochastic calculus are given in *SI Appendix*.

**Instantaneous-Update Control Dynamics.** The diffusive update rule for the control amplitude takes the form

$$\dot{\zeta} = \Delta_h \frac{\partial f_h(\Gamma, \zeta)}{\partial \zeta} + \chi_\zeta, \qquad [19]$$

with white noise of mean 0 and variance $\langle \chi_\zeta(t)\chi_\zeta(t')\rangle = D_\zeta\,\delta(t-t')$, similar to the stochastic trait dynamics given by [17]. Greedy update is defined by

$$\zeta_{\max}(t) = \arg\max_\zeta f_h(\Gamma(t), \zeta) + \chi_\zeta. \qquad [20]$$

The deterministic limit of these dynamics ($D_\zeta = 0$) is used in Fig. 4*A*; stochastic control paths are shown in *SI Appendix*, Fig. S2.

**Computational Control Dynamics.** Given deterministic pathogen evolution and landscapes $\psi_p(\Gamma, \zeta)$, $\psi_h(\Gamma, \zeta)$ without explicit time dependence, maximum-score control paths based on common mutations can be computed from a reduced payoff function,

$$\Delta \psi_\lambda(\Gamma, \zeta) \equiv \psi_h(\Gamma, \zeta) - \psi_h^* - \lambda H_\delta(\psi_h^* - \psi_h(\Gamma, \zeta)). \qquad [21]$$

Here, $\psi_h^*$ is the payoff at the computational equilibrium, Eq. 10, and we define $H_\delta(x) = 1$ for $x > \delta$ and $H_\delta(x) = 0$ otherwise. The optimal control amplitude as a function of the trait is given by

$$\zeta^*(\Gamma) = \arg\max_\zeta \frac{\Delta\psi_\lambda(\Gamma, \zeta)}{V_p(\Gamma, \zeta)}, \qquad [22]$$

where $V_\Gamma(\Gamma, \zeta) = D\,\partial\psi_p(\Gamma, \zeta)/\partial\Gamma$ is the deterministic evolutionary speed in [17]. The maximum payoff $\Delta\psi_\lambda^*(\Gamma) = \max_\zeta \Delta\psi_\lambda(\Gamma, \zeta)$ gives the reduced score, relative to the optimal stationary protocol, of the optimal control path linking an initial trait value $G$ with the equilibrium value $\Gamma^*$,

$$J(G) = \int_G^{\Gamma^*} \frac{\Delta\psi_\lambda^*(\Gamma)}{V_p(\Gamma, \zeta)}\, d\Gamma. \qquad [23]$$

These results are derived in *SI Appendix*.

**Pathogen Equilibrium States.** In the parameter regime of interest, the pathogen fitness landscape of [12] has two local fitness maxima. 1) The wild type $\Gamma = 0$ (vertical lines in Fig. 2) is assumed to be a boundary (i.e., the trait evolution generates only values $\Gamma \geq 0$). 2) The evolved state $\Gamma_e(\zeta)$ (inclined lines in Fig. 2) exists in the regime $c_p/q_{ph} < 1/4$, which provides a bound to the pathogen cost. Maximization of $f_p(G, \zeta)$ with respect to $G$ determines two characteristics of the evolved state. First, the trait value $\Gamma_e$ is shifted by an amount $r_e \equiv \Gamma_e(\zeta) - G_{1/2}(\zeta) = \log(q_{ph}/c_p)$ with respect to the half-binding point $G_{1/2}(\zeta)$ (i.e., the trait value that has the IC50 value $\zeta$). Second, the error $q_e$, which is defined as $q_e = P_{\text{bind}}$ for ecological control ($\epsilon = +1$) and as $q_e = 1 - P_{\text{bind}}$ for evolutionary control ($\epsilon = -1$), is given by

$$q_e = \exp(-r_e) = \frac{c_p}{q_{ph}}; \qquad [24]$$

here, we have used the exponential asymptotic form of the binding probability (11). These characteristics determine the trait of the evolved state

$$\Gamma_e(\zeta) = \epsilon \log \zeta + \log \frac{q_{ph}}{c_p} \qquad [25]$$

and the resulting fitness $f_{p,e}(\zeta) = f_p(\Gamma_e(\zeta), \zeta)$. The fitness ranking (i.e., the evolutionary stability) of the evolved state and the wild type depends on the control amplitude $\zeta$ (stable/metastable pathogen states are shown as solid/dashed line segments in Fig. 2). The rank switch, which marks the loss of the evolved trait, occurs at a value $\zeta_I$ determined by the condition $f_{p,e}(\zeta_I) = f_p(\zeta_I, 0)$ (horizontal lines in Fig. 2). This determines the cross-over point

$$(\Gamma_I, \zeta_I^\epsilon) = \left( \frac{q_{ph}}{c_p} - 1, \frac{c_p}{q_{ph}} \exp\left( \frac{q_{ph}}{c_p} - 1 \right) \right). \qquad [26]$$

**Computational Control Equilibria.** We compute the point of optimal stationary control, $(\Gamma^*, \zeta^*)$, by evaluating the host payoff with the pathogen at its conditional stable equilibrium, $f_h(\Gamma^*(\zeta), \zeta)$, and maximizing this function with respect to $\zeta$. For ecological control in the SC regime, the pathogen is in the wild type $\Gamma_{wt} = 0$; in the WC regime, the pathogen is in the evolved state $\Gamma_e(\zeta)$ given by [25]. We obtain the equilibrium point

$$(\Gamma^*, \zeta^*) = \begin{cases} \left( 0, \frac{c_p}{q_{ph}} \exp\left( \frac{q_{ph}}{c_p} - 1 \right) \right) & (\text{SC}, c_h < c_h^\bullet), \\[2mm] \left( \log \frac{q_{ph}q_{hp}}{c_h}, \frac{q_{hp}c_p}{c_h} \right) & (\text{WC}, c_h > c_h^\bullet); \end{cases} \qquad [27]$$

this determines the pathogen fitness $f_p^* = f_p(\Gamma^*, \zeta^*)$, the host payoff $\psi_h^* = \psi_h(\Gamma^*, \zeta^*)$, and the control efficiency

$$\eta^* = \begin{cases} 1 - \frac{c_p c_h}{q_{ph}^2 q_{hp}} \exp\left( \frac{q_{ph}}{c_p} - 1 \right) & (\text{SC}, c_h < c_h^\bullet), \\[2mm] \frac{c_p}{q_{ph}} \log \frac{q_{ph}q_{hp}}{c_h} & (\text{WC}, c_h > c_h^\bullet) \end{cases} \qquad [28]$$

(Fig. 3 and *SI Appendix*, Fig. S1). The transition between the SC and WC control regimes is determined by the condition $\Gamma_e = \Gamma_I$, from which we obtain the transition line $c_h^\bullet(c_p, q_{ph}, q_{hp})$, Eq. 13.

For evolutionary control, the pathogen has an evolved pathogen trait $\Gamma_e(\zeta)$ in the SC regime and a wild-type trait $\Gamma_{wt} = 0$ in the WC regime. We obtain the equilibrium point

$$(\Gamma^*, \zeta^*) = \begin{cases} \left( \frac{q_{ph}}{c_p} - 1, \frac{q_{ph}}{c_p} \exp\left( 1 - \frac{q_{ph}}{c_p} \right) \right) & (\text{SC}, c_h < c_h^\bullet), \\[2mm] (0, 0) & (\text{WC}, c_h > c_h^\bullet). \end{cases} \qquad [29]$$

The optimal SC control is now at $(\Gamma^* = \Gamma_I, \zeta^* = \zeta_I)$, as given by [26], where the pathogen has the maximum stable $\Gamma_e$ value. This point has the pathogen load $L_p^* = q_e f_p^*$ with $q_e = (1 - P_{\text{bind}})$ given by [24], the host payoff $f_h^*$, and the control efficiency

$$\eta^* = \begin{cases} 1 - \frac{c_p}{q_{ph}} - \frac{c_h q_{ph}}{c_p q_{hp}} \exp\left( 1 - \frac{q_{ph}}{c_p} \right) & (\text{SC}, c_h < c_h^\bullet), \\[2mm] 0 & (\text{WC}, c_h > c_h^\bullet) \end{cases} \qquad [30]$$

(Fig. 3 and *SI Appendix*, Fig. S1). The transition between the SC and WC control regimes is determined by the condition $\eta_{SC}^* = 0$, from which we obtain the transition line $c_h^\bullet(c_p, q_{ph}, q_{hp})$, Eq. **14**.

**Dosage Protocols for Adaptive Trait Formation.** In the protocols described above, an initial-phase dosage $\zeta$ elicits a gain-of-function mutation to a trait value $G$, which is the starting point for the subsequent breeding phase. The extended HJB equation for these two-phase protocols determines the relative score

$$\Delta\Omega(G, \zeta) = \Delta\Omega_{\text{in}}(G, \zeta) + J_{\text{br}}(G) \qquad [31]$$

with an initial-phase component $\Delta\Omega_{\text{in}}(G, \zeta)$ and a breeding-phase component $J_{\text{br}}(G)$ given by [23]. To compute $\Delta\Omega_{\text{in}}(G, \zeta)$, we assume that trait-changing mutations have an exponential rate distribution, $u(G) \sim \mu \exp(-G/\epsilon_0)$. The selective effects of these mutations are given by the pathogen fitness landscape, Eq. **12**. Prior to the start of control, all mutants with $G > 0$ evolve under negative selection of strength $s_0(G) = f_p(G, 0) - f_{p,0} < 0$. Under control, gain-of-function mutants have a positive selection coefficient against the wild type, $s(G, \zeta) = f_p(G, \zeta) - f_{p,0} > 0$, leading to an adaptive fixation rate

$$v(G, \zeta) = 2N_e u(G) s(G, \zeta) = \theta \exp(-G/\epsilon_0) s(G, \zeta). \qquad [32]$$

Because $s$ and $s_0$ are of similar magnitude, the number of mutants with trait amplitude $G$ that are present at the start of control and destined for fixation under the initial-phase protocol follows a Poisson distribution with expectation value close to $2N_e u(G)$. Hence, the expected time to high frequency of such mutants is

$$T_{\text{in}}(G, \zeta) = \frac{1 - \exp[-2N_e u(G)]}{s(G, \zeta)} + \frac{\exp[-2N_e u(G)]}{v(G, \zeta)}. \qquad [33]$$

The two terms on the right-hand side describe the contributions of standing variation at the start of control and of de novo mutations originating under control, respectively. The resulting initial-phase score reads

$$\Delta\Omega_{\text{in}}(G, \zeta) = (-q_{hp} - c_h\zeta - \psi_h^* - \lambda) T_{\text{in}}(G, \zeta). \qquad [34]$$

We evaluate the diversity $\theta = 2N_e\mu$ in the initial phase with $N_e = \bar{N}_0 \equiv f_0/c$. Combining Eqs. **31**–**34** displays the scaling form of the score, Eq. **15**, where the scaling function $\omega$ describes the cross-over in prevalence between de novo mutations ($\theta \lesssim 1$) and standing variation ($\theta \gtrsim 1$) driving the initial gain-of-function evolution. The breeding phase has a related cross-over from periodic selection ($\theta \lesssim 1$) to clonal interference ($\theta \gtrsim 1$) (34), which sets the effective population size in that phase (*SI Appendix*). Using Eqs. **23** and **34** together with a maximum-likelihood approximation for the gain-of-function amplitude,

$$G^*(\zeta) = \arg\min_G T_{\text{in}}(G, \zeta), \qquad [35]$$

the HJB Eq. **31** determines the global maximum-score protocol (*SI Appendix*, Fig. S4). This protocol has a gain-of-function dynamics

$$(G_{\text{in}}, \zeta_{\text{in}}) = \arg\max_\zeta \Omega(G^*(\zeta), \zeta) \qquad [36]$$

and a breeding phase described by [22] (Fig. 4 *D–G* and *SI Appendix*, Figs. S2 and S3).

**Dosage Protocols for Metastable Control.** For the two-state protocol, we evaluate the relative score per unit of time as a function of the escape amplitude $G$ and baseline dosage $\zeta$ for $\lambda = 0$,

$$\Delta\omega(G, \zeta) = \Delta\omega_{\text{base}}(\zeta) + v(G, \zeta)\,\Delta J_{\text{res}}(G, \zeta). \qquad [37]$$

Here, $v(G, \zeta)$ denotes the $\theta$-dependent establishment rate of escape mutations given by [32], $\Delta\omega_{\text{base}}(\zeta) = \psi_h(0, \zeta) - \psi_h^*$ is the baseline score per unit of time, and $\Delta J_{\text{res}}(G) = J_{\text{res}}(G) - T_{\text{res}}(G)\Delta\omega_{\text{base}}(\zeta)$ is the negative excess payoff per rescue boost. The optimal score per rescue, $J_{\text{res}}(G)$, is computed in a simple approximation: we apply a constant dosage $\zeta_{\text{res}}(G) = \exp(G + 2r_e)$, which generates positive selection of strength $s_{\text{res}}(G) \approx c_p G_{\text{res}}$ of the wild type against the escape mutant. The boost dosage is maintained over a period

$$T_{\text{res}}(G) = a\,\frac{\log(2N_e x_{\text{esc}} s_{\text{res}}(G))}{s_{\text{res}}(G)}, \qquad [38]$$

which is proportional to the expected time to loss of the mutant ($x_{\text{esc}}$ is the detection threshold frequency; the constant $a$ determines the expected failure rate of rescue). The resulting rescue score is

$$J_{\text{res}}(G) = [\psi_h(0, \zeta_{\text{res}}(G)) - \psi_h^*]\,T_{\text{res}}(G). \qquad [39]$$

Using again a maximum-likelihood approximation for escape mutations,

$$G^*(\zeta) = \arg\max_G v(G, \zeta), \qquad [40]$$

we obtain the optimal metastable protocol by maximization of $\Delta\omega$, Eq. **37** (*SI Appendix*, Fig. S4). This protocol has a baseline dosage

$$\zeta_{\text{base}} = \arg\max_\zeta \Delta\omega(G^*(\zeta), \zeta), \qquad [41]$$

which determines the escape trait $G_{\text{esc}} = G^*(\zeta_{\text{base}})$ and the rescue dosage $\zeta_{\text{res}} = \exp(G_{\text{res}} + 2r_e)$ (Fig. 4 *H–J*). The maximum score per unit of time, $j = \max_\zeta \Delta\omega(G^*(\zeta), \zeta)$, sets the long-term efficiency gain of metastable control, $\eta_{\text{ms}} - \eta^* = j/(q_{ph} q_{hp})$ (Fig. 4*K*).

1. A. S. Perelson *et al.*, HIV-1 dynamics in vivo: Virion clearance rate, infected cell lifespan, and viral generation time. *Science* **271**, 1582–1586 (1996).
2. T. W. Schoener, The newest synthesis: Understanding ecological dynamics. *Science* **331**, 426–429 (2011).
3. S. P. Carroll *et al.*, Applying evolutionary biology to address global challenges. *Science* **346**, 1245993 (2014).
4. A. F. Read, T. Day, S. Huijben, The evolution of drug resistance and the curious orthodoxy of aggressive chemotherapy. *Proc. Natl. Acad. Sci.* **108**, 10871–10877 (2011).
5. R. A. Gatenby *et al.*, Adaptive therapy. *Cancer Res* **69**, 4894–4903 (2010).
6. J. Zhang, J. J. Cunningham, J. S. Brown, R. A. Gatenby, Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. *Nat. Commun.* **8**, 1816 (2017).
7. M. Łuksza, M. Lässig, A predictive fitness model for influenza. *Nature* **507**, 57–61 (2014).
8. D. H. Morris *et al.*, Predictive modeling of influenza shows the promise of applied evolutionary biology. *Trends Microbiol.* **26**, 102–118 (2017).
9. D. Hughes, D. I. Andersson, Evolutionary trajectories to antibiotic resistance. *Annu. Rev. Microbiol.* **71**, 579–596 (2017).
10. M. Łuksza *et al.*, A neoantigen fitness model predicts tumour response to checkpoint blockade immunotherapy. *Nature* **551**, 517–520 (2017).
11. S. Wang *et al.*, Manipulating the selection forces during affinity maturation to generate cross-reactive HIV antibodies. *Cell* **160**, 785–797 (2015).
12. A. Nourmohammad, J. Otwinowski, J. B. Plotkin, Host-pathogen coevolution and the emergence of broadly neutralizing antibodies in chronic infections. *PLoS Genet.* **12**, e1006171 (2016).
13. V. Sachdeva, K. Husain, J. Sheng, S. Wang, A. Murugan, Tuning environmental timescales to evolve and maintain generalists. *Proc. Nat. Acad. Sci. USA* **117**, 12693–12699 (2020).
14. R. F. Stengel, *Optimal Control and Estimation* (Dover Publication Inc., New York, NY, 1994).
15. H. J. Kappen, "An introduction to stochastic control theory, path integrals and reinforcement learning" in *AIP Conference Proceedings* (AIP, 2007), pp. 149–181.
16. M. Lässig, V. Mustonen, A. M. Walczak, Predicting evolution. *Nat. Ecol. Evol.* **1**, 77 (2017).
17. J. M. A. Blair, M. A. Webber, A. J. Baylay, D. O. Ogbolu, L. J. V. Piddock, Molecular mechanisms of antibiotic resistance. *Nat. Rev. Microbiol.* **13**, 42–51 (2015).
18. T. Held, D. Klemmer, M. Lässig, Survival of the simplest in microbial evolution. *Nat. Commun.* **10**, 2472 (2019).
19. R. A. Neher, O. Hallatschek, Genealogies of rapidly adapting populations. *Proc. Natl. Acad. Sci.* **110**, 437–442 (2013).
20. V. Mustonen, M. Lässig, Fitness flux and ubiquity of adaptive evolution. *Proc. Natl. Acad. Sci.* **107**, 4248–4253 (2010).
21. U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.* **75**, 126001 (2012).
22. Y. Iwasa, Free fitness that always increases in evolution. *J. Theor. Biol.* **135**, 265–281 (1988).

23. J. Maynard Smith, *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, UK, 1982).
24. M. Manfred Eigen, J. McCaskill, P. Schuster, Molecular quasi-species. *J. Phys. Chem.* **92**, 6881–6891 (1988).
25. M. Das Thakur *et al.*, Modelling vemurafenib resistance in melanoma reveals a strategy to forestall drug resistance. *Nature* **494**, 251–255 (2013).
26. R. R. Regoes *et al.*, Pharmacodynamic functions: A multiparameter approach to the design of antibiotic treatment regimens. *Antimicrob. Agents Chemother.* **48**, 3670–3676 (2004).
27. M. Scott, C. W. Gunderson, E. M. Mateescu, Z. Zhang, T. Hwa, Interdependence of cell growth and gene expression: Origins and consequences. *Science* **330**, 1099–1102 (2010).
28. M. Lynch, G. K. Marinov, The bioenergetic costs of a gene. *Proc. Natl. Acad. Sci.* **112**, 15690–15695 (2015).
29. A. Fischer, I. Vázquez-García, V. Mustonen, The value of monitoring to control evolving populations. *Proc. Natl. Acad. Sci.* **112**, 1007–1012 (2015).
30. A. Mayer, V. Balasubramanian, T. Mora, A. M. Walczak, How a well-adapted immune system is organized. *Proc. Natl. Acad. Sci.* **112**, 5950–5955 (2015).
31. M. Caskey *et al.*, Antibody 10-1074 suppresses viremia in HIV-1-infected individuals. *Nat. Med.* **23**, 185–191 (2017).
32. M. G. J. de Vos, M. Zagorski, A. McNally, T. Bollenbach, Interaction networks, ecological stability, and collective antibiotic tolerance in polymicrobial infections. *Proc. Natl. Acad. Sci.* **114**, 10666–10671 (2017).
33. J. F. C. Kingman, On the genealogy of large populations. *J. Appl. Probab.* **19**, 27–43 (1982).
34. M. Desai, D. Fisher, Beneficial mutation-selection balance and the effect of linkage on positive selection *Genetics* **176**, 1759–1798 (2007).