

Optimal Evolutionary Control for Artificial Selection on Molecular Phenotypes

Armita Nourmohammad¹*

*Department of Physics, University of Washington,
3910 15th Ave Northeast, Seattle, Washington 98195, USA,
Max Planck Institute for Dynamics and Self-organization, am Faßberg 17, 37077 Göttingen, Germany,
and Fred Hutchinson Cancer Research Center,
1100 Fairview Ave North, Seattle, Washington 98109, USA*

Ceyhun Eksin

*Department of Industrial and Systems Engineering, Texas A&M University,
College Station, Texas 77845, USA*



(Received 24 May 2020; revised 2 December 2020; accepted 11 January 2021; published 4 March 2021)

Controlling an evolving population is an important task in modern molecular genetics, including directed evolution for improving the activity of molecules and enzymes, in breeding experiments in animals and in plants, and in devising public health strategies to suppress evolving pathogens. An optimal intervention to direct evolution should be designed by considering its impact over an entire stochastic evolutionary trajectory that follows. As a result, a seemingly suboptimal intervention at a given time can be globally optimal as it can open opportunities for desirable actions in the future. Here, we propose a feedback control formalism to devise globally optimal artificial selection protocol to direct the evolution of molecular phenotypes. We show that artificial selection should be designed to counter evolutionary trade-offs among multivariate phenotypes to avoid undesirable outcomes in one phenotype by imposing selection on another. Control by artificial selection is challenged by our ability to predict molecular evolution. We develop an information theoretical framework and show that molecular timescales for evolution under natural selection can inform how to monitor a population in order to acquire sufficient predictive information for an effective intervention with artificial selection. Our formalism opens a new avenue for devising artificial selection methods for directed evolution of molecular functions.

DOI: [10.1103/PhysRevX.11.011044](https://doi.org/10.1103/PhysRevX.11.011044)

Subject Areas: Biological Physics, Statistical Physics

I. INTRODUCTION

The concept of feedback control in molecular evolution was first advocated by Darwin and Wallace as a way of describing natural selection [1]. Wallace hypothesized that similar to the centrifugal governor of the steam engine, the action of natural selection is like a *controller* that balances organismic traits, such that weak feet are often accompanied with powerful wings [1]. Such evolutionary trade-offs are ubiquitous in natural fitness landscapes. For example, experiments on a protein transport system have shown that the fitness landscape for the underlying biochemical

network is tuned to exploit optimal control with feedback throughout evolution [2]. However, it remains to be determined whether these structures are solely reflective of biochemical constraints or have emerged as incidences of fitness landscapes that could accommodate for long-term evolutionary survival.

Evolution as a feedback control is also reflected in the inheritance strategies and phenotypic response of populations to time-varying environments. A prominent example of such adaptation is observed in bacteria where cells use phenotypic switches to produce slowly replicating bacteria with tolerance and persistence against antibiotics. Populations use this Lamarckian-type phenotypic response [3] to hedge their bets against fluctuating environments [4,5]—an optimal response that can be viewed as an evolutionary feedback control [6].

Another approach to evolutionary control is through external interventions with artificial selection to direct populations to acquire a desired trait. Figure 1 demonstrates artificial selection with a feedback control to breed “pink cows,” which are otherwise not favored by natural selection. Such selective breeding has long been used to

*Corresponding author.
armita@uw.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

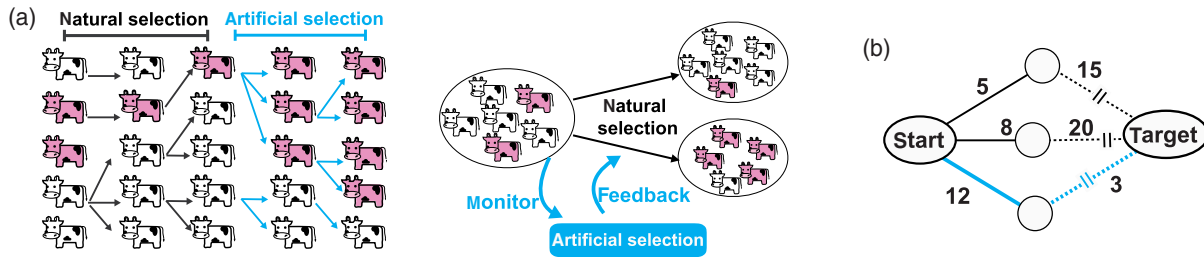


FIG. 1. Artificial selection as an optimal stochastic adaptive control strategy. (a) Artificial selection is an external intervention to select for a desired trait (i.e., pinkness of cows) in a population, which is otherwise not favored by natural selection. Artificial selection should be viewed as an optimal feedback control, whereby monitoring a stochastically evolving population informs the intervention protocol to optimally bias breeding and reproduction over generations. (b) The schematic graph shows different paths with indicated costs for a system to evolve from a start to a target state. Bellman's principle of optimality states that at each step an optimal decision is made, assuming that the following steps are also determined optimally. Although the first step (full line) of the blue path is more costly compared to the others (dotted lines), its cumulative cost is minimum, and hence, it should be chosen as the optimal path. This decision can be made best recursively, known algorithmically as dynamic programming.

domesticate animals or to improve agricultural yield in crops and became the basis for Mendelian genetics [7].

Another important avenue for artificial selection is to characterize intervention protocols against rapidly evolving pathogens, for example, to counter emergence of drug resistance in bacteria, escape of viral populations from immune challenge, or progression of evolving cancer tumors [8,9]. As such, control strategies have been suggested to direct inpatient somatic evolution of antibody secreting B cells to elicit potent broadly neutralizing antibodies against HIV [10–13]. Finding such control strategy involves optimization of the immunization cocktail, and the schedule for immunization, to direct somatic evolution of antibodies toward an outcome with high potency and breadth.

Artificial selection also plays a significant role in improving molecular functions through experimental directed evolution. Importantly, directed evolution in the lab is currently being employed to improve the activity and selectivity of molecules and enzymes [14–16], often desirable in industry or for pharmaceutical purposes. For example, experimental techniques like *morbidostat* have been designed to directly measure the growth rate of evolving microbial populations and accordingly tune the strength of selection induced by antibiotics in order to achieve continuous drug-induced inhibition in an experimental setup [17,18]. More recently, control protocols, such as proportional-integral-derivative (PID) control [19], have been experimentally implemented in high-throughput continuous directed evolution of proteins to automatically tune artificial selection based on the state of the population and optimize function [20–22]. Implementing optimal control in these automated and continuous directed evolution experiments will have significant impact in efficiently synthesizing molecules with desired function.

Designing any artificial selection protocol is limited by our ability to predict the outcome of evolution, which is often challenging due to a multitude of stochastic forces at

play, such as mutations, reproductive stochasticity (genetic drift), and environmental fluctuations [23,24]. In contrast to strongly divergent evolution at the genetic level, there is growing experimental evidence for convergent predictable evolution at the phenotypic level [17,25,26], including for complex molecular phenotypes like RNA polymerase function [27]. We will exploit this evolutionary predictability and focus on designing artificial selection for molecular phenotypes, which are key links between genotypic information, organismic functions, and evolutionary fitness [23].

Fitness and large-scale organismic traits are often encoded by a number of covarying molecular phenotypes, linked through genetic interactions; pigmentation patterns on the wings or body of fruit flies are among such multidimensional traits, impacted by the expression level of many interacting genes. A central issue in designing artificial selection for multivariate phenotypes is to avoid the undesirable (side) effects of selection, which can arise due to evolutionary trade-offs, e.g., between thermal stability and function of a protein [28,29]. Evolutionary interventions on multivariate phenotypes should be designed by assuming their impact over an entire evolutionary trajectory that follows. As a result, a locally optimal action at a given time point may be suboptimal once considering all the actions that are required to follow in order to direct the correlated evolution of the phenotypes toward their targets; see Fig. 1(b).

Finding a globally optimal protocol to direct a stochastic evolution is a topic of control theory, known for its impact in engineering, economics, and other fields [30]. Here, we introduce a formalism based on optimal control to devise a population-level artificial selection strategy and drive the stochastic evolution of multivariate molecular phenotypes toward a desired target. Importantly, we develop a framework to quantify how uncertainty and lack of evolutionary predictability can limit the efficacy of such artificial selection. By relating evolutionary predictability with control under artificial selection, we characterize how to

best monitor a population and acquire a sufficient predictive information in order to optimally intervene with its evolution.

II. RESULTS

A. Model of multivariate phenotypic evolution under optimal artificial selection

Molecular phenotypes are often polymorphic due to genetic variation in their encoding sequence within a population. Here, we primarily focus on phenotypes that are encoded by a relatively large number of genetic loci and, hence, are approximately normally distributed within a population—this Gaussian approximation, however, can be relaxed as discussed in Ref. [31]. In the case of normally distributed k -dimensional phenotypes, we characterize the population's phenotype statistics by the average $\mathbf{x} = [x_1, x_2, \dots, x_k]^\top$ and a symmetric covariance matrix K , where the diagonal elements $K_{ii}(\mathbf{x})$ indicate the variance of the i th phenotype and the off-diagonal entries $K_{ij}(\mathbf{x})$ indicate the covariance between different phenotypes. A nondiagonal covariance matrix reflects existence of mutational trade-offs between molecular phenotypes, for example, between the thermal stability and the catalytic activity of an enzyme or function of a protein [28,29], or between the affinity and the breadth (cross-reactivity) of an antibody [32,33].

The primary evolutionary forces that shape the composition of phenotypes within a population are selection, mutations, and genetic drift. Molecular phenotypes are often encoded in confined genomic regions of about a few 100 bps, and hence, are not strongly impacted by recombination, even in sexually reproducing populations. The impact of the evolutionary forces on phenotypes can be directly projected from the evolutionary dynamics in the high-dimensional space of the encoding genotypes [34,35]. For Gaussian distributed phenotypes, the change in mean phenotype $d\mathbf{x}$ over a short time interval dt simplifies to a stochastic process [36],

$$d\mathbf{x} = K \cdot \nabla F dt + \frac{1}{N} \Sigma \cdot d\mathbf{W}, \quad (1)$$

where F is the adaptive potential and ∇F is the corresponding adaptive force, reflecting the combined impact of natural selection and mutations during evolution. $d\mathbf{W}$ is a differential that reflects the stochastic effect of genetic drift by a multidimensional Wiener noise process [37]. The amplitude of the noise is proportional to Σ , which is the square root of the covariance matrix (i.e., $\Sigma^\top \Sigma \equiv K$), scaled by the effective population size N that adjusts the overall strength of the noise (see the Appendix A for details).

The stochastic evolution of the mean phenotype in Eq. (1) defines an ensemble of evolutionary trajectories. We can characterize the statistics of these evolutionary

paths by the dynamics of the underlying conditional probability density $P(\mathbf{x}', t' | \mathbf{x}, t)$ for a population to have a mean phenotype \mathbf{x}' at time t' , given its state \mathbf{x} at an earlier time $t < t'$. The dynamics of this probability density follows a high-dimensional Fokker-Planck equation [34]:

$$\frac{\partial}{\partial t} P(\mathbf{x}', t' | \mathbf{x}, t) = \left[\frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} - \nabla (K \cdot \nabla F) \right] P(\mathbf{x}', t' | \mathbf{x}, t). \quad (2)$$

Here, we measured time in units of effective population size ($t \rightarrow t/N$), which is the coalescence time in neutrality [38], and introduced the rescaled adaptive potential $NF \rightarrow F$. Moreover, we introduced the compact notation, $\text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} \equiv \sum_{ij} K_{ij} (\partial/\partial x_i) (\partial/\partial x_j)$.

Similar to the mean phenotype, the covariance matrix K is a time-dependent variable, impacted by evolution. However, fluctuations of covariance are much faster compared to the mean phenotype [31,34]. Moreover, even in the presence of moderately strong selection pressure, the phenotypic covariance depends only weakly on the strength of selection and is primarily determined by the supply of mutations in a population [34,35]. Therefore, we assume that the phenotypic covariance matrix remains approximately constant over time and equal to its stationary ensemble-averaged estimate throughout evolution (Appendix A).

B. Artificial selection to optimally direct evolution

Natural selection in Eq. (1) drives populations toward an optimum, which is a function of the organism's environmental and physiological constraints. Artificial selection aims to relax or tighten some of the natural constraints to drive evolution toward an alternative desired state \mathbf{x}^* . In general, we can formulate evolution subject to artificial selection as

$$d\mathbf{x} = \underbrace{(K \cdot \nabla F + \mathbf{u}(\mathbf{x}, t))}_{A(\mathbf{x}, t)} dt + \Sigma \cdot d\mathbf{W}, \quad (3)$$

where $\mathbf{u}(\mathbf{x}, t)$ is a time- and phenotype-dependent vector, which determines the impact of artificial selection and $A(\mathbf{x}, t)$ is the total force incurred by natural and artificial selection on the phenotypes.

Our goal is to find an optimal protocol for artificial selection $\mathbf{u}(\mathbf{x}, t)$ in order to reach the target \mathbf{x}^* by a desired time t_f , while minimizing the cost function,

$$\Omega(\mathbf{x}, \mathbf{u}, t) = V(\mathbf{x}, t) + \frac{1}{2} \mathbf{u}^\top B \mathbf{u}, \quad (4)$$

over an entire evolutionary trajectory. Here, $V(\mathbf{x}, t) \equiv V(|\mathbf{x}_t - \mathbf{x}^*|)$ is the cost for deviation of the phenotype state \mathbf{x}_t at time t from the desired target \mathbf{x}^* , and B is a matrix that characterizes the cost for imposing artificial

selection $\mathbf{u} \equiv \mathbf{u}(\mathbf{x}, t)$ and intervening with natural evolution of each phenotype. To solve the optimal control problem (i.e., to characterize an optimal artificial selection strategy), we define the *cost-to-go* function,

$$J(\mathbf{x}, t) = \min_{\mathbf{u}} \left\langle Q(\mathbf{x}, t_f) + \int_t^{t_f} ds \left(V(\mathbf{x}_s) + \frac{1}{2} \mathbf{u}_s^\top B \mathbf{u}_s \right) \right\rangle_{\text{evol}}, \quad (5)$$

where the angular brackets $\langle \cdot \rangle$ indicate expectation over stochastic evolutionary histories from time t until the target time t_f . Here, $Q(\mathbf{x}, t_f) \equiv Q(|\mathbf{x}_{t_f} - \mathbf{x}^*|)$ characterizes the cost of deviation from the target at the end of the evolutionary process t_f , which could be chosen to be different from the path cost $V(\mathbf{x})$.

An optimal artificial selection protocol should be designed by considering its impact over an entire stochastic evolutionary trajectory that follows. As a result, a seemingly suboptimal intervention at a given time can be globally optimal as it can open opportunities for more desirable actions in the future; see schematic Fig. 1(b). An optimal artificial selection protocol at each time point $\mathbf{u}^*(\mathbf{x}, t)$ assumes that the selection strategies implemented in the future are also optimal. This criterion is known as Bellman's "principle of optimality" [39], and allows us to express the optimal control problem in a recursive form, known as dynamic programming in computer science [39] (Appendix B). As a result, we can formulate a dynamical equation for the cost-to-go function, known as the Hamilton-Jacobi-Bellman (HJB) equation [40],

$$-\frac{\partial J(\mathbf{x}, t)}{\partial t} = \min_{\mathbf{u}} \left[\Omega(\mathbf{x}_t, \mathbf{u}_t) + A(\mathbf{x}_t) \cdot \nabla J + \frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} J \right], \quad (6)$$

with the boundary condition $J(\mathbf{x}, t_f) = Q(\mathbf{x}, t_f)$ at the end of the process (Appendix B). Importantly, the HJB equation (6) indicates that the cost to go $J(\mathbf{x}, t)$ is a potential function based on which the optimal artificial selection can be evaluated:

$$\mathbf{u}^*(\mathbf{x}, t) = -B^{-1} \cdot \nabla J(\mathbf{x}, t). \quad (7)$$

In other words, the cost-to-go function characterizes a time- and phenotype-dependent artificial fitness landscape that determines the strength of artificial selection $\mathbf{u}^*(\mathbf{x}, t)$.

The solution to the HJB equation (6) for the cost-to-go function $J(\mathbf{x}, t)$ and the artificial selection $\mathbf{u}^*(\mathbf{x}, t)$ can be complex time- and state-dependent functions, described by nonlinear evolutionary operators (Appendix B). Here, we consider a class of control problems, known as "path integral control" [41–43], where the cost matrix B for artificial intervention with evolution is inversely proportional to the phenotypic covariance K , i.e., $B = \lambda K^{-1}$,

where λ is a constant that determines the overall cost of artificial selection. This assumption implies that imposing artificial selection on highly conserved phenotypes is more costly than on variable phenotypes. This is intuitive as conserved phenotypes are often essential for viability of an organism and it is best to design a control cost function that limits the access to such phenotypes through artificial selection.

The path integral control assumption results in a significant mathematical simplification for the dynamics of the cost-to-go function $J(\mathbf{x}, t)$ and makes the inference of optimal artificial selection more tractable; see Appendixes B and C. We can characterize the evolution of the conditional distribution $P_u(\mathbf{x}', t' | \mathbf{x}, t)$ for a population under optimal artificial selection $\mathbf{u}^*(\mathbf{x}, t)$ to be in the phenotypic state \mathbf{x}' at time t' , given its state \mathbf{x} at time t , by

$$\begin{aligned} \frac{\partial}{\partial t} P_u(\mathbf{x}', t' | \mathbf{x}, t) &= \left[\frac{1}{2N} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} - \nabla(K \nabla F) - \frac{1}{\lambda} V(\mathbf{x}, t) \right] P_u(\mathbf{x}', t' | \mathbf{x}, t), \end{aligned} \quad (8)$$

with the initial condition $P_u(\mathbf{x}', t | \mathbf{x}, t) = \delta(\mathbf{x} - \mathbf{x}')$ (Appendix B). This conditional probability density can be used directly to compute the cost-to-go function $J(\mathbf{x}, t)$, and consequently the optimal control \mathbf{u}^* , as discussed in detail in Appendix B. Interestingly, the evolution of the optimally controlled conditional distribution $P_u(\mathbf{x}', t' | \mathbf{x}, t)$ resembles the natural evolutionary dynamics [Eq. (2) with $\mathbf{u} = 0$] with an extra annihilation term $V(\mathbf{x}, t)/\lambda$; see Appendix B and Ref. [42]. Therefore, artificial selection acts as an *importance sampling* protocol over each selection cycle (e.g., each generation) that removes (annihilates) individuals from the population with a rate proportional to their distance from the evolutionary target $\sim V(|\mathbf{x}_t - \mathbf{x}^*|)/\lambda$; see Fig. 2(a). Specifically, at each time point, this protocol generates a phenotypic distribution consistent with the evolutionary process under optimal artificial selection in Eq. (3) [Figs. 2(b) and 2(c)], without an explicit knowledge of the selection protocol $\mathbf{u}^*(\mathbf{x}, t)$. This result is highly practical for complex evolutionary processes, for which an analytical description of the optimal control protocol is inaccessible.

Although cost-to-go function $J(\mathbf{x}, t)$ and optimal control $\mathbf{u}^*(\mathbf{x}, t)$ are well-known concepts in the field of control theory, their connections to relevant evolutionary measures are far less explored. For evolutionary processes, the scaled cost-to-go-function $J(\mathbf{x}, t)/\lambda$ can be interpreted as a time- and phenotype-dependent fitness landscape associated with artificial selection $F_{\text{art}}(\mathbf{x}, t)$; see Eq. (7). Throughout an artificial selection process, populations evolve in an effective landscape $\hat{F}(\mathbf{x}, t) = F(\mathbf{x}) + F_{\text{art}}(\mathbf{x}, t)$, which is the superposition of the natural fitness landscape $F(\mathbf{x})$ and the

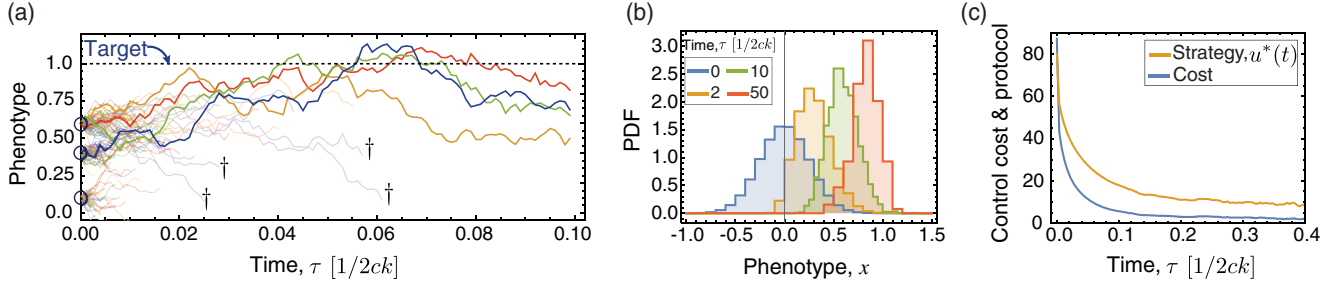


FIG. 2. Artificial selection with stochastic optimal annihilation. (a) Phenotypic trajectories starting from three distinct initial states (open circles) are shown for evolution under natural selection in a 1D quadratic adaptive landscape $F(x) = -cx^2$, where x is the phenotype centered around its optimum under natural selection. The trajectories are annihilated (\dagger and low opacity lines) with a rate proportional to the cost of their deviation from target throughout evolution (dotted line), $V(|x - x^*|)/\lambda$ [Eq. (8)]. At each time point, the survived trajectories characterize the ensemble of evolutionary states for evolution under optimal artificial selection to reach the target at $x^* = 1$. Control is designed under the assumption of infinite horizon. (b) The probability density function (PDF) shows the distribution of phenotypes for populations evolving subject to the annihilation protocol in (A). Populations start from an uncontrolled and natural state (blue distribution), and as a result of control annihilation, their distributions move toward the desired target at $x^* = 1$ (colors). (c) The expected control strategy $u^*(t) = -(1/\lambda)kg(x - x^*)$ (orange) and the cost of control $V(|x - x^*|) + \frac{1}{2}Bu^2$ (blue) are shown as a function of time, as populations are driven from their natural state toward the target state. Time is measured in units of the characteristic time for natural evolution ($1/2kc$). Parameters are $k = 0.4$, $c = 1$, $\lambda = 0.01$, $g = 4$.

artificial fitness landscape $F_{\text{art}}(\mathbf{x}, t)$. The overall scaling of the control cost λ determines the impact of artificial selection on evolution relative to natural selection, and when the control cost is small (i.e., $\lambda \ll 1$), artificial selection can dominate the course of evolution.

C. Artificial selection for multivariate phenotypes under stabilizing selection

Most of our analyses are applicable to general fitness and mutation (i.e., adaptive) landscapes (Appendixes B and C). However, we characterize in detail the features of artificial selection to direct evolution on high-dimensional quadratic adaptive landscapes ($F = -\mathbf{x}^\top \cdot C \cdot \mathbf{x}$), in which C is the adaptive pressure and \mathbf{x} is the shifted phenotype vector centered around the optimum under natural selection that the population approaches in stationary state [44]; see Appendix A. In addition, we assume a quadratic form for the cost function throughout the evolutionary process, $V(\mathbf{x}, t) = \frac{1}{2}(\mathbf{x}_t - \mathbf{x}^*)^\top G(\mathbf{x}_t - \mathbf{x}^*)$, and also at the end point, $Q(\mathbf{x}, t_f) = \frac{1}{2}(\mathbf{x}_{t_f} - \mathbf{x}^*)^\top \tilde{G}(\mathbf{x}_{t_f} - \mathbf{x}^*)$.

Characterizing an artificial selection protocol under such quadratic constraints falls within the class of standard stochastic control problems, known as linear-quadratic-Gaussian (LQG) control [30]. However, we will present our analyses based on the path integral control approach in Eq. (8), which is generalizable beyond LQG and can be applied to arbitrary cost functions and fitness landscapes (see Appendixes C and D for detailed derivation).

Let us imagine that our criterion is to drive evolution toward the optimum \mathbf{x}^* by time t_f , which implies that the path cost is zero $G = 0$ but the end point cost is nonzero $\tilde{G} > 0$; see Appendixes C and D for the general scenario including the case with $G > 0$. As time approaches the end

point, populations transition from evolving in their natural landscape $F(\mathbf{x})$ to the artificially induced fitness landscape $F_{\text{art}}(\mathbf{x}, t_f)$; see Fig. S1 in Supplemental Material [45] and Fig. 3(c) herein. Moreover, toward the end point, the fitness peak and the strength of selection approach their final values, determined by the target and the cost functions in Eq. (B3), in an exponentially rapid manner (Appendix D and Fig. S2 [45]). Interestingly, at the end point, the optimal artificial selection keeps the population close to the target with a strength,

$$\mathbf{u}^*(\tau \rightarrow 0) = -\frac{1}{\lambda}K\tilde{G}(\mathbf{x} - \mathbf{x}^*), \quad (9)$$

which resembles the breeder's equation [46] for artificial selection with a heritability factor, $h^2 = K\tilde{G}/\lambda$; see Appendix C for derivations and the general scenario including the case with $G > 0$.

One of the main issues in designing breeding experiments in plants and animals is the undesirable (side) effects of artificial selection on covarying phenotypes, primarily due to evolutionary trade-offs [47], e.g., between sturdiness and flavor of vegetables like tomatoes [48]. Similarly, trade-offs among covarying molecular phenotypes (e.g., function versus thermal stability of a protein) could lead to undesirable outcomes for artificial selection at the molecular level.

To demonstrate the consequences of phenotypic covariation, let us consider a simple example for artificial selection on two covarying phenotypes (x, y); the general solution to this problem in high dimensions is discussed in Appendix C. We aim to drive the phenotype x toward the target $x^* > 0$ by artificial selection while keeping the phenotype y at its stationary state value $y^* = 0$. An optimal

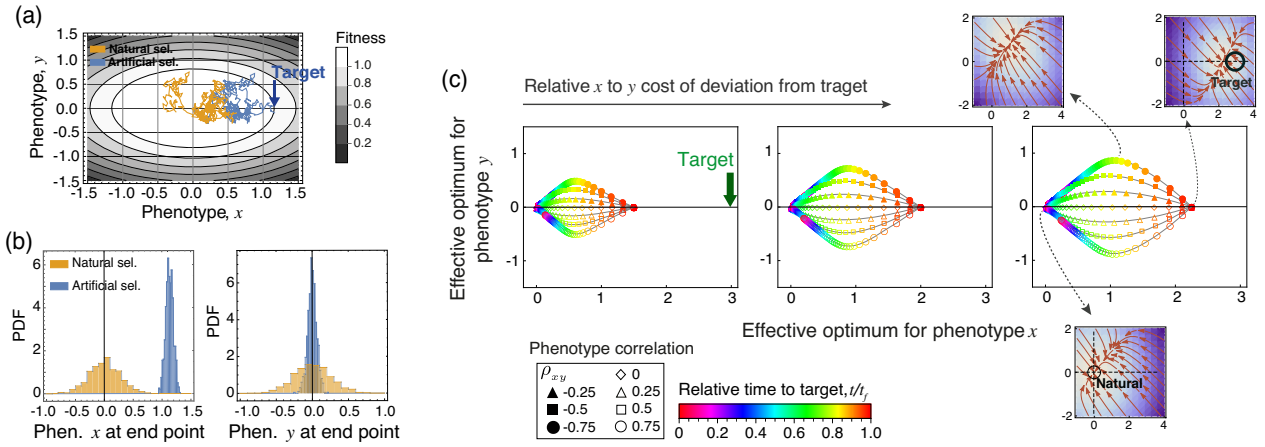


FIG. 3. Artificial selection on covarying phenotypes. (a) Trajectories for evolution under natural (orange) and artificial (blue) selection are shown for a 2D phenotype (x, y) , in a quadratic landscape. Parameters are $c_x = 2$, $c_y = 4$, $c_{xy} = 0$; $x^* = 1.2$, $y^* = 0$; $k_x = 0.02$; $k_y = 0.05$; $k_{xy} = 0$; $g_x = g_y = 2$; $\lambda = 0.01$. (b) The distribution of phenotypes at the end point of an artificial selection protocol (blue) is compared to the phenotypic distribution under natural selection (orange). Evolutionary parameters are the same as in (a). (c) The dynamics of the effective fitness peak is shown over time (colors) for 2D covarying phenotypes with correlations ρ_{xy} indicated by the shape of the markers. From left to right, panels show increasing end point cost of deviation from the target along the x phenotype, $g_x = 1, 2, 3$, with $g_y = 2$. Heat maps show the effective fitness landscapes associated with a specific fitness peak (indicated by the dotted arrows) for anticorrelated phenotypes at three different time points. The direction and length of the red arrows in each heat map indicate the direction and the strength of selection pressure toward the effective optimum. Parameters are $x^* = 3$, $y^* = 0$; $c_x = c_y = 5$, $c_{xy} = 0$; $k_x = k_y = 0.02$; $\lambda = 0.1$.

artificial selection protocol defines an effective two-dimensional quadratic fitness landscape that biases the evolutionary trajectories toward the target state [Fig. 3(a)]. As a result, the phenotype distributions at the end of the process become significantly distinct from the expectation under natural selection, and remain strongly restricted around their target values; see Fig. 3(b).

The peak of this fitness landscape (i.e., the effective optimum) changes from the natural state $(0, 0)$ to the target state (x^*, y^*) by the end of the selection process; see Fig. 3(c) and Fig. S3 [45]. The fitness peak moves monotonically along the x phenotype from the natural optimum 0 toward the target x^* , irrespective of the correlation ρ_{xy} between the two phenotypes; see Fig. 3(c). However, the dynamics of the fitness peak along the y phenotype is generally nonmonotonic and strongly dependent on the phenotypic correlation ρ_{xy} . An increase in x drives the positively (negatively) correlated y toward higher (lower) values. Therefore, in the beginning of the process, the optimal artificial selection protocol sets the fitness peak for the y phenotype at an opposite direction to counterbalance the effect of evolutionary forces due to phenotypic covariation. As the end point approaches, artificial selection becomes significantly strong with an effective fitness optimum set at the target for each phenotype x^* and y^* [Eq. (9)]. Therefore, the optimum y value should return to its target state ($y^* = 0$), resulting a nonmonotonic behavior in the dynamics of the fitness peak along the y phenotype; see Fig. 3(c). Moreover, the strength of selection changes over time and becomes stronger toward the target

phenotypes at the end point [heat maps in Fig. 3(c) and Fig. S4 [45]].

The optimal artificial selection protocol in Fig. 3 is sensitive to the structure of the phenotypic covariance matrix K [Eq. (9) and Appendix E]. Importantly, disregarding the phenotypic covariance in designing a control protocol would result in an increase in the associated cost and failure to reach the desired phenotype targets, as the controller misjudges the response of the covarying phenotypes to the designed interventions (Fig. S2 [45] and Appendix E).

The optimal artificial selection protocol requires an accurate description of the evolutionary dynamics, which may not be available in certain scenarios. In such settings, we can devise a naive but intuitive control strategy that is informed by the structure of the optimal control. For instance, an optimal controller with quadratic cost on a quadratic adaptive landscape (LQG) is proportional to the difference between the phenotypic state and the target $(\mathbf{x}_t - \mathbf{x}^*)$ with a prefactor (i.e., a relative strength) that exponentially increases as the time approaches the end point (Appendix D and Fig. S2 [45]). Using the general structure of this optimal strategy, we can devise proportional controllers with exponentially growing weights even when the accurate model for the underlying dynamics is not available (Appendix E). The protocol inspired by the optimal control outperforms a completely naive proportional controller that only intervenes with the evolutionary dynamics as the deviation of the population from the target passes a threshold but does not tune its relative strength as

time approaches the end point; see Fig. S5 [45] and Appendix E. However, it should be noted that the parameters of the exponential-proportional controller (e.g., the exponential weights) should be tuned to achieve the desired target states which may be time-consuming in practice. Moreover, tuning these weights according to certain objectives, e.g., achieving the desired phenotype traits, may have undesirable consequences such as high variability or large cost of control (Fig. S5 [45] and Appendix E).

D. Artificial selection with intermittent monitoring

Imposing artificial selection based on continuous feedback from the state of the population (Fig. 1) requires complete monitoring and the possibility of continuous evolutionary intervention—a criterion that is often not met in real conditions. In general, discrete selection protocols based on limited information can be inefficient for devising evolutionary feedback control [8,49]. Here, we develop a principled approach to characterize the limits of discrete optimal interventions based on the evolutionary response of the population to artificial selection. We consider a simple scenario where in a stationary state we aim to keep a population at the target phenotype \mathbf{x}^* , using discrete monitoring and interventions at time points ($i = 1, \dots, M$) with a time separation $\tau \equiv t_{i+1} - t_i$. We define a stationary cost-to-go function,

$$J(\mathbf{x}, t_m; \tau) = \min_{\mathbf{u}} \lim_{M \rightarrow \infty} \frac{1}{(M-m)\tau} \left\langle \sum_{i=m}^M \mathbf{u}_i^\top B \mathbf{u}_i + \int_{t_i}^{t_m} V(\mathbf{x}_t) dt \right\rangle_{\text{evol}}, \quad (10)$$

where the division by the total time $M\tau$ assures that the cost to go remains finite. To further simplify, we only consider one-dimensional phenotype x with intrapopulation variance k , the cost of deviation $V(x) = g(x - x^*)^2/2$ from target x^* , and the cost of intervention $\beta u^2/2$ with artificial selection u . However, our analyses can be easily generalized to multivariate phenotypes.

In the stationary state and in the regime of small perturbations ($gk/\lambda \ll 1$), the optimal artificial selection protocol u^* should be a variant of the case with full information with a strength of selection α_τ dependent on the time window τ , $u_\tau^* = -k\alpha_\tau(x - x^*)$; see Appendix F. We can characterize the optimal strength of artificial selection α_τ by minimizing the cost-to-go function in Eq. (10),

$$\alpha_\tau = \alpha_0 \left[\frac{(1 - e^{-2\tau}) + 8c(x^*)^2(1 - e^{-\tau})}{2\tau(1 + 4c(x^*)^2)} \right] + \mathcal{O}[(k\gamma/\lambda)^2], \quad (11)$$

where $\alpha_0 = g/\lambda$ is the optimal selection strength under continuous monitoring. Here, time τ is measured in units of

the characteristic time for evolution under natural selection, i.e., $(2kc)^{-1}$.

The partially informed artificial selection α_τ depends on most of the evolutionary parameters similar to selection with complete monitoring α_0 . Importantly, the ratio α_τ/α_0 depends only weakly on the strength of natural selection c [Fig. 4(a), top] and the target for artificial selection x^* [Fig. 4(a), bottom] and it is insensitive to the phenotypic diversity k and the parameter λ [Eq. (11)].

However, the optimal artificial selection α_τ strongly depends on the time interval τ and it decays as the time interval τ broadens [Fig. 4(a)]. This decay is linear and relatively slow up to the characteristic time for evolution under natural selection $(2kc)^{-1}$. This is the timescale over which an intermittent artificial selection can still contain the population around the desired target x^* . If interventions are separated further in time (i.e., $\tau \gg 1$), the optimal selection strength decays rapidly as $\sim \tau^{-1}$. Imposing a strong artificial selection in this regime is highly ineffective as populations can easily drift away from the target and toward their natural state within each time interval, and any artificial selection would only contribute to the intervention cost $\sim u^2$ without offering any benefits.

E. Information cost for artificial selection

Artificial selection is an evolutionary drive that shifts the equilibrium state of the population under natural selection to a new state around the target. As a result, the phenotypic histories $\mathbf{x}_{t_0, \dots, t_f}$ over the period of (t_0, \dots, t_f) are statistically distinct for evolution under natural and artificial selection [Fig. 3(a)]. This deviation can be quantified by the Kullback-Leibler distance $D_{\text{KL}}(\mathcal{P}_u(\mathbf{x}) \parallel \mathcal{P}(\mathbf{x}))$ between the distribution of histories under artificial selection $\mathcal{P}_u(\mathbf{x}) \equiv \mathcal{P}_u(\mathbf{x}_{t_0, \dots, t_f})$ and under natural selection $\mathcal{P}(\mathbf{x})$. In the stationary state, the Kullback-Leibler distance quantifies the impact of artificial selection on evolution and can be interpreted as the amount of work $W_{t_0}^{t_f}(\mathbf{u})$ done by external forces [50] (i.e., the artificial selection) to shift the population from its natural equilibrium to the artificial equilibrium,

$$W_{t_0}^{t_f}(\mathbf{u}) = D_{\text{KL}}(\mathcal{P}_u(\mathbf{x}) \parallel \mathcal{P}(\mathbf{x})) = \int d\mathbf{x}_{t_0}^{t_f} \mathcal{P}_u(\mathbf{x}) \log \left[\frac{\mathcal{P}_u(\mathbf{x})}{\mathcal{P}(\mathbf{x})} \right]. \quad (12)$$

The cumulative work is related to the cost of artificial selection, and for the case of path integral control, it is equivalent to the cumulative cost of control $W_{t_0}^{t_f}(\mathbf{u}) = \langle \frac{1}{2} \int \mathbf{u}^\top K^{-1} \mathbf{u} dt \rangle = \frac{1}{2\lambda} \langle \int \mathbf{u}^\top B \mathbf{u} dt \rangle$, where the angular brackets $\langle \cdot \rangle$ denote expectation over the ensemble of evolutionary histories under artificial selection; see Refs. [51,52] and Appendix G. The power (i.e., work per unit time), associated with intermittent artificial

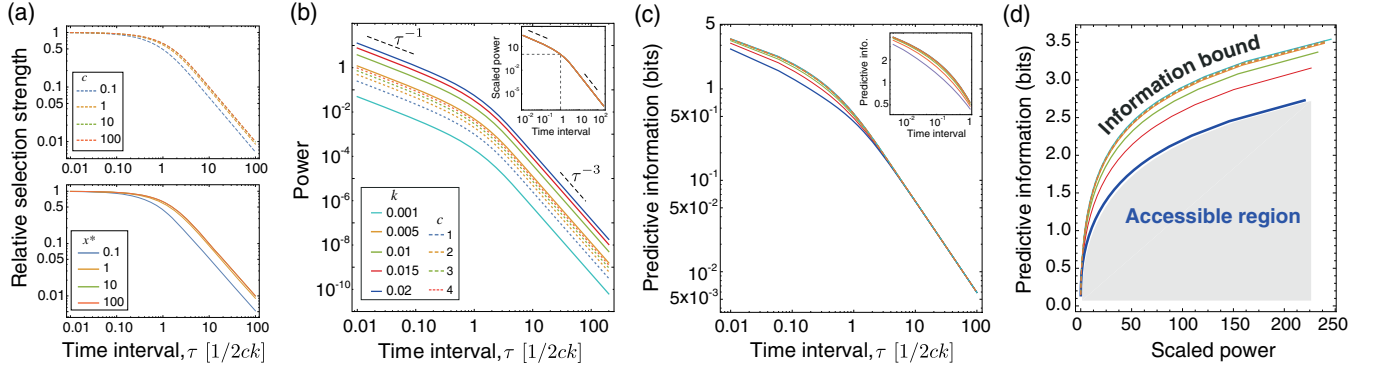


FIG. 4. Artificial selection with limited information. (a) Relative strength of artificial selection α_τ/α_0 [Eq. (11)] is shown as a function of the time interval for monitoring and intervention τ , measured in units of the characteristic time for evolution under natural selection ($1/2kc$). The selection ratio is shown for various strengths of natural selection c (top, with $x^* = 1$) and for various targets of artificial selection x^* (bottom, with $c = 1$). (b) Power [Eq. (13)] is shown as a function of the time interval τ for a range of parameters for the phenotypic diversity k (full line) and the strength of natural selection c (dotted line). The inset shows a collapse plot for power scaled with the expectation at the characteristic time for natural selection $\text{power}(\tau)/\text{power}(\tau = 1)$. (c) Predictive mutual information $\mathcal{I}(\tau)$ [Eq. (14)] is shown to decay with the time interval τ for a wide range of parameters (k, c). Inset shows an enlargement of a narrower range for the time interval $\tau < 1$. (d) Predictive information [Eq. (14)] is shown as a function of the scaled power for optimized artificial selection for a range of τ values [Eq. (13)]. Each curve sets an information bound for artificial selection for a given set of evolutionary parameters (k, c). A nonoptimal selection intervention should lie below the information curve, shown by the gray shaded area as the accessible region associated with the dark blue curve. Color code in (c) and (d) is similar to (b). Other parameters are $\lambda = 0.6$, $x^* = 3$, $g = 2$.

selection, can be expressed as the amount of work per time interval τ :

$$\text{power}(\tau) = \lim_{M \rightarrow \infty} \frac{1}{M\tau} \sum_{i=1}^M W(t_i) = \frac{1}{2\tau} \langle \mathbf{u}_\tau^\top K^{-1} \mathbf{u}_\tau \rangle. \quad (13)$$

The expected work, and hence the power, depend on the time interval τ through various factors. Work depends quadratically on the strength of artificial selection α_τ and on the expected population's deviation from the target $\langle (x_\tau - x^*)^2 \rangle$. On the one hand, the optimal strength of artificial selection α_τ decays with increasing the time interval; see Fig. 4(a) and Eq. (11). On the other hand, as the time interval broadens, populations deviate from the target toward their natural state, resulting in an increase in the expected work by artificial selection. Moreover, since interventions are done once per cycle, the power has an overall scaling with the inverse of the cycle interval $\sim \tau^{-1}$. These factors together result in a reduction of the expected power associated with artificial selection as the time interval widens; see Fig. 4(b).

Power depends strongly on the parameters of natural evolution including the strength of natural selection (c) and the phenotypic diversity within a population (k); see Fig. 4(b). This is due to the fact that steering evolution under strong natural selection (i.e., with large k and c) is more difficult and would require a larger power by artificial selection. However, the dependence of power on the evolutionary parameters (k, c) remains approximately separated from its dependence on the time interval τ .

Thus, power rescaled by its expectation at the characteristic time $\tau = 1$ shows a universal time-decaying behavior, independent of the evolutionary parameters [Fig. 4(b)].

F. Predictive information as a limit for efficient artificial selection

Artificial selection can only be effective to the extent that an intervention is predictive of the state of the population in the future. The mutual information between artificial selection and the future state of the population quantifies the amount of predictive information [53] by artificial selection, or alternatively, the memory of the population from the selection intervention. We characterize the predictive information I_τ as a time-averaged mutual information $I(x_t, x_0)$ between an intervention (at time $t = 0$) and the state of the population at a later time $t(0 < t < \tau)$, during each intervention cycle in the stationary state,

$$\begin{aligned} \mathcal{I}_\tau &= \frac{1}{\tau} \int_0^\tau dt I(x_t, x_0) \\ &= \frac{1}{\tau} \int dt \int dx_0 dx_t P(x_t, x_0) \log \left[\frac{P(x_t | x_0)}{P(x_t)} \right]. \end{aligned} \quad (14)$$

The predictive mutual information monotonically decreases as the time interval τ increases and the population evolves away from the selection target; see Fig. 4(c).

Predictive information in Eq. (14) quantifies the impact of artificial selection on the future of a population. The information theoretical measure of power in Eq. (13), on the other hand, quantifies how the optimal artificial

selection protocol distinguishes a population's past evolutionary history from the expectation under natural selection. The efficacy of any intervention (i.e., power) is tightly bounded by the impact it may have on the state of the population in the future (i.e., predictive information); see Fig. 4(d). Any nonoptimal artificial selection effort should lie below this bound and within the accessible region of the information-power plane [Fig. 4(d)].

Phenotypic diversity k characterizes the rate at which a population evolves away from the target and toward its natural state during an intervention interval [Eq. (1)]. As a result, the information bound for artificial selection is tighter in more diverse populations, which can rapidly evolve away from the target and toward their natural state during each time interval τ .

As interventions become more frequent, predictive mutual information increases but more slowly than the amount of power necessary to induce an effective artificial selection [Fig. 4(d)]. Importantly, the gain in predictive information becomes much less efficient for time intervals shorter than the characteristic time of natural selection ($\tau \ll 1$).

We postulate that trading power with information provides a guideline for scheduling of control interventions of stochastic processes in general, and for evolutionary control, in particular. The characteristic time for evolution under natural selection is a powerful gauge for scheduling the interventions. Importantly, setting the time interval within the range of the characteristic evolutionary time $\tau \sim 1$ could provide a sufficient power-to-information ratio for an optimal artificial selection protocol. However, as information becomes less predictive or the inferred selection protocol becomes suboptimal, it would be necessary to monitor and intervene more frequently.

Predictive information quantifies how the state of the system in the past is informative of its future, whereas the control power measures the cost associated with an intervention due to its impact on the future state of the system. The connection between predictive information and control in the context of directed evolution relates the past and the future of an evolutionary process, subject to external interventions. Indeed, predictive information sets the limit for an effective control in general stochastic processes, but the interpretation of power and predictive information would be specific to the problem in hand.

III. DISCUSSION

An optimal intervention should be designed by considering its impact over an entire evolutionary trajectory that follows. Here, we infer an artificial selection strategy as an optimal control with feedback to drive multivariate molecular phenotypes toward a desired target. This selection protocol is optimal over the defined time course and may seem suboptimal on short timescales as it drives one phenotype away from its target while driving another

toward the target to overcome trade-offs [Fig. 3(c)]. Monitoring and feedback from the state of a population are key for imposing an effective artificial selection strategy. We show that the schedule for monitoring should be informed by the molecular timescales of evolution under natural selection, which set the rate at which a population loses memory of artificial interventions by evolving toward its favorable state under natural selection.

Being able to control evolution could have significant applications in designing novel molecular functions or in suppressing the emergence of undesirable resistant strains of pathogens or cancers. Interventions that select for desired phenotypes have become possible in molecular engineering [22,54,55], in targeted immune-based therapies against evolving pathogens [56], and in immunogen design for optimal vaccination protocols against rapidly evolving viruses like HIV [10–13].

One class of experiments for artificial selection uses targeted mutations that are inferred to be beneficial, using machine learning techniques to characterize genotype-phenotype maps [57]. Another class of experiments relies on implementing feedback control to tune artificial selection during continuous evolution of molecules and proteins, to direct them toward a desired target [17,18,20–22]. For example, implementing proportional-integral-derivative control [19], which is known for its role in cruise control during driving, has shown significant improvements for molecular optimization by directed continuous evolution [22]. PID control is simple to implement in practice and it is relatively robust to errors, as its proportional term corrects for spontaneous error, the integrator reduces the impact of long run error, and derivative term would suppress overshooting by anticipating the impact of control [30]. Still, PID is far from optimal and requires fine-tuning of control parameters for the system, which is often tedious and subject to uncertainty. Therefore, implementing a more principled optimal control approach for molecular evolution, such as the path integral control introduced here, can further optimize the continuous directed evolution experiments.

The efficacy of these actions is limited by our ability to monitor and predict the evolutionary dynamics in response to interventions. Evolution is shaped by a multitude of stochastic effects, including the stochasticity in the rise of novel beneficial mutations and fluctuations in the environment, which at best raise skepticism about predicting evolution [23,24]. However, evolutionary predictability is not an absolute concept and it depends strongly on the time window and the molecular features that we are interested in. For example, despite a rapid evolutionary turnover in the influenza virus, a number of studies have successfully forecasted the dominant circulating strains for a one year period [58,59]. Similarly, phenotypic convergence across parallel evolving populations has been reported as an evidence for phenotypic predictability, despite a wide genotypic divergence

[17,25–27]. Therefore, to exploit the evolutionary predictability for the purpose of control, it is essential to identify the relevant degrees of freedom (e.g., phenotypes versus genotypes) that confer predictive information and to characterize the evolutionary timescales over which our observations from a population can inform our interventions to drive the future evolution.

We focus on modeling control of molecular phenotypes. Phenotypic diversity within a population provides standing variation that selection can act upon. To allow for a continuous impact of artificial selection over many generations, we have limited our analyses to a regime of moderate time- and phenotype-dependent artificial selection to sustain the phenotypic variability in a population. However, it would be interesting to allow for stronger artificial selection to significantly impact the standing variation and the structure of the phenotypic covariance within a population over time. Indeed, allowing a population to rapidly collapse as it approaches a desired target is a common strategy in evolutionary optimization algorithms [60]—a strategy that could accelerate the design of new functions with directed molecular evolution.

In this work, we assume a stochastic model for evolution of multivariate molecular phenotypes, which has been powerful in describing a range of biological systems, including the evolution of gene expression levels [61]. Indeed, optimal control protocols are often designed by assuming a specific model for the underlying dynamics. However, in most biological systems, we lack a knowledge of the details and the relevant parameters of the underlying evolutionary process. Optimal control strategies can inform *ad hoc* (albeit suboptimal) control approaches to drive evolution toward the desired target (Fig. S5 [45] and Appendix E). In addition, if one can at least approximately design a control scenario that satisfies the criteria for path integral control, an effective artificial selection protocol could be achieved through annihilation of evolutionary trajectories with a rate proportional to the cost of their deviation from the desired target (Fig. 2), and without a detailed knowledge of the underlying dynamics. Nonetheless, the ultimate goal is to simultaneously infer an effective evolutionary model based on the accumulating observations and to design an optimal intervention to control the future evolution—an involved optimization problem known as dual adaptive control [62].

ACKNOWLEDGMENTS

We are thankful to Bert Kappen, Jakub Otwinowski, and Olivier Rivoire for discussions and providing valuable feedback. This work has been supported by Royalty Research fund from the University of Washington (A. N.), the DFG grant (No. SFB1310) for Predictability in Evolution (A. N., C. E.), the MPRG funding through the Max Planck Society (A. N.), and the National Science Foundation Grant No. NSF ECCS-1953694 (C. E.).

APPENDIX A: EVOLUTION OF MULTIVARIATE MOLECULAR PHENOTYPES

In the case of normally distributed k -dimensional phenotypes, we characterize the population's phenotype statistics by the average $\mathbf{x} = [x_1, x_2, \dots, x_k]^\top$ and a symmetric covariance matrix K , where the diagonal elements $K_{ii}(\mathbf{x})$ indicate the variance of the i th phenotype and the off-diagonal entries $K_{ij}(\mathbf{x})$ indicate the covariance between different phenotypes. To model the evolution of such multivariate phenotype, we consider the three primary evolutionary forces: natural selection, mutations, and genetic drift. The effect of selection on the mean phenotype is proportional to the covariance between fitness and phenotype within a population [63]. For Gaussian distributed phenotypes, the change in mean phenotype $d\mathbf{x}$ over a short time interval dt simplifies to a stochastic process [36],

$$d\mathbf{x} = (K \cdot \nabla F + \nabla M)dt + \frac{1}{N}\Sigma \cdot d\mathbf{W}, \quad (\text{A1})$$

where F and M are fitness and mutation potentials, respectively. The gradient functions (denoted by ∇F and ∇M) determine the forces acting on the phenotypes by selection and mutation, respectively [34]. $d\mathbf{W}$ is a differential that reflects the stochastic effect of genetic drift by a multidimensional Wiener noise process [37]. The amplitude of the noise is proportional to Σ , which is the square root of the covariance matrix (i.e., $\Sigma^\top \Sigma \equiv K$), scaled by the effective population size N that adjusts the overall strength of the noise. In other words, the fluctuations of the mean phenotype across realizations of an evolutionary process is proportional to the intrapopulation variance K and inversely scales with the effective population size (i.e., the sample size) N .

The fitness potential F can be simply approximated by the mean fitness of a population [34]. The mutational potential M , however, can depend on the underlying genotype-phenotype map. To characterize M , let us assume a general map from a genotypic sequence with ℓ loci $\vec{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_\ell)$ and the encoded k -dimensional phenotype $\vec{E}(\vec{\sigma})$:

$$E^\alpha = \underbrace{\sum_i J_i^\alpha \sigma_i}_{E_{(1)}^\alpha} + \underbrace{\sum_{i \neq j} J_{ij}^\alpha \sigma_i \sigma_j}_{E_{(2)}^\alpha} + \dots + E_{(r_{\max})}^\alpha, \quad (\text{A2})$$

where $\alpha = 1, \dots, k$ refers to the α coordinate of the k -dimensional phenotype \vec{E} and the subscript $r = 1, 2, \dots, r_{\max}$ is the degree of genetic interaction. Without a loss of generality we assume that genotypic loci are biallelic with $\sigma_i \in \{-1, 1\}$ indicating the state of locus i .

Mutations occur with a rate μ per site per generation, resulting in a site flip: $-1 \rightleftharpoons 1$. Thus, the change in the α component of the phenotype due to mutations follows

$$\begin{aligned}
\Delta E^\alpha &= \sum_{r=1}^{r_{\max}} \Delta E_{(r)}^\alpha \\
&= -2\mu \sum_{i_1} J_{i_1}^\alpha \sigma_{i_1} - 4\mu \sum_{i_1, i_2} J_{i_1 i_2}^\alpha \sigma_{i_1} \sigma_{i_2} + \dots \\
&\quad - 2r\mu \sum_{i_1, \dots, i_{r_{\max}}} J_{i_1 \dots i_{r_{\max}}}^\alpha \sigma_{i_1} \dots \sigma_{i_{r_{\max}}} + \mathcal{O}(\mu^2) \\
&= -2\mu \sum_r r E_{(r)}^\alpha + \mathcal{O}(\mu^2), \tag{A3}
\end{aligned}$$

where we assumed that the mutation rate per locus per generation is small such that we can neglect terms of order μ^2 and higher. Similarly, the change in the mean α phenotype $\mathbf{x}^\alpha = \overline{E}^\alpha$ due to mutations up to $\mathcal{O}(\mu^2)$ follows $\Delta \mathbf{x}^\alpha = -2\mu \sum_r r \mathbf{x}_{(r)}^\alpha$, where $\mathbf{x}_{(r)}^\alpha = \overline{E}_{(r)}^\alpha$ is the intrapopulation mean of the α component of the phenotype. In this case, we can define a mutation potential,

$$M = -K^{-1} \mu \sum_{\alpha=1}^k \left(\sum_{r=1}^{r_{\max}} r \mathbf{x}_{(r)}^\alpha \right)^2, \tag{A4}$$

such that the change in phenotype due to mutations follows

$$d\mathbf{x} = K \cdot \nabla M. \tag{A5}$$

If the magnitude of mutation rate per locus per generation μ is small such that double mutations are significantly less likely than single mutations, a mutational potential function can be defined for phenotypic evolution. Dominance of double (or higher order) mutations results in mutational curls [64,65] and, thus, a breakdown of potential approximation.

Most of our analyses are applicable to general fitness and mutation landscapes. However, we characterize in detail the features of artificial selection to direct evolution on quadratic fitness and mutation landscapes, where phenotypes evolve by natural selection toward an evolutionary optimum [44]. In this case, the impacts of selection and mutation follow linear functions in the high-dimensional phenotypic space, $\nabla F = -2C_0 \cdot \mathbf{x}$, $\nabla M = -2L \cdot \mathbf{x}$, where \mathbf{x} denotes the shifted phenotype vector centered around its stationary state and C_0 and L are selection and mutation matrices, respectively— L is a generalization of per-locus mutation rate in high dimensions. The evolutionary model in this case assumes a linear genotype-phenotype map [i.e., $r_{\max} = 1$ in Eq. (A2)] and a nonlinear quadratic phenotype-fitness map, which resembles biophysical models of global epistasis [61,66–71]. We can formulate the evolution of mean phenotypes by

$$d\mathbf{x} = -2KC\mathbf{x}dt + \Sigma d\mathbf{W}, \tag{A6}$$

where $C \equiv N(C_0 + K^{-1}L)$ is the effective adaptive pressure, scaled by the population size, which quantifies the

potential of a phenotype to accumulate mutations under selection. The adaptive potential could in principle be measured directly using lineage tracking evolution experiments, in which impacts of a large number of adaptive mutations can be simultaneously probed [72].

In this work, we use F as a shorthand for the adaptive landscape under natural selection, whose gradient characterizes the adaptive pressure, $\nabla F = -2C\mathbf{x}$, in Eq. (A6). We have also rescaled time with the effective population size (i.e., $t \rightarrow Nt$), which is the coalescence time in neutrality [38].

Similar to the mean, the covariance matrix K is a time-dependent variable, impacted by evolution. However, fluctuations of covariance are much faster compared to the mean phenotype, and therefore, covariance can be approximated by its stationary ensemble-averaged estimate [31,34]. Moreover, even in the presence of moderately strong selection pressure, the phenotypic covariance depends only weakly on the strength of selection and is primarily determined by the supply of mutations in a population [34,35]. Therefore, we also assume that the phenotypic covariance matrix remains approximately constant over time, throughout evolution. With these approximations, evolution of the mean phenotype can be described as a stochastic process with a constant adaptive pressure that approaches its stationary state over a characteristic equilibration time $\sim (2KC)^{-1}$.

The stochastic evolution of the mean phenotype in Eq. (A6) defines an ensemble of evolutionary trajectories. We can characterize the statistics of these evolutionary paths by the dynamics of the underlying conditional probability density $P(\mathbf{x}', t' | \mathbf{x}, t)$ for a population to have a mean phenotype \mathbf{x}' at time t' , given its state \mathbf{x} at an earlier time $t < t'$. The dynamics of this probability density follows a high-dimensional Fokker-Planck equation [34],

$$\frac{\partial}{\partial t} P(\mathbf{x}', t' | \mathbf{x}, t) = \left[\frac{1}{2N} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} - \nabla (K \cdot \nabla F) \right] P(\mathbf{x}', t' | \mathbf{x}, t), \tag{A7}$$

where we introduced the compact notation $\text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} \equiv \sum_{ij} K_{ij} (\partial/\partial x_i) (\partial/\partial x_j)$. As a result, the conditional distribution of phenotypes follows an Ornstein-Uhlenbeck process, described by a time-dependent multivariate Gaussian distribution.

APPENDIX B: HAMILTON-JACOBI-BELLMAN EQUATION FOR OPTIMAL CONTROL

We define a general stochastic evolutionary process for a population of size N with an evolutionary drive due to natural selection and mutations $A(\mathbf{x}, t)$ and an external artificial selection $\mathbf{u}(\mathbf{x}, t)$:

$$d\mathbf{x} = [A(\mathbf{x}) + \mathbf{u}(\mathbf{x}, t)]dt + \Sigma(\mathbf{x})d\mathbf{W}. \tag{B1}$$

Here, time t is measured in units of the coalescence time N (i.e., the effective population size). $d\mathbf{W}$ is a differential random walk due to an underlying Wiener process with an amplitude Σ , which is the square root of the phenotypic covariance matrix K : $\Sigma^\top \Sigma \equiv K$. The stochastic evolution in Eq. (B1) defines an ensemble of phenotypic trajectories, the statistics of which can be characterized by a conditional probability density $P(\mathbf{x}, t | \mathbf{x}', t')$ for a population to have a phenotype \mathbf{x} at time t , given its state \mathbf{x}' at a previous time $t' < t$. For a given artificial selection protocol $\mathbf{u}(\mathbf{x}, t)$, the conditional probability density evolves according to a Fokker-Planck equation [73],

$$\frac{\partial}{\partial t} P(\mathbf{x}, t | \mathbf{x}', t') = \left[\frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} - \nabla_{\mathbf{x}} \cdot (A(\mathbf{x}) + \mathbf{u}(\mathbf{x}, t)) \right] \times P(\mathbf{x}, t | \mathbf{x}', t'), \quad (\text{B2})$$

where we have used the shorthand notation, $\nabla_{\mathbf{x}} \cdot f = \sum_i (\partial / \partial x_i) f$, and $\text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} f = \sum_{i,j} K_{ij} (\partial^2 / \partial x_i \partial x_j) f$, as operators that act on the function f in front of them.

The purpose of artificial selection is to minimize a cost function,

$$\Omega(\mathbf{x}, \mathbf{u}, t) = V(\mathbf{x}, t) + \frac{1}{2} \mathbf{u}^\top B \mathbf{u}, \quad (\text{B3})$$

where $V(\mathbf{x}, t) \equiv V(|\mathbf{x}_t - \mathbf{x}^*|)$ is the cost for deviating from the desired target \mathbf{x}^* during evolution and B is the cost for intervening with natural evolution and applying artificial selection $\mathbf{u} \equiv \mathbf{u}(\mathbf{x}, t)$.

We define the cost-to-go function $J(\mathbf{x}, t)$ as the expected value for the cumulative cost from time t to the end of the process t_f , subject to the evolutionary dynamics and under an optimal control $\mathbf{u}_{t \rightarrow t_f}^*$:

$$J(\mathbf{x}, t) = \min_{\mathbf{u}_{t \rightarrow t_f}} \left\langle Q(\mathbf{x}, t_f) + \int_t^{t_f} \Omega(\mathbf{x}_s, \mathbf{u}_s) ds \right\rangle. \quad (\text{B4})$$

Here, $Q(\mathbf{x}, t_f) \equiv Q(|\mathbf{x}_{t_f} - \mathbf{x}^*|)$ is the cost of deviation from the target at the end point t_f , which in general can be distinct from the path cost $V(\mathbf{x}_t)$. We can formulate a recursive relation for the cost-to-go function $J(\mathbf{x}, t)$,

$$\begin{aligned} J(\mathbf{x}, t) &= \min_{\mathbf{u}_{t \rightarrow t_f}} \left\langle Q(\mathbf{x}_{t_f}) + \int_t^{t_f} \Omega(\mathbf{x}_s, \mathbf{u}_s) ds \right\rangle \\ &= \lim_{\delta t \rightarrow 0} \min_{\mathbf{u}_{t \rightarrow t_f}} \left\langle Q(\mathbf{x}_{t_f}) + \int_t^{t+\delta t} \Omega(\mathbf{x}_s, \mathbf{u}_s) ds + \int_{t+\delta t}^{t_f} \Omega(\mathbf{x}_s, \mathbf{u}_s) ds \right\rangle \\ &= \lim_{\delta t \rightarrow 0} \min_{\mathbf{u}_{t \rightarrow t_f}} \left\langle J(\mathbf{x}_{t+\delta t}, t + \delta t) + \int_t^{t+\delta t} \Omega(\mathbf{x}_s, \mathbf{u}_s) ds \right\rangle \\ &= J(\mathbf{x}_t, t) + \min_{\mathbf{u}_{t \rightarrow t_f}} \left\langle \Omega(\mathbf{x}_s, \mathbf{u}_s) \delta t + \left[\frac{\partial}{\partial t} J(\mathbf{x}_t, t) + (A(\mathbf{x}_t) + \mathbf{u})^\top (\nabla J) + \frac{1}{2} \sum_{ij} K_{ij} \frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} J \right] \delta t \right\rangle, \quad (\text{B5}) \end{aligned}$$

where we used Ito calculus to expand the cost-to-go function, $J(\mathbf{x}_{t+\delta t}, t + \delta t)$; see, e.g., Ref. [37]. By reordering the terms in Eq. (B5), we arrive at the Hamilton-Jacobi-Bellman equation:

$$\begin{aligned} -\frac{\partial}{\partial t} J(\mathbf{x}, t) &= \min_{\mathbf{u}} \left[\Omega(\mathbf{x}_t, \mathbf{u}_t) + (A(\mathbf{x}_t) + \mathbf{u})^\top \cdot \nabla J \right. \\ &\quad \left. + \frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} J \right] \\ &= \min_{\mathbf{u}} \left[\frac{1}{2} \mathbf{u}^\top B \mathbf{u} + \mathbf{u}^\top \cdot \nabla J \right] \\ &\quad + V(\mathbf{x}) + A(\mathbf{x}_t)^\top \cdot \nabla J + \frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} J. \quad (\text{B6}) \end{aligned}$$

The functional form for the optimal artificial selection \mathbf{u}^* follows by minimizing the right-hand side of Eq. (B6) with respect to \mathbf{u} :

$$\mathbf{u}^* = -B^{-1} \nabla J. \quad (\text{B7})$$

Therefore, the time- and phenotype-dependent solution for the cost-to-go function $J(\mathbf{x}, t)$ determines the optimal protocol for artificial selection $\mathbf{u}^*(\mathbf{x}, t)$. By substituting the form of the optimal control \mathbf{u}^* in Eq. (B6), we arrive at a nonlinear partial differential equation for the cost-to-go function,

$$\begin{aligned} -\frac{\partial}{\partial t} J(\mathbf{x}, t) &= -\frac{1}{2} (\nabla J)^\top B^{-1} \nabla J + V(\mathbf{x}) \\ &\quad + A(\mathbf{x}_t)^\top \cdot \nabla J + \frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} J, \quad (\text{B8}) \end{aligned}$$

which should be solved with a boundary condition $J(\mathbf{x}, t_f) = Q(\mathbf{x}, t_f)$ at the end point. We introduce a new variable $\Psi = \exp[-J/\lambda]$ as the exponential of the cost-to-go function. The dynamics of Ψ follows

$$\frac{\lambda}{\Psi} \frac{\partial}{\partial t} \Psi = -\frac{\lambda^2}{2\Psi^2} (\nabla\Psi)^\top B^{-1} \nabla\Psi + V(\mathbf{x}) - \frac{\lambda}{\Psi} A(\mathbf{x}_t)^\top \cdot (\nabla\Psi) - \frac{\lambda}{2} K \left[\frac{-1}{\Psi^2} (\nabla\Psi)^\top \cdot \nabla\Psi + \frac{1}{\Psi} \nabla_{\mathbf{x}\mathbf{x}} \Psi \right]. \quad (\text{B9})$$

The dynamics of Ψ linearizes if and only if there exists a scalar λ that relates the control cost to the covariance matrix such that $B = \lambda K^{-1}$. This criterion is known as the path integral control condition [41,42] by which we can map a generally nonlinear control problem onto a linear stochastic process. The path integral control condition implies that the cost of artificial selection on each phenotype should be inversely proportional to the phenotype's fluctuations. In other words, artificially tweaking with highly conserved phenotypes should be more costly than with variable

phenotypes. In this case, the HJB equation for the transformed cost-to-go function Ψ follows

$$\frac{\partial}{\partial t} \Psi = -A(\mathbf{x})^\top \cdot \nabla\Psi - \frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} \Psi + \frac{1}{\lambda} V(\mathbf{x}) \Psi \equiv -L^\dagger \Psi, \quad (\text{B10})$$

where L^\dagger is a linear operator acting on the function Ψ . Equation (B10) has the form of a backward Fokker-Planck equation with the boundary condition at the end point, $\Psi(\mathbf{x}, t_f) = \exp[-J(\mathbf{x}, t_f)/\lambda] = \exp[Q(\mathbf{x}, t_f)/\lambda]$. We can define a conjugate function P_u that evolves forward in time according to the Hermitian conjugate of the operator L^\dagger . This conjugate operator L can be characterized by evaluating the inner product of the two functions,

$$\begin{aligned} \langle LP_u | \Psi \rangle &= \langle P_u | L^\dagger \Psi \rangle = \int d\mathbf{x} P_u(\mathbf{x}, t) L^\dagger \Psi(\mathbf{x}, t) \\ &= \int d\mathbf{x} P_u(\mathbf{x}, t) \left(A(\mathbf{x})^\top \cdot \nabla\Psi + \frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} \Psi - \frac{1}{\lambda} V(\mathbf{x}) \Psi \right) \Psi(\mathbf{x}, t) \\ &= \int d\mathbf{x} \left(-\frac{1}{\lambda} V(\mathbf{x}) P_u(\mathbf{x}, t) - \nabla A(\mathbf{x}) P_u + \frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} P_u \right)^\top \Psi(\mathbf{x}, t), \end{aligned} \quad (\text{B11})$$

where we performed integration by part and assumed that the function P_u vanishes at the boundaries. This formulation suggests a forward evolution by the operator L^\dagger for the function $P_u(\mathbf{x}', t' | \mathbf{x}, t)$,

$$\begin{aligned} \frac{\partial}{\partial t} P_u(\mathbf{x}', t' | \mathbf{x}, t) &= LP_u(\mathbf{x}', t' | \mathbf{x}, t) \\ &= \underbrace{\left[\frac{1}{2} \text{Tr} K \nabla_{\mathbf{x}\mathbf{x}} - \nabla A(\mathbf{x}) \right]}_{L_0} - \frac{1}{\lambda} V(\mathbf{x}) \Big] P_u(\mathbf{x}', t' | \mathbf{x}, t), \end{aligned} \quad (\text{B12})$$

with a boundary condition at the initial time point $P(\mathbf{x}', t | \mathbf{x}, t) = \delta(\mathbf{x} - \mathbf{x}')$. Importantly, the linear operator L resembles the forward Fokker-Planck operator L_0 for evolution under natural selection [i.e., the dynamics in Eq. (B2) with $\mathbf{u} = 0$] with an extra annihilation term $V(\mathbf{x})/\lambda$. The evolutionary dynamics with the L_0 operator under natural selection conserves the probability density. The annihilation term, on the other hand, eliminates the evolutionary trajectories with a rate proportional to their cost $V(\mathbf{x}, t)/\lambda$ at each time point.

Since Ψ evolves backward in time according to L^\dagger and P_u evolves forward in time according to L , their inner product $\langle P_u | \Psi \rangle = \int d\mathbf{x}' P_u(\mathbf{x}', t' | \mathbf{x}, t) \Psi(\mathbf{x}', t')$ remains time invariant [74]. Therefore, the inner product of the two

functions at the initial and the final time points are equal, which follows

$$\begin{aligned} \langle P_u(t) | \Psi(t) \rangle &= \langle P_u(t_f) | \Psi(t_f) \rangle \rightarrow \int d\mathbf{x}' P_u(\mathbf{x}', t | \mathbf{x}, t) \Psi(\mathbf{x}', t) \\ &= \int d\mathbf{x}' P_u(\mathbf{x}', t_f | \mathbf{x}, t) \Psi(\mathbf{x}', t_f) \\ &\rightarrow \Psi(\mathbf{x}, t) = \int d\mathbf{x}' P_u(\mathbf{x}', t_f | \mathbf{x}, t) \exp[-Q(\mathbf{x}', t_f)/\lambda], \end{aligned} \quad (\text{B13})$$

where we substituted the boundary condition for P_u at the initial time t and for Ψ at the final time t_f . Thus, the cost-to-go function follows

$$\begin{aligned} J(\mathbf{x}, t) &= -\lambda \log \Psi(\mathbf{x}, t) \\ &= -\lambda \log \int d\mathbf{x}' P_u(\mathbf{x}', t_f | \mathbf{x}, t) \exp[-Q(\mathbf{x}', t_f)/\lambda]. \end{aligned} \quad (\text{B14})$$

APPENDIX C: PATH INTEGRAL SOLUTION TO STOCHASTIC ADAPTIVE CONTROL

Given the structure of the linear forward operator L [Eq. (B12)], we can either exactly compute the conditional function $P_u(\mathbf{x}', t_f | \mathbf{x}, t)$ or use approximation methods

common for path integrals (e.g., the semiclassical method) to evaluate cost-to-go function in Eq. (B14). To formulate a path integral for $P_u(\mathbf{x}', t_f | \mathbf{x}, t)$, we discretize the time window $[t: t_f]$ into n small time slices of length ϵ (t_0, t_1, \dots, t_n), with $n\epsilon = t_f - t$. The conditional probability $P_u(\mathbf{x}', t_f | \mathbf{x}, t)$ can be written as an integral over all trajectories that start at the phenotypic state \mathbf{x} at time $t_0 \equiv t$ and end at \mathbf{x}' at time $t_n \equiv t_f$:

$$P_u(\mathbf{x}', t_f | \mathbf{x}, t) \sim \int \prod_{i=1}^n d\mathbf{x}_i P_u(\mathbf{x}_i, t_i | \mathbf{x}_{i-1}, t_{i-1}) \delta(\mathbf{x}_n - \mathbf{x}'). \quad (\text{C1})$$

The short-time propagator $P_u(\mathbf{x}_i, t_i | \mathbf{x}_{i-1}, t_{i-1})$ follows a simple Gaussian form [73],

$$\begin{aligned} P_u(\mathbf{x}_i, t_i | \mathbf{x}_{i-1}, t_{i-1}) &\sim \exp \left\{ -\frac{1}{\lambda} \left[(\mathbf{x}_i - \mathbf{x}_{i-1} - \epsilon A(\mathbf{x}_i))^\top \frac{\lambda K^{-1}}{2\epsilon} (\mathbf{x}_i - \mathbf{x}_{i-1} - \epsilon A(\mathbf{x}_i)) + V(\mathbf{x}_i) \epsilon \right] \right\} \\ &= \exp \left\{ -\frac{\epsilon}{\lambda} \left[\left(\frac{\mathbf{x}_i - \mathbf{x}_{i-1}}{\epsilon} - A(\mathbf{x}_i) \right)^\top \frac{B}{2} \left(\frac{\mathbf{x}_i - \mathbf{x}_{i-1}}{\epsilon} - A(\mathbf{x}_i) \right) + V(\mathbf{x}_i) \right] \right\}, \end{aligned} \quad (\text{C2})$$

where we used $K = \lambda B^{-1}$. We can express the cost-to-go function [Eq. (B14)] as a path integral,

$$\begin{aligned} e^{-J(\mathbf{x}, t)/\lambda} &= \int d\mathbf{x}' P_u(\mathbf{x}', t_f | \mathbf{x}, t) \exp[-Q(\mathbf{x}', t_f)/\lambda] \\ &\sim \int \prod_{i=1}^n d\mathbf{x}_i \exp \left\{ -\frac{\epsilon}{\lambda} \left[\left(\frac{\mathbf{x}_i - \mathbf{x}_{i-1}}{\epsilon} - A(\mathbf{x}_i) \right)^\top \frac{B}{2} \left(\frac{\mathbf{x}_i - \mathbf{x}_{i-1}}{\epsilon} - A(\mathbf{x}_i) \right) + V(\mathbf{x}_i) + \frac{Q(\mathbf{x}_n)}{\epsilon} \right] \right\} \\ &\sim \int \mathcal{D}(\mathbf{x}) \exp \left[-\frac{1}{\lambda} \left(Q(\mathbf{x}(t_f)) + \int_t^{t_f} dt \left[\left(\frac{d\mathbf{x}(t)}{dt} - A(\mathbf{x}(t), t) \right)^\top \frac{B}{2} \left(\frac{d\mathbf{x}(t)}{dt} - A(\mathbf{x}(t), t) \right) + V(\mathbf{x}, t) \right] \right) \right] \\ &\equiv \int \mathcal{D}(\mathbf{x}) \exp \left[-\frac{1}{\lambda} S_{\text{path}}(\mathbf{x}(t \rightarrow t_f)) \right], \end{aligned} \quad (\text{C3})$$

where $S_{\text{path}}(\mathbf{x}(t \rightarrow t_f))$ is a corresponding action and $\mathcal{D}(\mathbf{x}) \sim \prod_{i=1}^n d\mathbf{x}_i$ is the integral measure over all the trajectories that start at $\mathbf{x}_0 = \mathbf{x}(0)$. Numerically, this formulation provides a way to generate evolutionary trajectories under artificial selection as an exponentially weighted ensemble from trajectories under natural selection [41,42]. Moreover, if λ is small, the integral is dominated by the trajectories that are close to the most likely (classical) trajectory $\hat{\mathbf{x}}(t \rightarrow t_f)$, and the path integral can be approximated using the semiclassical method; see Ref. [41].

APPENDIX D: CONTROL OF MOLECULAR PHENOTYPES WITH QUADRATIC COST

In the case that the path cost is zero $V(\mathbf{x}) = 0$, the artificially and naturally selected trajectories become distinct only due to the contributions from the end point cost at $t = t_f$. For the choice of a linear evolutionary force, $A(\mathbf{x}) = -2KC\mathbf{x}$, and a quadratic end point cost, $Q(\mathbf{x}, t_f) = \frac{1}{2}(\mathbf{x}_{t_f} - \mathbf{x}^*)^\top \tilde{G}(\mathbf{x}_{t_f} - \mathbf{x}^*)$, evolution follows an Ornstein-Uhlenbeck process and the solution to Eq. (B12) takes a Gaussian form (see, e.g., Ref. [37]),

$$\begin{aligned} P_u(\mathbf{x}, t) &= \int d\mathbf{x}_{t_f} P_0(\mathbf{x}_{t_f}, t_f | \mathbf{x}, t) P_0(\mathbf{x}, t) \exp[-Q(\mathbf{x}_{t_f})/\lambda] \\ &\sim \int d\mathbf{x}_{t_f} \exp \left[\frac{-1}{2} (\mathbf{x}_{t_f} - \mu(\mathbf{x}, \tau))^\top K_\tau^{-1} \right. \\ &\quad \left. \times (\mathbf{x}_{t_f} - \mu(\mathbf{x}, \tau)) \right] P_0(\mathbf{x}, t) \exp[-Q(\mathbf{x}_{t_f})/\lambda], \end{aligned} \quad (\text{D1})$$

where $P_0(\mathbf{x}, t)$ denotes the marginal phenotype distribution in the uncontrolled process at time t and $P_0(\mathbf{x}_{t_f}, t_f | \mathbf{x}, t)$ indicates the conditional probability density in the uncontrolled process, which is a Gaussian distribution with a time-dependent mean,

$$\mu(\mathbf{x}, \tau) = \exp[-2KC\tau]\mathbf{x}, \quad (\text{D2})$$

and a covariance matrix,

$$K_\tau = \int_t^{t_f} dt' \exp[-2KC(t_f - t')] K \exp[-2CK(t_f - t')], \quad (\text{D3})$$

with $\tau = t_f - t$.

In this case, the cost to go in Eq. (B14) can be evaluated by a Gaussian integral to marginalize over the end state \mathbf{x}_{t_f} ,

$$\begin{aligned} \exp[-J/\lambda] &\sim \int d\mathbf{x}_{t_f} \exp\left[\frac{-1}{2\lambda}(\mathbf{x}_{t_f} - \mu(\mathbf{x}, \tau))^\top \lambda K_\tau^{-1} \times (\mathbf{x}_{t_f} - \mu(\mathbf{x}, \tau))\right] \exp\left[\frac{-1}{2\lambda}(\mathbf{x}_{t_f} - \mathbf{x}^*)^\top \tilde{G}(\mathbf{x}_{t_f} - \mathbf{x}^*)\right] \\ &\sim \exp\left[\frac{1}{2\lambda}(\mu(\mathbf{x}, \tau) - \mathbf{x}^*)^\top (\tilde{G}[\lambda K_\tau^{-1} + \tilde{G}]^{-1} \tilde{G} - \tilde{G})(\mu(\mathbf{x}, \tau) - \mathbf{x}^*)\right], \end{aligned} \quad (\text{D4})$$

resulting in an optimal artificial selection protocol:

$$\begin{aligned} \mathbf{u}^* &= -B^{-1}\nabla J = -\frac{K}{\lambda}\nabla J = -\frac{K}{\lambda}[\nabla\mu(\mathbf{x}, \tau)]^\top [\tilde{G} - \tilde{G}[\lambda K_\tau^{-1} + \tilde{G}]^{-1}\tilde{G}](\mu(\mathbf{x}, \tau) - \mathbf{x}^*) \\ &= -\frac{K}{\lambda}\exp[-2CK\tau]\tilde{G}\left[I - \frac{K_\tau}{\lambda}\left[I + \frac{K_\tau}{\lambda}\tilde{G}\right]^{-1}\tilde{G}\right](e^{-2CK\tau}\mathbf{x} - \mathbf{x}^*). \end{aligned} \quad (\text{D5})$$

When the goal is to drive the population toward a target by an end point t_f , the effective fitness $\hat{F}(\mathbf{x}, t)$ remains close to the natural landscape for an extended period. As time approaches the end point, populations transition from evolving in their natural landscape $F(\mathbf{x})$ to the artificially induced fitness landscape $F_{\text{art}}(\mathbf{x}, t_f)$; see Fig. 3 herein and Figs. S1 and S2 in Supplemental Material [45] for evolution in one and two dimensions. Moreover, toward the end point, the fitness peak and the strength of selection approach their final values, determined by the target and the cost functions in Eq. (B3), in an exponentially rapid manner (Figs. S1 and S2 [45] and Fig. 3).

As the time approaches the end point ($t \rightarrow t_f$ or $\tau \rightarrow 0$), optimal artificial selection acts as a linear attractive force (i.e., a stabilizing selection),

$$\mathbf{u}^*(\tau \rightarrow 0) = \frac{-1}{\lambda}K\tilde{G} \cdot (\mathbf{x} - \mathbf{x}^*) + \mathcal{O}(\tau), \quad (\text{D6})$$

to maintain the population close to the phenotypic target, with an effective strength of artificial stabilizing selection \tilde{G}/λ .

APPENDIX E: MODEL MISSPECIFICATION AND ALTERNATIVE APPROACHES TO ARTIFICIAL SELECTION

Devising an optimal control strategy relies on the knowledge of system dynamics, which in some cases may not be available with high precision. In such settings, *ad hoc* selection protocols, guided by optimal control, can be used to drive evolution toward a desired target. Here, we consider such alternative (*ad hoc*) selection protocols and compare their performance with the optimal artificial selection strategy for the 2D covariate phenotypes, in terms of total cost and performance variability.

(a) Control without the knowledge of phenotypic correlation. Optimal control on multivariate phenotypes should be devised by considering the phenotypic correlation, which could lead to nonmonotonic artificial selection strategies in the course of evolution [Fig. 3(c)]. Here, we can quantify the importance of phenotypic covariance on devising optimal artificial selection protocols. To do so, we assume that the covariance matrix is misspecified and that phenotypes evolve independently (i.e., assuming the covariance matrix K is diagonal). The devised protocol in this case does not take into account the synergistic or antagonistic interactions between the phenotypes and their response to interventions. Therefore, as correlation between phenotypes increases, the cost function [i.e., $Q(\mathbf{x}, t_f) + \int_{t_f}^t ds[V(\mathbf{x}_s) + \frac{1}{2}\mathbf{u}_s^\top B\tilde{\mathbf{u}}_s]$] with the misspecified control protocol $\tilde{\mathbf{u}}$ soars (Fig. S2 [45]). It should be noted that although phenotypic cost with the misspecified control protocol always exceeds the cost under optimal control, its impact could still be favorable over the uncontrolled system (Fig. S2 [45]), in certain parameter regimes.

(b) *Ad hoc* proportional artificial selection protocol on multivariate phenotypes. We consider a naive (and intuitive) approach to artificial selection. Specifically, for 2D covariate phenotypes, the strategy is to intervene with a proportional control (i.e., artificial selection in a quadratic landscape) as long as the phenotypes are outside a specified neighborhood (with range δ) of their targets, and to relax intervention once phenotypes are close enough to their targets,

$$\mathbf{u} = \begin{cases} \mathbf{u}_x = -\kappa_t(x_t - x^*), \mathbf{u}_y = 0 & \text{if } |x_t - x^*| > \delta \\ \mathbf{u}_x = 0, \mathbf{u}_y = -\kappa_t(y_t - y^*) & \text{if } |x_t - x^*| < \delta, |y_t - y^*| > \delta \\ \mathbf{u}_x = 0, \mathbf{u}_y = 0 & \text{otherwise,} \end{cases} \quad (\text{E1})$$

where κ_t is the strength of artificial selection. In this strategy, artificial selection is preferentially applied to phenotype “ x ” and once x is close enough to its target, artificial interventions would select for phenotype “ y .” In a sense, this protocol presents an intuitive approach to a control problem similar to the one considered in Fig. 3, in which phenotype x is driven toward a target far from its natural state, while keeping phenotype y close to its natural state.

We implement two versions of Eq. (E1). In the first version, we select κ_t to be equal to a constant $\kappa > 0$ over all times. We denote this version as proportional control. In the second version, we implement an exponentially increasing artificial selection strength as evolution approaches the end time T , i.e., $\kappa_t = \kappa \exp^{-\tau+1}$. We denote this protocol as proportional control with exponential weights. We note that the choice of exponentially increasing artificial selection strength mimics the increase in control gains of optimal artificial selection strategy as remaining time approaches, in accordance with Eq. (D5). We tune the constants in both control strategies such that the mean values of the ensemble of end point values of phenotype x ($\langle x_T \rangle$) are approximately equal to the mean ensemble value obtained from the optimal artificial selection protocol, and close to the desired target x^* .

We observe that the focal phenotype x_t at the end point is much more widely distributed under both versions of the proportional control protocol [Eq. (E1)] compared to the optimal protocol (Fig. S5A [45]). Similarly, both proportional control strategies obtain similar y_T values near the target $y^* = 0$, with a higher ensemble variability than the one obtained using the optimal artificial selection strategy (Fig. S5A [45]). The variability in end point values tends to be higher in proportional control with exponential weights compared to the rest of the protocols. We also note that the peak magnitude of the artificial selection is much smaller for the optimal control compared to the proportional control protocols; see Fig. S6 for a comparison of control magnitudes between the optimal strategy and the proportional control strategy with exponential weights [45].

When we compare the costs of control, the proportional control strategy with exponential weights has an order of magnitude smaller cost than the proportional control with a constant weight (Fig. S5B [45]). This is expected because the earlier control actions incur control costs but their effects are overwritten by the uncontrolled selection dynamics. In addition, the proportional control with exponential weights is on average preferred over doing nothing in terms of the total cost. Still, the accumulated cost for the proportional control strategy with exponential weights is at least an order of magnitude higher than the costs incurred under the optimal strategy (Fig. S5B [45]).

In summary, while the *ad hoc* control strategies can be tuned to drive evolution toward a desired phenotypic target on average, they would suffer from an increased variability across evolutionary realizations and an elevated cost of control.

APPENDIX F: ARTIFICIAL SELECTION WITH INTERMITTENT MONITORING

Here, we assume that artificial selection is imposed in discrete steps with time interval τ . Similar to the continuous control, the cost function has two components: the cost of control at the end of each intervention and the cumulative cost of deviation from the optimum throughout each interval. The stationary cost-to-go function follows

$$J(\mathbf{x}, t_m; \tau) = \min_{\mathbf{u}} \lim_{M \rightarrow \infty} \frac{1}{(M - m)\tau} \times \left\langle \sum_{i=m}^M \mathbf{u}_i^\top B \mathbf{u}_i + \int_{t_i}^{t_M} V(\mathbf{x}_t) dt \right\rangle_{\text{evol}}, \quad (\text{F1})$$

where we have normalized the path cost by the interval τ to assure that the cost to go remains finite.

To further simplify, we consider only one-dimensional phenotype x with intrapopulation variance k , the cost of deviation $V(x) = g(x - x^*)^2/2$ from target x^* , and the cost of intervention $\beta u^2/2$ with artificial selection u . In the stationary state and in the regime of small interventions ($gk/\lambda < 1$), we assume that the optimal artificial selection protocol u^* should be a variant of the case with full information with a strength of selection α_τ dependent on the time window τ , $u_\tau^* = -k\alpha_\tau(x - x^*)$. Our goal is to characterize the strength of optimal artificial selection α_τ .

The total cost over an interval τ in the stationary state follows

$$\begin{aligned} \Omega_\tau(x) &= \frac{\beta}{2} \langle u^2 \rangle + \frac{1}{\tau} \left\langle \int_{t=t_i}^{t_i+\tau} V(x_t) dt \right\rangle \\ &= \left\langle \frac{\beta}{2} k^2 \alpha_\tau^2 (x_\tau - x^*)^2 + \frac{1}{2\tau} g \int_{t=0}^\tau (x_t - x^*)^2 dt \right\rangle. \end{aligned} \quad (\text{F2})$$

We are interested in the regime of moderate to weak interventions ($gk/\lambda < 1$), for which the linear response theory can characterize the change in the state of the system after each intervention. In this regime, evolution under artificial selection can be approximated as a perturbation from the dynamics under natural selection. The evolutionary dynamics of the phenotype distribution is governed by a time-dependent Fokker-Planck operator, $L(x, t)$,

$$\frac{\partial}{\partial t} P_{u;\tau}(x, t) = [L_0(\mathbf{x}) + L_u(\mathbf{x})Y(t)]P_{u;\tau}(x, t), \quad (\text{F3})$$

where $P_{u;\tau}(\mathbf{x}, t)$ is the full probability density under intermittent artificial selection, which can be split into the stationary solution under natural selection and the time-dependent deviation due to artificial selection: $P_{u;\tau}(x, t) = P_0(\mathbf{x}) + P_u(\mathbf{x}, t; \tau)$. $L_0(\mathbf{x})$ is the Fokker-Planck operator associated with evolution under natural selection [i.e., the dynamics in Eq. (B2) with $\mathbf{u} = 0$], $L_u(\mathbf{x}) = \partial_x k \alpha_\tau (x - x^*)$ is the state-dependent operator associated with artificial

selection, and $Y(t) = \lim_{M \rightarrow \infty} \sum_{i=1}^M \delta(t - t_i)$ characterizes a time-dependent component due to artificial selection interventions at the end of each time interval.

In the regime of linear response, where the impact of artificial selection is small, the deviation $\langle \Delta z \rangle$ of an expected value of an arbitrary function $\langle z(\mathbf{x}) \rangle$ from its stationary state (i.e., under natural selection) follows (see, e.g., Ref. [73])

$$\begin{aligned} \langle \Delta z(t) \rangle &= \int z(\mathbf{x}) P_u(x, t) d\mathbf{x} \\ &\equiv \int_{-\infty}^{\infty} R_{z, L_u}(t - t') Y(t') dt', \end{aligned} \quad (\text{F4})$$

where $R_{z, L_u}(t)$ is the response function to artificial selection L_u :

$$R_{z, L_u}(t) = \begin{cases} \int z(\mathbf{x}) e^{L_0(\mathbf{x}) \cdot t} L_u(\mathbf{x}) P_0(\mathbf{x}) d\mathbf{x} & \text{for } t \geq 0 \\ 0 & \text{for } t < 0. \end{cases} \quad (\text{F5})$$

At end of each time interval, artificial selection imposes a shock-type perturbation to the evolutionary process. The immediate response of the population to this selection pressure can be characterized by the instantaneous response function [i.e., with $Y(t') = \delta(t - t')$], resulting in the change in a given phenotypic statistics z (see, e.g., Ref. [73]),

$$\begin{aligned} \langle \Delta z(t) \rangle &= \int z(x) L_u(\mathbf{x}) P_0(x) dx \\ &= \frac{1}{Z} \int dx z(x) \frac{\partial}{\partial x} \left(k\alpha_\tau (x - x^*) \exp \left[-\frac{x^2}{2\text{var}_{\text{st}_0}} \right] \right) \\ &= k\alpha_\tau \left(\langle z(x) \rangle_{\text{st}_0} - \frac{1}{\text{var}_{\text{st}_0}} [\langle z(x)x^2 \rangle_{\text{st}_0} - x^* \langle z(x)x \rangle_{\text{st}_0}] \right), \end{aligned} \quad (\text{F6})$$

where $P_0(x)$ is the Gaussian stationary distribution for the mean phenotype under natural selection, $\text{var}_{\text{st}_0} = 1/4c$ is the

stationary ensemble variance for the mean phenotype under natural selection, Z is the normalization factor for the underlying stationary distribution, and $\langle \cdot \rangle_{\text{st}_0}$ indicates expectation values under the stationary distribution.

At the beginning of each interval $t = 0$, the deviation of the mean phenotype from its expectation under natural selection $\langle x \rangle_{\text{st}_0} = 0$ follows

$$\langle \Delta x \rangle = \langle x(t=0) \rangle = k\alpha_\tau x^*. \quad (\text{F7})$$

Similarly, the deviation in the second moment of the phenotypic distribution from the expectation under natural selection $\langle x^2 \rangle_{\text{st}_0} = \text{var}_{\text{st}_0}$ follows

$$\begin{aligned} \langle \Delta x^2 \rangle &= \langle x^2(t=0) \rangle - \text{var}_{\text{st}_0} \\ &= \alpha_\tau \left(\text{var}_{\text{st}_0} - \frac{1}{\text{var}_{\text{st}_0}} \langle x^4 \rangle_{\text{st}_0} \right) \\ &= -2k\alpha_\tau \text{var}_{\text{st}_0} = -k\alpha_\tau / 2c. \end{aligned} \quad (\text{F8})$$

Thus, the phenotypic variance at the beginning of each interval follows

$$\begin{aligned} \text{var}_u(t=0) &= \langle x^2(t=0) \rangle - \langle x(t=0) \rangle^2 \\ &= \text{var}_{\text{st}_0} [1 - 2k\alpha_\tau] - [k\alpha_\tau x^*]^2. \end{aligned} \quad (\text{F9})$$

Following an intervention at time $t = 0$, populations evolve according to natural selection until the next intervention. Therefore, the phenotype statistics during each time interval ($0 < t < \tau$) follow

$$\langle x(t) \rangle = \alpha_\tau x^* e^{-2kct}, \quad (\text{F10})$$

$$\begin{aligned} \text{var}(t) &= [\text{var}_{\text{st}_0} (1 - 4kc) - (k\alpha_\tau x^*)^2] e^{-4kct} \\ &\quad + \text{var}_{\text{st}_0} (1 - e^{-4kct}) \\ &= \text{var}_{\text{st}_0} (1 - 2k\alpha_\tau e^{-4kct}) - (k\alpha_\tau x^*)^2 e^{-4kct}. \end{aligned} \quad (\text{F11})$$

We can now evaluate the cumulative cost function [Eq. (F2)],

$$\begin{aligned} \Omega_\tau(x) &= \frac{\beta}{2} \langle u^2 \rangle + \frac{1}{2\tau} \gamma \left\langle \int_{t=0}^{\tau} (x_t - x^*)^2 dt \right\rangle \\ &= \left\langle \frac{\beta}{2} k^2 \alpha_\tau^2 (x_\tau - x^*)^2 + \frac{1}{2\tau} \gamma \int_{t=0}^{\tau} (x_t - x^*)^2 dt \right\rangle \\ &= \frac{\beta}{2} k^2 \alpha_\tau^2 [(\langle x_\tau \rangle - x^*)^2 + \text{var}(\tau)] + \frac{1}{2\tau} \gamma \int_{t=0}^{\tau} [(\langle x_t \rangle - x^*)^2 + \sigma^2(t)] dt \\ &= \frac{\beta}{2} k^2 \alpha_\tau^2 [\text{var}_{\text{st}_0} (1 - 2k\alpha_\tau e^{-4kc\tau}) + (x^*)^2 (1 - 2k\alpha_\tau e^{-2kc\tau})] \\ &\quad + \frac{1}{2\tau} \gamma \left[\frac{-\alpha_\tau}{2c} ((1 - e^{-4kc\tau}) \text{var}_{\text{st}_0} + 2(1 - e^{-2kc\tau})(x^*)^2) + (\text{var}_{\text{st}_0} + (x^*)^2) \tau \right], \end{aligned} \quad (\text{F12})$$

where we have used the time-dependent expectation and variance in Eqs. (F10) and (F11).

The optimal strength of artificial selection α_τ^* for intermittent interventions can be characterized by minimizing the cost function [Eq. (F12)] with respect to α_τ ,

$$\alpha_\tau^* = \frac{\gamma}{\lambda} \left(\frac{(1 - e^{-\tau})(1 + 8c(x^*)^2 + e^{-\tau})}{2\tau(1 + 4c(x^*)^2)} \right) + \mathcal{O}[(k\gamma/\lambda)^2], \quad (\text{F13})$$

which in the limit of small separation time ($\tau \rightarrow 0$) approaches the expectation under continuous monitoring in the stationary state [Eq. (D6)], $\alpha^*(\tau \rightarrow 0) = \gamma/\lambda$.

APPENDIX G: WORK PERFORMED BY ARTIFICIAL SELECTION

Artificial selection changes the distribution of phenotypic trajectories $\mathbf{x}_{t_0}^{t_f} \equiv (\mathbf{x}_{t_0}, \dots, \mathbf{x}_{t_f})$ from $P_0(\mathbf{x}_{t_0}^{t_f})$ in the stationary state under natural selection to a configuration closer to the desired target $P_u(\mathbf{x}_{t_0}^{t_f})$. In analogy to thermodynamics, we can associate a free energy to these distributions, as $F_0 = \log P_0(\mathbf{x}_{t_0}^{t_f})$ and $F_u = \log P_u(\mathbf{x}_{t_0}^{t_f})$ [50]. The expected difference between the free energy of the

two phenotypic configurations can be interpreted as the amount of work done by artificial selection, which corresponds to the Kullback-Leibler distance between the two distributions,

$$\begin{aligned} W_u &\equiv \langle F_u \rangle - \langle F_0 \rangle = \int \mathcal{D}\mathbf{x} P_u(\mathbf{x}_{t_0}^{t_f}) \log \left[\frac{P_u(\mathbf{x}_{t_0}^{t_f})}{P_0(\mathbf{x}_{t_0}^{t_f})} \right] \\ &\equiv D_{\text{KL}}(P_u(\mathbf{x}_{t_0}^{t_f}) || P_0(\mathbf{x}_{t_0}^{t_f})), \end{aligned} \quad (\text{G1})$$

where $\mathcal{D}\mathbf{x}$ is the integral measure over all trajectories. The estimate of work in Eq. (G1), however, should not be interpreted as a physical work, but rather as an information theoretical measure of discrimination between the two phenotype distributions due to artificial selection.

The evolution of the distribution for phenotype trajectories $P_u(\mathbf{x}_{t_0}^{t_f})$ under a given artificial selection protocol $\mathbf{u}_{t_0}^{t_f}$ is Markovian [Eqs. (B1) and (B2)]. To characterize this path probability density, we follow the path integral formulation in Eq. (C1) and discretize the time window $[t_0 : t_f]$ into n small time slices of length ϵ , (t_0, t_1, \dots, t_n) , with $n\epsilon = t_f - t_0$. The probability of a given trajectory $P_u(\mathbf{x}_{t_0}^{t_f})$ can be written as a product of short-term propagators (i.e., conditional probabilities) (see Ref. [73]),

$$\begin{aligned} P_u(\mathbf{x}_{t_0}^{t_f}) &= \lim_{\epsilon \rightarrow 0} \prod_{s=t_0}^{t_f} (P_u(\mathbf{x}_{s+\epsilon} | \mathbf{x}_s)) \\ &= \lim_{n \rightarrow \infty} \prod_{i=1}^n \frac{1}{Z_i} \exp \left[-(\mathbf{x}_{i+1} - \mathbf{x}_i - \epsilon(A(\mathbf{x}_i) + \mathbf{u}(\mathbf{x}_i)))^\top \frac{K^{-1}}{2\epsilon} (\mathbf{x}_{i+1} - \mathbf{x}_i - \epsilon(A(\mathbf{x}_i) + \mathbf{u}(\mathbf{x}_i))) \right] \\ &\sim P_0(\mathbf{x}_{t_0}^{t_f}) \lim_{n \rightarrow \infty} \prod_{i=1}^n e^{\mathbf{u}^\top(\mathbf{x}_i) K^{-1} (\mathbf{x}_{i+1} - \mathbf{x}_i - \epsilon A(\mathbf{x}_i))} \times e^{-(\epsilon/2) \mathbf{u}(\mathbf{x}_i)^\top K^{-1} \mathbf{u}(\mathbf{x}_i)} \\ &= P_0(\mathbf{x}_{t_0}^{t_f}) \exp \left[- \int_{t_0}^{t_f} dt \frac{1}{2} \mathbf{u}^\top(\mathbf{x}, t) K^{-1} \mathbf{u}(\mathbf{x}, t) + \int_{t_0}^{t_f} \mathbf{u}^\top(\mathbf{x}, t) K^{-1} (d\mathbf{x}_t - A(\mathbf{x}_t) dt) \right]. \end{aligned} \quad (\text{G2})$$

The Kullback-Leibler distance between the two distributions follows

$$\begin{aligned} D_{\text{KL}}(P_u(\mathbf{x}_{t_0}^{t_f}) || P_0(\mathbf{x}_{t_0}^{t_f})) &= \int \mathcal{D}\mathbf{x} P_u(\mathbf{x}_{t_0}^{t_f}) \left[- \int_{t_0}^{t_f} dt \left(\frac{1}{2} \mathbf{u}^\top(\mathbf{x}, t) K^{-1} \mathbf{u}(\mathbf{x}, t) \right) + \int_{t_0}^{t_f} \mathbf{u}^\top(\mathbf{x}, t) K^{-1} (d\mathbf{x}_t - A(\mathbf{x}_t) dt) \right] \\ &= \int_{t_0}^{t_f} dt \int \mathcal{D}\mathbf{x} P_u(\mathbf{x}_{t_0}^{t_f}) \left(\frac{1}{2} \mathbf{u}^\top(\mathbf{x}, t) K^{-1} \mathbf{u}(\mathbf{x}, t) \right) \equiv \left\langle \frac{1}{2} (\mathbf{x}_{t_0}^{t_f})^\top K^{-1} \mathbf{u}_{t_0}^{t_f} \right\rangle, \end{aligned} \quad (\text{G3})$$

where we have used $d\mathbf{x}_t - A(\mathbf{x}_t) dt = \mathbf{u}(\mathbf{x}_t, t) dt + d\mathbf{W}_t$, with $d\mathbf{W}_t$ as the stochastic differential measure for a multivariate Wiener process (see Ref. [37]). Importantly, with the criteria of path integral control (i.e., $K^{-1} = B/\lambda$), the Kullback-Leibler distance between the artificially and naturally selected phenotype distributions is equivalent to the cumulative cost of intervention, divided by the overall cost of artificial selection λ ,

$$D_{\text{KL}}(P_u(\mathbf{x}_{t_0}^{t_f}) || P_0(\mathbf{x}_{t_0}^{t_f})) = \frac{1}{2\lambda} \langle (\mathbf{u}_{t_0}^{t_f})^\top B \mathbf{u}_{t_0}^{t_f} \rangle, \quad (\text{G4})$$

which can intuitively be interpreted as the amount of work done by artificial selection.

- [1] C. R. Darwin and A. R. Wallace, *On the Tendency of Species to form Varieties and on the Perpetuation of Varieties and Species by Natural Means of Selection*, *J. Proc. Linn. Soc. London Zool.* **3**, 45 (1858).
- [2] R. Chakrabarti, H. Rabitz, S. L. Springs, and G. L. McLendon, *Mutagenic Evidence for the Optimal Control of Evolutionary Dynamics*, *Phys. Rev. Lett.* **100**, 258103 (2008).
- [3] J. B. de Lamarch, *Philosophie Zoologique, ou Exposition des Considérations Relatives à l'Histoire Naturelle des Animaux* (J. B. Bailliere, Paris, 1830).
- [4] N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, and S. Leibler, *Bacterial Persistence as a Phenotypic Switch*, *Science* **305**, 1622 (2004).
- [5] K. Paarporn, C. Eksin, and J. S. Weitz, *Information Sharing for a Coordination Game in Fluctuating Environments*, *J. Theor. Biol.* **454**, 376 (2018).
- [6] O. Rivoire and S. Leibler, *A Model for the Generation and Transmission of Variations in Evolution*, *Proc. Natl. Acad. Sci. U.S.A.* **111**, E1940 (2014).
- [7] G. Mendel, *Versuche über Pflanzen-Hybriden*, *Verhandlungen des naturforschenden Vereines in Brünn* **5**, 3 (1862).
- [8] A. Fischer, I. Vazquez-Garcia, and V. Mustonen, *The Value of Monitoring to Control Evolving Populations*, *Proc. Natl. Acad. Sci. U.S.A.* **112**, 1007 (2015).
- [9] M. Lässig and V. Mustonen, *Eco-Evolutionary Control of Pathogens*, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 19694 (2020).
- [10] S. Wang, J. Mata-Fink, B. Kriegsman, M. Hanson, D. J. Irvine, H. N. Eisen, D. R. Burton, K. D. Wittrup, M. Kardar, and A. K. Chakraborty, *Manipulating the Selection Forces during Affinity Maturation to Generate Cross-Reactive HIV Antibodies*, *Cell* **160**, 785 (2015).
- [11] J. S. Shaffer, P. L. Moore, M. Kardar, and A. K. Chakraborty, *Optimal Immunization Cocktails Can Promote Induction of Broadly Neutralizing Abs against Highly Mutable Pathogens*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, E7039 (2016).
- [12] K. G. Sprenger, J. E. Louveau, P. M. Murugan, and A. K. Chakraborty, *Optimizing Immunization Protocols to Elicit Broadly Neutralizing Antibodies*, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 20077 (2020).
- [13] V. Sachdeva, K. Husain, J. Sheng, S. Wang, and A. Murugan, *Tuning Environmental Timescales to Evolve and Maintain Generalists*, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 12693 (2020).
- [14] M. Eigen and W. Gardiner, *Evolutionary Molecular Engineering Based on RNA Replication*, *Pure Appl. Chem.* **56**, 967 (1984).
- [15] K. Chen and F. H. Arnold, *Tuning the Activity of an Enzyme for Unusual Environments: Sequential Random Mutagenesis of Subtilisin E for Catalysis in Dimethylformamide*, *Proc. Natl. Acad. Sci. U.S.A.* **90**, 5618 (1993).
- [16] K. M. Esvelt, J. C. Carlson, and D. R. Liu, *A System for the Continuous Directed Evolution of Biomolecules*, *Nature (London)* **472**, 499 (2011).
- [17] E. Toprak, A. Veres, J.-B. Michel, R. Chait, D. L. Hartl, and R. Kishony, *Evolutionary Paths to Antibiotic Resistance under Dynamically Sustained Drug Selection*, *Nat. Genet.* **44**, 101 (2012).
- [18] E. Toprak, A. Veres, S. Yildiz, J. M. Pedraza, R. Chait, J. Paulsson, and R. Kishony, *Building a Morbidostat: An Automated Continuous-Culture Device for Studying Bacterial Drug Resistance under Dynamically Sustained Drug Inhibition*, *Nat. Protoc.* **8**, 555 (2013).
- [19] K. H. Ang, G. Chong, and Y. Li, *PID Control System Analysis, Design, and Technology*, *IEEE Trans. Control Syst. Technol.* **13**, 559 (2005).
- [20] B. G. Wong, C. P. Mancuso, S. Kiriakov, C. J. Bashor, and A. S. Khalil, *Precise, Automated Control of Conditions for High-Throughput Growth of Yeast and Bacteria with eVOLVER*, *Nat. Biotechnol.* **36**, 614 (2018).
- [21] Z. J. Heins, C. P. Mancuso, S. Kiriakov, B. G. Wong, C. J. Bashor, and A. S. Khalil, *Designing Automated, High-Throughput, Continuous Cell Growth Experiments Using eVOLVER*, *J. Vis. Exp.* **147**, e59652 (2019).
- [22] Z. Zhong, B. G. Wong, A. Ravikumar, G. A. Arzumanyan, A. S. Khalil, and C. C. Liu, *Automated Continuous Evolution of Proteins In Vivo*, *ACS Synth. Biol.* **9**, 1270 (2020).
- [23] A. Nourmohammad, T. Held, and M. Lässig, *Universality and Predictability in Molecular Quantitative Genetics*, *Curr. Opin. Genet. Dev.* **23**, 684 (2013).
- [24] M. Lässig, V. Mustonen, and A. M. Walczak, *Predicting Evolution*, *Nat. Ecol. Evol.* **1**, 0077 (2017).
- [25] S. Kryazhinskiy, D. P. Rice, E. R. Jerison, and M. M. Desai, *Global Epistasis Makes Adaptation Predictable Despite Sequence-Level Stochasticity*, *Science* **344**, 1519 (2014).
- [26] J. Barroso-Batista, J. Demengeot, and I. Gordo, *Adaptive Immunity Increases the Pace and Predictability of Evolutionary Change in Commensal Gut Bacteria*, *Nat. Commun.* **6**, 8945 (2015).
- [27] O. Tenaillon, A. Rodríguez-Verdugo, R. L. Gaut, P. McDonald, A. F. Bennett, A. D. Long, and B. S. Gaut, *The Molecular Diversity of Adaptive Convergence*, *Science* **335**, 457 (2012).
- [28] B. K. Shoichet, W. A. Baase, R. Kuroki, and B. W. Matthews, *A Relationship between Protein Stability and Protein Function*, *Proc. Natl. Acad. Sci. U.S.A.* **92**, 452 (1995).
- [29] K. B. Zeldovich, P. Chen, and E. I. Shakhnovich, *Protein Stability Imposes Limits on Organism Complexity and Speed of Molecular Evolution*, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 16152 (2007).
- [30] D. P. Bertsekas, *Dynamic Programming and Optimal Control* (Athena Scientific, Belmont, MA, 1995).
- [31] A. Nourmohammad, J. Otwinowski, and J. B. Plotkin, *Host-Pathogen Coevolution and the Emergence of Broadly Neutralizing Antibodies in Chronic Infections*, *PLoS Genet.* **12**, e1006171 (2016).
- [32] G. J. Wedemayer, P. A. Patten, L. H. Wang, P. G. Schultz, and R. C. Stevens, *Structural Insights into the Evolution of an Antibody Combining Site*, *Science* **276**, 1665 (1997).
- [33] S. A. Frank, *Immunology and Evolution of Infectious Disease* (Princeton University Press, Princeton, NJ, 2002).
- [34] A. Nourmohammad, S. Schiffels, and M. Lässig, *Evolution of Molecular Phenotypes under Stabilizing Selection*, *J. Stat. Mech.* (2013) P01012.
- [35] T. Held, A. Nourmohammad, and M. Lässig, *Adaptive Evolution of Molecular Phenotypes*, *J. Stat. Mech.* (2014) P09029.

- [36] R. Lande, *Natural Selection and Random Genetic Drift in Phenotypic Evolution*, *Evolution* **30**, 314 (1976).
- [37] C. Gardiner, *Handbook of Stochastic Methods: For Physics, Chemistry and the Natural Sciences*, 3rd ed. (Springer, New York, 2004).
- [38] J. F. C. Kingman, *On the Genealogy of Large Populations*, *J. Appl. Probab.* **19**, 27 (1982).
- [39] R. E. Bellman, *Dynamic Programming* (Princeton University Press, Princeton, NJ, 1957).
- [40] R. E. Bellman, *Dynamic Programming and a New Formalism in the Calculus of Variations*, *Proc. Natl. Acad. Sci. U.S.A.* **40**, 231 (1954).
- [41] H. J. Kappen, *Linear Theory for Control of Nonlinear Stochastic Systems*, *Phys. Rev. Lett.* **95**, 200201 (2005).
- [42] H. J. Kappen, *Path Integrals and Symmetry Breaking for Optimal Control Theory*, *J. Stat. Mech.* (2005) P11011.
- [43] E. Todorov, *Linearly-Solvable Markov Decision Problems*, in *Advances in Neural Information Processing Systems 19*, edited by B. Scholkopf, J. Platt, and T. Hoffman (MIT Press, Cambridge, MA, 2007), pp. 1369–1376.
- [44] R. A. Fisher, *The Genetical Theory of Natural Selection*, 1st ed. (Oxford University Press, New York, 1930), Vol. 19.
- [45] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.11.011044> for Supplemental Figs. S1–S6.
- [46] J. L. Lush, *Animal Breeding Plans* (Iowa State Press, Ames, IA, 1937).
- [47] R. Lande and S. J. Arnold, *The Measurement of Selection on Correlated Characters*, *Evolution* **37**, 1210 (1983).
- [48] D. Tieman, G. Zhu, M. F. R. Resende, T. Lin, C. Nguyen, D. Bies, J. L. Rambla, K. S. O. Beltran, M. Taylor, B. Zhang, H. Ikeda, Z. Liu, J. Fisher, I. Zemach, A. Monforte, D. Zamir, A. Granell, M. Kirst, S. Huang, and H. Klee, *A Chemical Genetic Roadmap to Improved Tomato Flavor*, *Science* **355**, 391 (2017).
- [49] O. Rivoire and S. Leibler, *The Value of Information for Populations in Varying Environments*, *J. Stat. Phys.* **142**, 1124 (2011).
- [50] L. D. Landau and E. M. Lifshitz, *Statistical Physics; Part 1*, 1st ed. (Addison-Wesley, Reading, MA, 1958).
- [51] H. J. Kappen and H. C. Ruiz, *Adaptive Importance Sampling for Control and Inference*, *J. Stat. Phys.* **162**, 1244 (2016).
- [52] E. A. Theodorou and E. Todorov, *Relative Entropy and Free Energy Dualities: Connections to Path Integral and KL Control*, in *Proceedings of 51st IEEE Conference on Decision and Control, Maui, 2012* (IEEE, New York, 2012), pp. 1466–1473.
- [53] W. Bialek and N. Tishby, *Predictive Information*, [arXiv: cond-mat/9902341](https://arxiv.org/abs/cond-mat/9902341).
- [54] R. E. Cobb, N. Sun, and H. Zhao, *Directed Evolution as a Powerful Synthetic Biology Tool*, *Methods* **60**, 81 (2013).
- [55] G. T. Hess, J. Tycko, D. Yao, and M. C. Bassik, *Methods and Applications of CRISPR-Mediated Base Editing in Eukaryotic Genomes*, *Mol. Cell* **68**, 26 (2017).
- [56] K. Naran, T. Nundalall, S. Chetty, and S. Barth, *Principles of Immunotherapy: Implications for Treatment Strategies in Cancer and Infectious Diseases*, *Front. Microbiol.* **9**, 3158 (2018).
- [57] S. Sinai and E. D. Kelsic, *A Primer on Model-Guided Exploration of Fitness Landscapes for Biological Sequence Design*, [arXiv:2010.10614v2](https://arxiv.org/abs/2010.10614v2).
- [58] M. Luksza and M. Lässig, *A Predictive Fitness Model for Influenza*, *Nature (London)* **507**, 57 (2014).
- [59] R. A. Neher, C. A. Russell, and B. I. Shraiman, *Predicting Evolution from the Shape of Genealogical Trees*, *eLife* **3**, e03568 (2014).
- [60] J. Otwinowski, C. H. LaMont, and A. Nourmohammad, *Information-Geometric Optimization with Natural Selection*, *Entropy* **22**, 967 (2020).
- [61] A. Nourmohammad, J. Rambeau, T. Held, V. Kovacova, J. Berg, and M. Lässig, *Adaptive Evolution of Gene Expression in Drosophila*, *Cell Rep.* **20**, 1385 (2017).
- [62] B. Wittenmark, *Adaptive Dual Control Methods: An Overview*, *IFAC Proc. Vol.* **28**, 67 (1995).
- [63] G. R. Price, *Selection and Covariance*, *Nature (London)* **227**, 520 (1970).
- [64] G. J. Baxter, R. A. Blythe, and A. J. McKane, *Exact Solution of the Multi-Allelic Diffusion Model*, *Math. Biosci.* **209**, 124 (2007).
- [65] V. Mustonen and M. Lässig, *Fitness Flux and Ubiquity of Adaptive Evolution*, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 4248 (2010).
- [66] V. Mustonen and M. Lässig, *Evolutionary Population Genetics of Promoters: Predicting Binding Sites and Functional Phylogenies*, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 15936 (2005).
- [67] V. Mustonen, J. Kinney, C. G. J. Callan, and M. Lässig, *Energy-Dependent Fitness: A Quantitative Model for the Evolution of Yeast Transcription Factor Binding Sites*, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 12376 (2008).
- [68] M. Manhart and A. V. Morozov, *Protein Folding and Binding Can Emerge as Evolutionary Spandrels through Structural Coupling*, *Proc. Natl. Acad. Sci. U.S.A.* **112**, 1797 (2015).
- [69] J. Otwinowski, D. M. McCandlish, and J. B. Plotkin, *Infering the Shape of Global Epistasis*, *Proc. Natl. Acad. Sci. U.S.A.* **115**, E7550 (2018).
- [70] J. Otwinowski, *Biophysical Inference of Epistasis and the Effects of Mutations on Protein Stability and Function*, *Mol. Biol. Evol.* **5**, e23156 (2018).
- [71] K. Husain and A. Murugan, *Physical Constraints on Epistasis*, *Mol. Biol. Evol.* **37**, 2865 (2020).
- [72] J. R. Blundell and S. F. Levy, *Beyond Genome Sequencing: Lineage Tracking with Barcodes to Study the Dynamics of Evolution, Infection, and Cancer*, *Genomics* **104**, 417 (2014).
- [73] H. Risken, C. Schmid, and W. Weidlich, *Fokker-Planck Equation, Distribution and Correlation Functions for Laser Noise*, *Z. Phys.* **194**, 337 (1966).
- [74] The inner product of the two conjugate functions $\langle P_u | \Psi \rangle$ is time invariant:

$$\begin{aligned}
 \langle P_u(t') | \Psi(t') \rangle &\equiv \int d\mathbf{x}' P_u(\mathbf{x}', t' | \mathbf{x}, t) \Psi(\mathbf{x}', t') \\
 &= \langle e^{L(t'-t)} P_u(t) | e^{-L^\dagger(t'-t)} \Psi(t) \rangle \\
 &= \langle P_u(t) | e^{L^\dagger(t'-t)} e^{-L(t'-t)} \Psi(t) \rangle \\
 &\equiv \langle P_u(t) | \Psi(t) \rangle.
 \end{aligned}$$