Open access • Posted Content • DOI:10.1101/867143

Controlling the speed and trajectory of evolution with counterdiabatic driving

Shamreen Iram, Emily Dolson, Joshua Chiel, Julia Pelesko ...+9 more authors

Institutions: Case Western Reserve University, Cleveland Clinic, University of California, Berkeley, University of Maryland, College Park ...+1 more institutions

Published on: 06 Dec 2019 - bioRxiv (Cold Spring Harbor Laboratory)

Topics: Fitness landscape, Population and Context (language use)

Controlling the speed and trajectory of evolution with counterdiabatic driving

Shamreen Iram[1,*], Emily Dolson[2,*], Joshua Chiel[1,*], Julia Pelesko[1], Nikhil Krishnan[3], Özenç Güngör[1], Benjamin Kuznets-Speck[1,6], Sebastian Deffner[4], Efe Ilker[5], Jacob G. Scott[1,2,3,†], and Michael Hinczewski[1,†]

[1]Department of Physics, Case Western Reserve University, Cleveland, OH, 44106, USA
[2]Translational Hematology Oncology Research, Cleveland Clinic, Cleveland OH, 44106, USA
[3]Case Western Reserve University School of Medicine, Cleveland, OH, 44106, USA
[4]Department of Physics, University of Maryland, Baltimore County, Baltimore, MD 21250, USA
[5]Physico-Chimie Curie UMR 168, Institut Curie, PSL Research University, 75248 Paris Cedex 05, France
[6]Biophysics Graduate Group, University of California, Berkeley, CA 94720, USA
[*]equal contribution
[†]to whom correspondence should be addressed: scottj10@ccf.org, michael.hinczewski@case.edu

ABSTRACT

The pace and unpredictability of evolution are critically relevant in a variety of modern challenges: combating drug resistance in pathogens and cancer, understanding how species respond to environmental perturbations like climate change, and developing artificial selection approaches for agriculture. Great progress has been made in quantitative modeling of evolution using fitness landscapes, allowing a degree of prediction for future evolutionary histories. Yet fine-grained control of the speed and the distributions of these trajectories remains elusive. We propose an approach to achieve this using ideas originally developed in a completely different context – counterdiabatic driving to control the behavior of quantum states for applications like quantum computing and manipulating ultra-cold atoms. Implementing these ideas for the first time in a biological context, we show how a set of external control parameters (i.e. varying drug concentrations / types, temperature, nutrients) can guide the probability distribution of genotypes in a population along a specified path and time interval. This level of control, allowing empirical optimization of evolutionary speed and trajectories, has myriad potential applications, from enhancing adaptive therapies for diseases, to the development of thermotolerant crops in preparation for climate change, to accelerating bioengineering methods built on evolutionary models, like directed evolution of biomolecules.

The quest to control evolutionary processes in areas like agriculture and medicine predates our understanding of evolution itself. Recent years have seen growing research efforts toward this goal, driven by rapid progress in quantifying genetic changes across a population[1–3] as well as a global rise in challenging problems like therapeutic drug resistance[4–6]. New approaches that have arisen in response include prospective therapies that steer evolution of pathogens toward maximized drug sensitivity[7,8], typically requiring multiple rounds of selective pressures and subsequent evolution under them. Since we cannot predict the exact progression of mutations that occur in the course of the treatment, the best we can hope for is to achieve control over probability distributions of evolutionary outcomes. However, our lack of precise control over the timing of these outcomes poses a major practical impediment to engineering the course of evolution. This naturally raises a question: *Rather than being at the mercy of evolution's unpredictability and pace, what if we could simultaneously control the speed and the distribution of genotypes over time?*

Controlling an inherently stochastic process like evolution has close parallels to problems in other disciplines. Quantum information protocols crucially depend on coherent control over the time evolution of quantum states under external driving[9,10], in many cases requiring that a system remain in an instantaneous ground state of a time-varying Hamiltonian in applications like cold atom transport[11] and quantum adiabatic computation[12]. The adiabatic theorem of quantum mechanics facilitates such control when the driving is infinitely slow, but over finite time intervals control becomes more challenging, because fast driving can induce random transitions to undesirable excited states. Overcoming this challenge—developing fast processes that mimic the perfect control of infinitely slow ones—has led to a whole subfield of techniques called "shortcuts to adiabaticity"[13–18]. One such method in particular, known as transitionless, or counterdiabatic (CD) driving, involves adding an auxiliary control field to the system to inhibit transitions to excited states[19–21]. Intriguingly, the utility of CD driving is not limited

to quantum contexts: requiring a quantum system to maintain an instantaneous ground state under driving is mathematically analogous to demanding a classical stochastic system remains in an instantaneous equilibrium state as external control parameters are changed[22,23]. Extending CD driving ideas to the classical realm has already led to proof-of-concept demonstrations of accelerated equilibration in optical tweezer[24] and atomic force microscope[25] experimental frameworks, and is closely related to optimal, finite-time control problems in stochastic thermodynamics[26,27].

Here we demonstrate the first biological application of CD driving, by using it to control the distribution of genotypes in a Wright-Fisher model[28] describing evolution in a population of organisms. The auxiliary CD control field (implemented for example through varying drug concentrations or other external parameters that affect fitness) allows us to shepherd the system through a chosen sequence of genotype distributions, moving from one evolutionary equilibrium state to another in finite time. We validate the CD theory through numerical simulations using an agent-based model of evolving unicellular populations, focusing on a system where sixteen possible genotypes compete via a drug dose-dependent fitness landscape derived from experimental measurements.

1 Theory

1.1 Evolutionary model

We develop our CD driving theory in the framework of a Wright-Fisher diffusion model for the evolution of genotype frequencies in a population (see Methods for details). Let us consider M possible genotypes, where the ith genotype comprises a fraction x_i of a population. Since $\sum_{i=1}^{M} x_i = 1$, we can describe the state of the system through $M-1$ independent values of x_i, or equivalently through a frequency vector $\boldsymbol{x} = (x_1, \ldots, x_{M-1})$. Without loss of generality, we will take the Mth genotype to be the reference (the "wild type") with respect to which the relative fitnesses of the others will be defined. Let $1+s_i$ be the relative fitness of genotype $i = 1, \ldots, M-1$ compared to the wild type, where s_i is a selection coefficient, defining the ith component of a vector $\boldsymbol{s}$. We assume fitnesses are influenced by some time-dependent control parameter $\lambda(t)$, which we write as a scalar quantity, though it could in principle be a vector, reflecting a set of control parameters. These parameters could involve any environmental quantity amenable to external control: in the examples below we consider the concentration of a single drug applied to a population of unicellular organisms. However we could have more complicated drug protocols (switching between multiple drugs)[7] or other perturbations in fitness secondary to microenvironmental change (e.g. nutrient or oxygenation levels). Our control protocol $\lambda(t)$ from initial time t_0 to final time t_f defines a trajectory of the selection coefficient vector, $\boldsymbol{s}\big(\lambda(t)\big)$, shown schematically in Fig. 1A. Our population thus evolves under a time-dependent fitness landscape, or so-called "seascape"[29]. Note that all time variables, unless otherwise noted, are taken to be in units of Wright-Fisher generations.

For simplicity, the total population is assumed to be fixed at a value N, corresponding to a scenario where the system stays at a time-independent carrying capacity over the time interval of interest. (Our approach is easily generalized to more complicated cases with time-dependent $N(t)$, as shown in the Supplementary Information [SI]). The final quantity characterizing the dynamics is an $M \times M$ dimensional mutation rate matrix m, where each off-diagonal entry $m_{\beta\alpha}$ represents the mutation probability (per generation) from the αth to the βth genotype. For later convenience, the αth diagonal entry of m is defined as the opposite of the total mutation rate out of that genotype, $m_{\alpha\alpha} \equiv -\sum_{\beta\neq\alpha} m_{\beta\alpha}$. As in the case of N, we assume the matrix m is time-independent, though this assumption can be relaxed.

1.2 Driving the genotype frequency distribution

Given the system described above, we focus on $p(\boldsymbol{x},t)$, the probability to find genotype frequencies $\boldsymbol{x}$ at time t, calculated over an ensemble of possible evolutionary trajectories. The dynamics of this probability for the WF model can be described to an excellent approximation through a Fokker-Planck equation:

$$\partial_t p(\boldsymbol{x},t) = \mathcal{L}\big(\lambda(t)\big) p(\boldsymbol{x},t), \tag{1}$$

where $\partial_t \equiv \partial/\partial t$ and $\mathcal{L}\big(\lambda(t)\big)$ is a differential operator, acting on functions of $\boldsymbol{x}$, described in the Methods. This operator involves N, m, and $\boldsymbol{s}\big(\lambda(t)\big)$, and we highlight the dependence on $\lambda(t)$. In setting up the analogy to driving in quantum mechanics, Eq. (1) corresponds to the Schrödinger equation, with $p(\boldsymbol{x},t)$ playing the role of the wavefunction and $\mathcal{L}\big(\lambda(t)\big)$ the time-dependent Hamiltonian operator. The full analogy between quantum and evolutionary dynamics is described in more detail in Box 1 of the Methods. Though for our purposes we only employ this analogy qualitatively, in fact there exists in certain cases an explicit mapping from the Fokker-Planck to the

Schrödinger equation (though not vice versa)[30–32]. For a particular value of the control parameter λ, the analogue of the quantum ground state wavefunction is the eigenfunction $\rho(\boldsymbol{x};\lambda)$ with eigenvalue zero, the solution of the equation

$$\mathcal{L}(\lambda)\rho(\boldsymbol{x};\lambda) = 0. \tag{2}$$

In the evolutionary context, $\rho(\boldsymbol{x};\lambda)$ has an additional meaning with no direct quantum correspondence: it is the *equilibrium probability distribution of genotypes*. If one fixes the control parameter $\lambda(t) = \lambda$, the distribution $p(\boldsymbol{x},t)$ obeying Eq. (1) will approach $\rho(\boldsymbol{x};\lambda)$ in the limit $t \to \infty$.

Consider the following control protocol, where we start at one control parameter value, $\lambda(t) = \lambda_0$ for $t \leq t_0$, and finish at another value, $\lambda(t) = \lambda_f$ for $t \geq t_f$, with some arbitrary driving function $\lambda(t)$ in the interval $t_0 < t < t_f$. We assume the system starts in one equilibrium distribution, $p(\boldsymbol{x},t_0) = \rho(\boldsymbol{x};\lambda_0)$, and we know that it will eventually end at a new equilibrium, $p(\boldsymbol{x},t) \to \rho(\boldsymbol{x};\lambda_f)$ for $t \gg t_f$. But what happens at intermediate times? If $\lambda(t)$ changes infinitesimally slowly during the driving (and hence $t_f \to \infty$) then the system would remain at each moment in the corresponding instantaneous equilibrium (IE) distribution, $p(\boldsymbol{x},t) = \rho\big(\boldsymbol{x};\lambda(t)\big)$ for all t. This result, derived in the SI, is the analogue of the quantum adiabatic theorem[33] applied to the ground state: for a time-dependent Hamiltonian that changes extremely slowly, a quantum system that starts in the ground state of the Hamiltonian always remains in the same instantaneous ground state (assuming that at all times there is a gap between the ground state energy and the rest of the energy spectrum). Fig. 1B shows schematic snaphsots of $\rho\big(\boldsymbol{x};\lambda(t)\big)$ at three times, with the control parameter shifting them across the genotype frequency space.

When the driving occurs over finite times ($t_f < \infty$), the above results break down: $p(\boldsymbol{x},t) \neq \rho\big(\boldsymbol{x};\lambda(t)\big)$ for $0 < t < t_f$, but is instead a linear combination of many instantaneous eigenfunctions of the Fokker-Planck operator, just as the corresponding quantum system under faster driving will generically evolve into a superposition of the instantaneous ground state and excited states. This will manifest itself as a lag, with $p(\boldsymbol{x},t)$ moving towards but not able to catch up with $\rho\big(\boldsymbol{x};\lambda(t)\big)$, as illustrated in Fig. 1C. For $t > t_f$, once $\lambda(t)$ stops changing, the system will eventually settle into equilibrium at $\rho(\boldsymbol{x};\lambda_f)$ in the long time limit.

1.3 Control and counterdiabatic driving

This lag can be an obstacle if one wants to control the evolution of the system over finite time intervals. Since evolutionary trajectories are stochastic, we cannot necessarily guarantee that the system starts and ends at precise genotype frequencies, but we can attempt to specify initial and final target frequency distributions. At the end of the driving $t = t_f$, we would like our system to arrive at the target distribution, and then stay there as long as the control parameter is fixed. In this way we complete one stage of the control protocol, and have a known starting point for the next stage, since in practice we could imagine the interval $t_0 < t < t_f$ as just one step of a multi-stage protocol involving distinct interventions (i.e. a sequence of different drugs). Completing each stage as quickly as possible, while accurately hitting each target, would for example be a crucial prerequisite to translating certain evolutionary medicine approaches to clinical settings (see SI Sec. 8 for a fuller discussion). Thus, if we were enumerating the characteristics of an ideal control mechanism, at the very least it should be able to drive the system from one equilibrium distribution, $p(\boldsymbol{x},t_0) = \rho(\boldsymbol{x};\lambda_0)$, to another, $p(\boldsymbol{x},t_f) = \rho(\boldsymbol{x};\lambda_f)$, over a finite time $t_f - t_0$.

In the context of quantum adiabatic computing[12], the typical focus is on the initial ground state (which has to be easy to realize experimentally) and the final ground state (since it encodes the solution to the computational problem). In the evolutionary case, we can imagine additional desired characteristics for our driving, beyond the start and end-point distributions. There are many ways to go from an initial fitness landscape, $\boldsymbol{s}(\lambda_0)$, to a final fitness landscape, $\boldsymbol{s}(\lambda_f)$, corresponding to different possible trajectories in the selection coefficient space of Fig. 1A that share initial and final values. Depending on how we empirically implement the control, many of these trajectories may be physically inaccessible. But among the remaining set of realizable trajectories, some may be more desirable than others (i.e. have different evolutionary consequences[34], or trade-offs[35]). Each trajectory defines a continuous sequence of IE distributions $\rho\big(\boldsymbol{x};\lambda(t)\big)$, and for each distribution there is a mean genotype frequency $\overline{\boldsymbol{x}}\big(\lambda(t)\big)$, illustrated in the lower half of Fig. 1B. We may, for example, want protocols that minimize the chances of our system visiting certain problematic genotypes: in practice this could translate to demanding that the curve $\overline{\boldsymbol{x}}\big(\lambda(t)\big)$ for $t_0 < t < t_f$ stays far away from certain regions of the genotype frequency space. This in turn restricts the $\boldsymbol{s}\big(\lambda(t)\big)$ trajectories and hence the protocols $\lambda(t)$ of practical interest. In simpler terms, we would ideally like to control not just the distributions at the beginning and end of the driving, but also if possible along the way.

We formulate this ideal control problem in the following way: we demand that $p(\boldsymbol{x},t) = \rho\big(\boldsymbol{x};\lambda(t)\big)$ for some chosen control protocol $\lambda(t)$ between $t_0 < t < t_f$. The protocol $\lambda(t)$ is determined with the above considerations in mind, and thus defines a particular path through the space of genotype frequency distributions over which we would like to

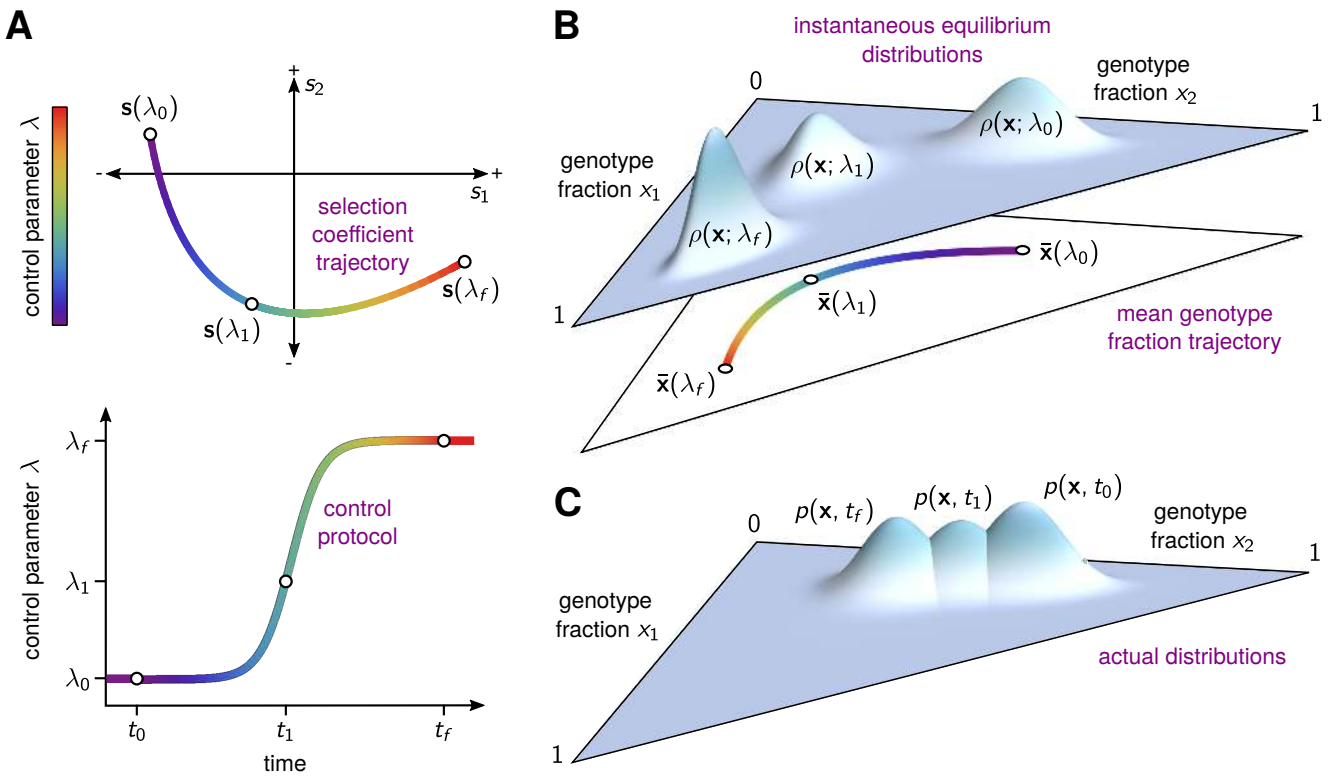

Figure 1. **(A)** A schematic illustration for three genotypes ($M = 3$) showing the trajectory of the selection coefficient vector $\boldsymbol{s}\big(\lambda(t)\big)$ as a function of a time-varying control parameter $\lambda(t)$ depicted in the bottom of the panel. This represents the fitness "seascape" under which the population evolves during driving. Three time points are highlighted: the initial time t_0, an intermediate time t_1, and the final time t_f, where the corresponding control parameter values are λ_0, λ_1, and λ_f. The amplitude of the control parameter along the trajectory is represented through a color gradient. **(B)** The instantaneous equilibrium (IE) distribution of genotypes $\rho\big(\boldsymbol{x};\lambda(t)\big)$ for the three highlighted values of the control parameter from panel **A**. These distributions are probability densities on the 2D simplex defined by $x_1 + x_2 \leq 1$ and $x_1,\, x_2 \geq 0$. In the lower part of the panel we show the curve of mean IE genotype frequencies $\overline{\boldsymbol{x}}\big(\lambda(t)\big)$. **(C)** For driving over finite times, the actual distribution of genotypes $p(\boldsymbol{x},t)$ will generally lag behind the IE distribution while the control parameter is changing. Thus at t_f the distribution $p(\boldsymbol{x},t_f)$ is still far from $\rho(\boldsymbol{x};\lambda_f)$, and will only catch up with it at times $t \gg t_f$ as the system re-equilibrates.

guide our system. Clearly we will not achieve success by just directly implementing $\lambda(t)$, since $p(\boldsymbol{x},t)$ obeying Eq. (1) will generally lag behind $\rho\big(\boldsymbol{x};\lambda(t)\big)$[36]. The resolution of this problem in the quantum case through CD driving is to add a specially constructed auxiliary time-dependent Hamiltonian to the original Hamiltonian[19–21]. For a specific choice of this auxiliary Hamiltonian, we can guarantee that our new system always remains in the instantaneous ground state of the original Hamiltonian. The evolutionary analogue of CD is to replace the Fokker-Planck operator $\mathcal{L}\big(\lambda(t)\big)$ in Eq. (1) with a different operator $\widetilde{\mathcal{L}}\big(\lambda(t),\dot{\lambda}(t)\big)$, which depends on both $\lambda(t)$ and its time derivative $\dot{\lambda}(t) \equiv d\lambda(t)/dt$. This CD operator satisfies

$$\partial_t \rho\big(\boldsymbol{x};\lambda(t)\big) = \widetilde{\mathcal{L}}\big(\lambda(t),\dot{\lambda}(t)\big)\rho\big(\boldsymbol{x};\lambda(t)\big). \tag{3}$$

Thus by construction, $p(\boldsymbol{x},t) = \rho\big(\boldsymbol{x};\lambda(t)\big)$ is a solution to the Fokker-Planck equation with the new operator. Additionally, to be consistent with the slow adiabatic driving limit discussed above, $\widetilde{\mathcal{L}}\big(\lambda(t),0\big) = \mathcal{L}\big(\lambda(t)\big)$, so we recover the original Fokker-Planck operator when the speed of driving $\dot{\lambda}(t) \to 0$ and $t_f \to \infty$.

Of course defining $\widetilde{\mathcal{L}}\big(\lambda(t),\dot{\lambda}(t)\big)$ in this way is the easy part: figuring out how to implement a new control protocol to realize $\widetilde{\mathcal{L}}\big(\lambda(t),\dot{\lambda}(t)\big)$ is more challenging. In the Methods, we show how the most general solution to go from $\mathcal{L}$ to $\widetilde{\mathcal{L}}$ is to replace the original selection coefficient trajectory $\boldsymbol{s}\big(\lambda(t)\big)$ with a frequency-dependent version, $\tilde{\boldsymbol{s}}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)$.

Implementing a particular frequency dependent fitness seascape is a degree of control that is generally impossible in realistic scenarios. Fortunately, we show that in one important parameter regime the CD seascape becomes approximately frequency-independent, $\tilde{\boldsymbol{s}}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big) \approx \tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)$. This occurs in the large population, frequent mutation regime: if the typical mutation rate scale is μ, meaning $m_{\beta\alpha} \sim \mathcal{O}(\mu)$ for all nonzero mutation rates where $\alpha \neq \beta$, then this corresponds to $\mu N \gg 1$, $N \gg 1$[37–40]. In this regime multiple genotypes can generally coexist in the population at equilibrium (though one may be quite dominant), which is particularly relevant for pathogenic populations, especially ones spreading through space[41–43]. Remarkably, there is a simple analytical expression that provides an excellent approximation to $\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)$ in this case:

$$\tilde{s}_i\big(\lambda(t),\dot{\lambda}(t)\big) \approx s_i\big(\lambda(t)\big) + \frac{d}{dt}\ln\frac{\overline{x}_i\big(\lambda(t)\big)}{\overline{x}_M\big(\lambda(t)\big)}, \qquad i = 1,\ldots,M-1, \tag{4}$$

where $\overline{x}_M\big(\lambda(t)\big) \equiv 1 - \sum_{i=1}^{M-1}\overline{x}_i\big(\lambda(t)\big)$. We see that the new selection coefficient protocol is defined through the target mean genotype frequency trajectory $\overline{\boldsymbol{x}}\big(\lambda(t)\big)$, and reduces to the original protocol when $\dot{\lambda}(t) \to 0$. Moreover, as we show in the examples below for specific systems, Eq. (4) can at least in certain cases be implemented through physically realistic manipulations of the environment, like time-varying drug dosages. While we focus on the frequent mutation regime in the current work, the applicability of CD ideas is not limited to just this regime: for the opposite case of infrequent mutations, $\mu N \ll 1$, where the evolutionary dynamics can be modeled as a sequence of mutant fixations, one can also formulate a CD theory based on a discrete Markov state description[44] (see our follow-up article[45]).

2 Results

2.1 Two genotypes

The simplest example of our CD theory is for a two genotype ($M = 2$) system, where the dynamics are one dimensional, described by a single frequency x_1 and selection coefficient $s_1(\lambda(t))$. As shown in the Methods, Eq. (4) in this case can be evaluated analytically. To illustrate driving, we assume a control protocol $\lambda(t)$ such that the selection coefficient increases according to a smooth ramp (the original protocol in Fig. 2B). This starts from zero at t_0 (both genotypes have equal fitness) and increases until reaching a plateau at a final selection coefficient that favors genotype 1. Fig. 2A shows $p(x_1,t)$ from a numerical solution of Eq. (1) using this protocol, compared against the IE distribution $\rho\big(x_1;\lambda(t)\big)$, solved using Eq. (2), at three time snapshots. To validate the Fokker-Planck approach, we also designed an agent-based model (ABM), described in the Methods section 4.5, which simulates the individual life trajectories of an evolving population of cells. Because there exists a mapping between the parameters of the ABM and the equivalent Fokker-Planck equation (Methods section 4.5.3), one can directly compare the $p(x_1,t)$ results from the ABM simulations (circles) to the Fokker-Planck numerical solution of Eq. (1) (curves), which show excellent agreement. In the absence of CD driving, as expected, $p(x_1,t)$ lags behind $\rho\big(x_1;\lambda(t)\big)$, with the latter shifting rapidly to larger x_1 frequencies as the fitness of genotype 1 increases.

To eliminate this lag, we implement the alternative selection coefficient trajectory of Eq. (20). Fig. 2B shows a comparison between $\tilde{s}_1\big(\lambda(t),\dot{\lambda}(t)\big)$ and the original $s_1\big(\lambda(t)\big)$. We see that the CD intervention requires a transient overshoot of the selection coefficient during the driving, nudging $p(x_1,t)$ to keep up with $\rho\big(x_1;\lambda(t)\big)$. Panel C shows the same snapshots as in panel A, but now with CD driving: we see the actual and IE distributions nearly perfectly overlap at all times. To quantify the effectiveness of the CD protocol, we measure the degree of overlap through the Kullback-Leibler (KL) divergence[36,46], defined for any two probability distributions $p(\boldsymbol{x})$ and $q(\boldsymbol{x})$ as $D_{\mathrm{KL}}(p||q) = \int d\boldsymbol{x}\, p(\boldsymbol{x})\log_2\big(p(\boldsymbol{x})/q(\boldsymbol{x})\big)$. Expressed in bits, the KL divergence is always ≥ 0, and equals 0 for identical distributions. Panel D shows $D_{\mathrm{KL}}(\rho||p)$ for both the original and CD protocols, with the latter dramatically reducing the divergence across the time interval of driving.

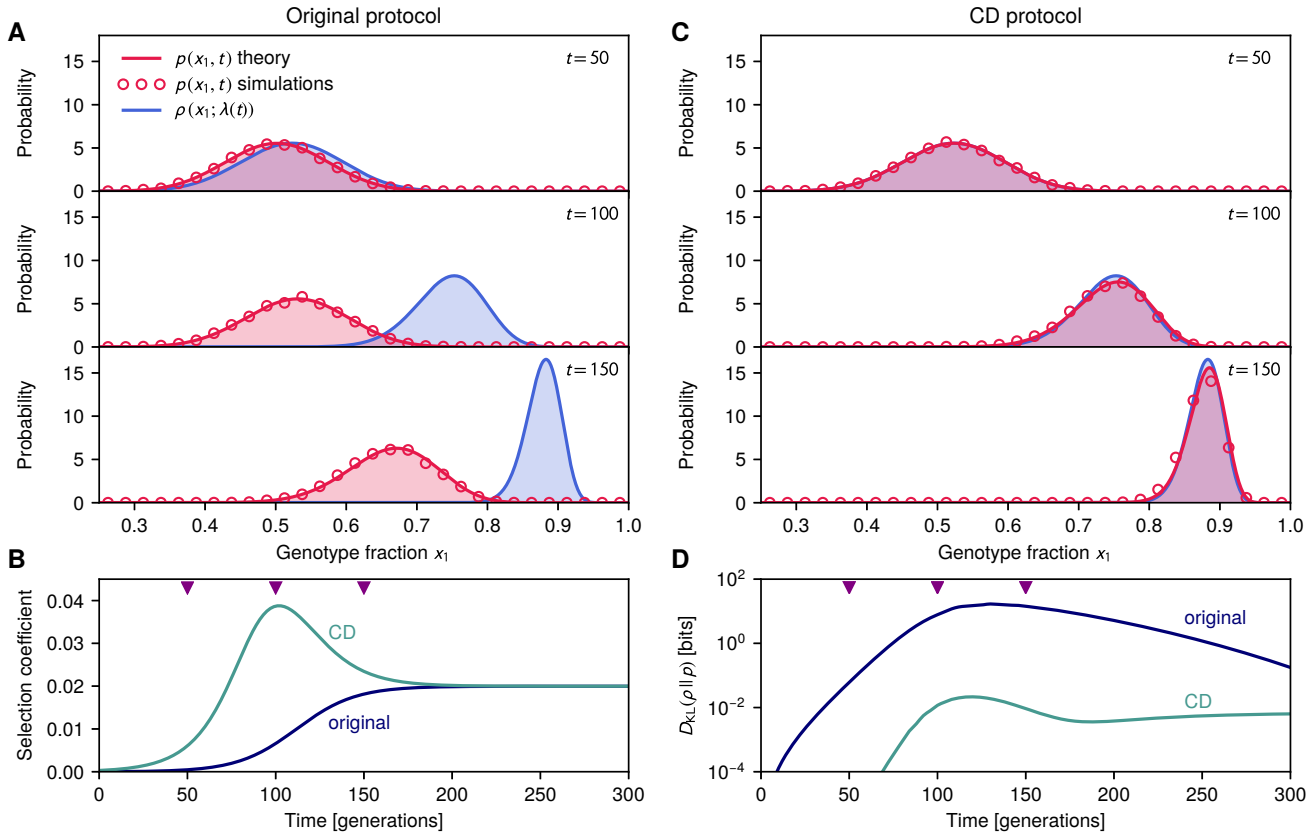

Figure 2. Results for two genotypes. (A,C) We plot three time snapshots of the actual probability distribution $p(x_1,t)$ versus the IE distribution $\rho\big(x_1;\lambda(t)\big)$ for driving with the original control protocol **(A)** and with the CD driving protocol **(C)**, where x_1 is the fraction of genotype 1 in the population. Solid red curves are numerical solutions of the Fokker-Planck Eq. (1) for $p(x_1,t)$, red circles are agent-based simulations. Without CD driving, the actual distribution always lags behind the IE. **B)** The selection coefficient trajectory s_1 for the original control protocol (dark blue) versus the corresponding CD prescription (green) $\tilde{s}_1$ from Eq. (20). The three snapshot times (50, 100, 150 generations) are indicated by triangles. **(D)** Kullback-Leibler divergence between actual and IE distributions versus time, with and without CD driving, calculated using the numerical Fokker-Planck solution.

2.2 Multiple genotypes via agent-based modeling

The ABM simulations also allow us to test the CD theory in more complex scenarios. To this end we considered a system with 16 genotypes (4 alleles), with selection coefficients based on a well characterized experimental system: the fitness effects of the anti-malarial drug pyrimethamine at varying concentrations on all possible combinations of four different drug-resistance alleles[2,3]. Our control parameter $\lambda(t)$ is the drug concentration, and we implement the seascape by increasing the drug over time (after an initial equilibration period), eventually saturating at a concentration of 10^{-4} M (the protocol labeled "original" in Fig. 3E). With our choice of simulation parameters (Methods section 4.5.2), a number of the genotypes have sufficient resistance to survive even at higher drug dosages, so the overall population remains at carrying capacity. What changes as the dosage increases is the distribution of genotypes. Fig. 3A and 3B show the results in the absence of CD driving, with each genotype labeled by a 4-digit binary sequence. The population goes from being dominated by 1110 (with smaller fractions of other genotypes) to eventually becoming dominated by 1111. However there is a dramatic lag behind the IE distribution, taking more than 1500 generations to resolve. This is quantified in the KL divergence $D_{\mathrm{KL}}(\rho||p)$ in Fig. 3F, which rapidly increases by 5 orders of magnitude as the drug ramp starts showing its effects (around generation 500). Equilibration to the higher drug dosages brings the divergence back down over time, but it only achieves relatively small amplitudes after generation 2000. Note that the scale of the KL divergences for 15-dimensional probability distributions is larger than for the 1D example in the previous section: this reflects both the greater sensitivity of the KL measure to small discrepancies in a 15-dimensional space, as well as the fact that distributions estimated from an ensemble of simulations (1000 indepedent runs in this case) will always have a degree of sampling error. Thus it is more instructive to look at the relative change of the KL with driving rather than the absolute magnitudes.

To reduce the lag through CD driving, one should in principle implement the selection coefficients according to Eq. (4). However this involves guiding the system along a fitness trajectory in a 15-dimensional space, and in this case we have a single tuning knob (the concentration of pyrimethamine) to perturb fitnesses. In such scenarios one then looks at the closest approximation to CD driving that can be achieved with the experimentally accessible control parameters. In this particular case the genotypes which dominate the population at small and large drug concentrations are 1110 ($i = 15$) and the wild-type 1111 ($i = 16$), so the selection coefficient s_{15} which encodes their fitness relative to each other plays the most important role in the dynamics. We thus choose a CD drug dosage by numerically solving for the concentration that most closely approximates the $i = 15$ component of Eq. (4) at each time. Because in real-world scenarios there will be limits on the maximum allowable dosage, we constrain the CD concentrations to be below a certain cutoff. The approximation described here, where two different genotypes dominate at different times during the driving, is just a special case of a more general approximation approach where we seek to achieve the closest possible protocol to the one described by Eq. (4), given the experimental constraints. In SI Sec. 9 we illustrate how this general strategy works in two additional 16-genotype seascape examples (including the empirical seascape for the drug cycloguanil[2]) where more than two genotypes dominate during driving.

Fig. 3E shows CD drug protocols with three different cutoffs: 10^{-2}, 10^{-3}, 5×10^{-4} M (all within the experimentally measured dosage range). The higher the cutoff, the better the approximation to CD driving. We can directly quantify the overall reduction in lag time Δt due to CD from the KL divergence results of Fig. 3F, as explained in Methods section 4.5.4. For a cutoff of 10^{-2} M the lag is reduced by $\Delta t = 1210$ generations. Notably, though the approximation is based on the top two genotypes (1110 and 1111), it reduces the lag time across the board for all genotypes (see Fig. 3C for four representative genotype trajectories at 10^{-2} M cutoff, with other genotypes shown in the snapshots of Fig. 3D). This is because driving of the top two also entrains the dynamics of the subdominant genotypes whose populations are sustained by mutations out of and into the dominant ones. Even with the more restrictive constraint of 5×10^{-4} M there is still a substantial benefit, with the lag reduced by $\Delta t = 656$ generations. This highlights the robustness of the CD approach: even if one cannot implement the solution of Eq. (4) exactly, we can still arrive at the target distribution faster through an approximate CD protocol.

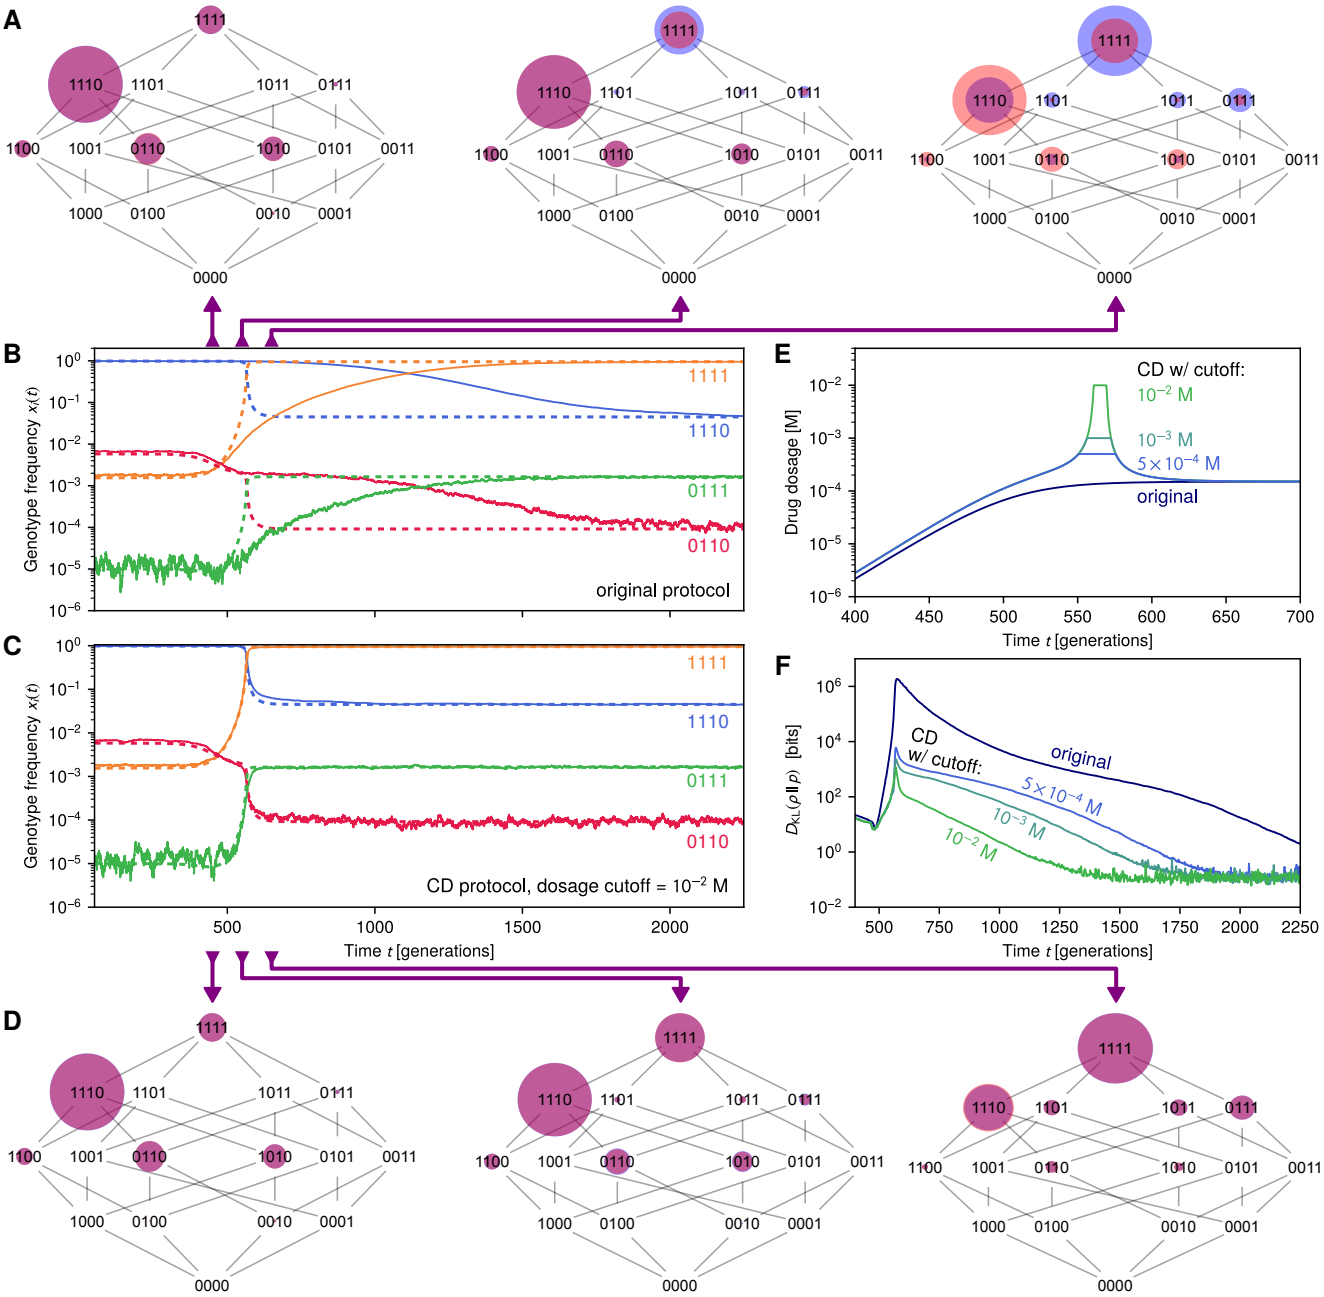

Figure 3. CD driving eliminates evolutionary lag in 16 genotype simulation. **(A,D)** We plot three snapshots of our evolving agent-based population model without **(A)** and with **(D)** CD driving. Each of the three 16 genotype (4 binary alleles) hypercubic graphs ("tesseracts") has vertices with log-scaled radii representing the fraction of each genotype in the total population at a given time. Orange is the actual fraction, blue the IE fraction, and the overlap appears purple. The CD driving in this case is implemented approximately through a drug dosage protocol (panel **E**) with cutoff 10^{-2} M. **(B,C)** Corresponding sample simulation trajectories (solid lines) versus IE expectation (dashed lines) for the fraction of 4 representative genotypes without **(B)** and with **(C)** CD driving. The latter significantly reduces the nearly 1500 generation lag. **(E)** Original drug protocol versus the CD protocol with different possible dosage cutoffs. **(F)** Kullback-Leibler divergence between actual and IE distributions versus time, with and without CD driving. For the latter, increasing the dosage cutoff value makes the protocol more closely approximate the true CD solution, and hence decreases the divergence.

3 Discussion and Conclusions

Our demonstration of the CD driving approach in a population model with empirically-derived drug-dependent fitnesses shows that we can accelerate evolution toward a target distribution *in silico*. As new technologies progressively allow us to assemble ever more extensive fitness landscapes for various organisms as a function of external perturbations like drugs[1–3], the next step is implementing CD driving in the laboratory. This would be a necessary milestone on the path to a range of potential applications (the latter discussed in more detail in SI Sec. 8). Thus it is worth considering the challenges and potential workarounds that will be involved in experimental applications.

One salient issue is the range of control parameters available in laboratory settings. Our examples have focused on the simplest cases of one-dimensional control, but to access the full power of the CD approach presented here we should explore a richer parameter space: not only single drugs, but combinations, along with varying nutrients, metabolites, oxygen levels, osmotic pressure, and temperature. The eventual goal would be to have for every system a library of well-characterized interventions that could be applied in tandem, allowing us the flexibility to map out desired target trajectories through a multidimensional fitness landscape. In other words for a given system we would have access to a selection coefficient function $\boldsymbol{s}(\boldsymbol{\lambda}(t))$, where $\boldsymbol{\lambda}(t) = (\lambda_1(t), \lambda_2(t), \ldots)$ is a multidimensional vector of control parameters at time t: $\lambda_1(t)$ the concentration of one drug, $\lambda_2(t)$ the concentration of another drug (or nutrient), and so on. More generally, one could explore how fundamental differences among fitness landscapes (i.e. the difficulties in reaching local optima in so-called "hard" landscapes[47]) influences the types of interventions needed to achieve driving and their effectiveness.

Even given accurately measured fitnesses, one might be hampered by imperfect estimation of other system parameters, for example mutation rates. To determine how large the margin for error is, we tested the CD driving prescription calculated using incorrect mutation rates, varying the degree of discrepancy over two orders of magnitude (see SI for details). While such discrepancies do reduce the efficacy of CD driving, leading to deviations between actual and IE distributions at intermediate times, populations driven with an incorrect protocol still reached the target distribution faster than in the absence of driving. As in the case of the dosage cutoff discussed above, the CD approach has a degree of robustness to errors in the protocol, which increases its chances of success in real-world settings.

But what if we lacked measurements of the underlying fitness seascape? Interestingly, there might still be some utility of the CD method even in this case. We could first do a preliminary quasi-adiabatic experimental trial: vary external parameter(s) λ extremely gradually, and use sequencing at regularly spaced intervals to determine the quasi-equilibrium mean genotype fractions $\overline{x}_i(\lambda)$ as a function of λ. If we now wanted to guide the system through the same sequence of evolutionary distributions but much faster, we have enough information to approximately evaluate the CD perturbation in Eq. (4), which just depends on $\overline{x}_i(\lambda)$ and the rate $\dot{\lambda}(t)$ that we would like to implement. So at the very least the CD prescription could be estimated, providing a blueprint, and the remaining challenge would be figuring out what combination of external perturbations would yield the right sorts of fitness perturbations to achieve CD driving.

"Nothing makes sense in biology except in the light of evolution" is an oft-quoted maxim which was the title of a 1973 essay by Theodosius Dobzhansky[48]. However evolution is not just the fundamental paradigm through which we can understand living systems, but also a framework by which we can shape and redesign nature at a variety of scales: from engineering new proteins[49] and aptamers[50,51] to combating drug resistance in pathogens and cancer[7] to the development of crops that can withstand climate stress[52]. In all its manifestations, natural and synthetic, evolution is a stochastic process that occurs across a wide swath of timescales. Our work represents a significant step toward more precise control of both the distribution of possible outcomes and the timing of this fundamental process.

4 Methods

Box 1: Analogies between quantum physics and evolutionary dynamics

Here we summarize the connections between quantum and evolutionary concepts used in our theory. Each numbered item in the quantum column on the left has its analogue in the evolutionary column on the right.

Quantum physics

1. **Wavefunction:** describes the state of a quantum system. For a simple quantum particle in a spatial region described by coordinates $\boldsymbol{x}$, this is a function $\psi(\boldsymbol{x},t)$ whose squared amplitude $|\psi(\boldsymbol{x},t)|^2$ is the probability density of finding the particle at $\boldsymbol{x}$ at time t.

2. **Hamiltonian operator:** a differential operator $\mathcal{H}(\lambda(t))$ depending on control parameters $\lambda(t)$ defined below. It describes how the wavefunction changes in time through the **time-dependent Schrödinger equation**, $i\hbar\partial_t\psi(\boldsymbol{x},t) = \mathcal{H}(\lambda(t))\psi(\boldsymbol{x},t)$, where $\hbar$ is the reduced Planck constant. $\mathcal{H}(\lambda(t))$ involves terms that correspond to the kinetic and potential energies of the quantum particle.

3. **Control parameters:** a set of parameters $\lambda(t)$ that can be manipulated over time by an experimentalist. These parameters modify the kinetic/potential energy terms in the Hamiltonian, and thus influence the quantum dynamics. An example of this would be the magnitude of an externally applied electromagnetic field.

4. **Ground state:** the lowest energy state of a quantum system. In general, for a Hamiltonian $\mathcal{H}(\lambda)$ and given parameter values λ, the energy states (labeled by $n = 0,1,2,\ldots$) correspond to solutions of the time-independent Schrödinger equation: $\mathcal{H}(\lambda)\psi_n(\boldsymbol{x};\lambda) = E_n(\lambda)\psi_n(\boldsymbol{x};\lambda)$. Here $E_n(\lambda)$ and $\psi_n(\boldsymbol{x};\lambda)$ are the energy and wavefunction respectively of the nth state. The energies $E_0 < E_1 < \cdots$, and the ground state corresponds to $n = 0$. If λ is fixed, a system whose wavefunction is $\psi_n(\boldsymbol{x};\lambda)$ will be stationary, with its wavefunction not changing in time.

5. **Adiabatic theorem:** if we start in the nth energy state, $\psi(\boldsymbol{x},t=0) = \psi_n(\boldsymbol{x};\lambda(0))$ for some initial control parameters $\lambda(0)$, and then vary $\lambda(t)$ infinitesimally slowly (adiabatically) the theorem states that the wavefunction at later times remains in nth energy state corresponding to the instantaneous value of the parameters, $\psi(\boldsymbol{x},t) = \psi_n(\boldsymbol{x};\lambda(t))$. This is true so long as there is always a nonzero difference between $E_n(\lambda(t))$ and any $E_m(\lambda(t))$ for $m \neq n$ at all t.

Evolutionary dynamics

1. **Genotype probability distribution:** the distribution $p(\boldsymbol{x},t)$ of genetic variants (genotypes) in a population of organisms at time t, where $\boldsymbol{x}$ is a vector of genotype fractions.

2. **Fokker-Planck operator:** a differential operator $\mathcal{L}(\lambda(t))$ depending on control parameters $\lambda(t)$ defined below. It describes how the genotype probability $p(\boldsymbol{x},t)$ changes in time through the **Fokker-Planck equation**, $\partial_t p(\boldsymbol{x},t) = \mathcal{L}(\lambda(t))p(\boldsymbol{x},t)$. The full form of $\mathcal{L}(\lambda(t))$ [Eqs. (5)-(7)] involves terms that describe the mean change in genotype fractions due to mutations and selection, as well as fluctuations due to genetic drift.

3. **Control parameters:** a set of parameters $\lambda(t)$ that can be manipulated over time by an experimentalist. These parameters modify genotype fitnesses, and hence influence evolutionary dynamics through the selection terms in the Fokker-Planck operator. An example would be the concentration of a drug applied to a microbial population, where different genotypes exhibit different degrees of resistance against the drug depending on the concentration.

4. **Equilibrium state:** for a given set of parameter values λ, this is the genotype probability $\rho(\boldsymbol{x};\lambda)$ which would remain unchanged in time (stationary) during evolutionary dynamics under fixed λ. In general a Fokker-Planck operator $\mathcal{L}(\lambda)$ has a set of eigenfunctions $\psi_n(\boldsymbol{x};\lambda)$ and eigenvalues $-\kappa_n(\lambda) \leq 0$ for $n = 0,1,2\ldots$ defined through the following equation: $\mathcal{L}(\lambda)\psi_n(\boldsymbol{x};\lambda) = -\kappa_n(\lambda)\psi_n(\boldsymbol{x};\lambda)$. The equilibrium state corresponds to $n = 0$, with eigenvalue $-\kappa_0(\lambda) = 0$ and $\rho(\boldsymbol{x};\lambda) \equiv \psi_0(\boldsymbol{x};\lambda)$.

5. **Adiabatic theorem:** If we start at equilibrium $p(\boldsymbol{x},t=0) = \rho(\boldsymbol{x};\lambda(0))$ for some initial control parameters $\lambda(0)$, and then vary $\lambda(t)$ infinitesimally slowly (adiabatically) the theorem (derived in SI Sec. 1) states that at later times we will always remain in the equilibrium state corresponding to the instantaneous value of the parameters, $p(\boldsymbol{x},t) = \rho(\boldsymbol{x};\lambda(t))$.

4.1 Fokker-Planck description of Wright-Fisher evolutionary model

The underlying evolutionary dynamics of our model are based on the canonical haploid Wright-Fisher (WF) model with mutation and selection, and we adopt the formalism of recent approaches[29,53] that generalized Kimura's original two-allele diffusion theory[54] to the case of multiple genotypes. A convenient feature of the WF formalism is that other, more detailed descriptions of the population dynamics (for example agent-based models that track the life histories of individual organisms) can often be mapped onto an effective WF form, as we illustrate below.

The starting point of the Fokker-Planck diffusion approximation[29,53,54] for evolutionary population dynamics is the assumption that genotype frequencies change only by small amounts in each generation. Thus we can take the genotype frequency vector $\boldsymbol{x}$ to be a continuous variable that follows a stochastic trajectory. The key quantities describing these stochastic dynamics are the lowest order moments of $\delta\boldsymbol{x}$, the change in genotype frequency per generation. We will denote the mean of the change in the ith genotype, δx_i, taken over the ensemble of possible trajectories, as $v_i\big(\boldsymbol{x};\lambda(t)\big) \equiv \langle \delta x_i \rangle$. Note that in general $v_i\big(\boldsymbol{x};\lambda(t)\big)$ will be a function of the genotype frequencies $\boldsymbol{x}$ at the current time step, and also have a dependence on the control parameter $\lambda(t)$ through the selection coefficient vector $\boldsymbol{s}\big(\lambda(t)\big)$ (which influences $\langle \delta x_i \rangle$).

In non-evolutionary contexts $v_i\big(\boldsymbol{x};\lambda(t)\big)$ is called the drift function, but here we will call it a velocity function to avoid confusion with genetic drift. Similarly we will introduce an $(M-1)\times(M-1)$ diffusivity matrix $D_{ij}(\boldsymbol{x})$ to describe the covariance of the genotype changes, defined through $2D_{ij}(\boldsymbol{x}) \equiv \langle \delta x_i \delta x_j \rangle - \langle \delta x_i \rangle \langle \delta x_j \rangle$. As shown in the SI, to lowest order approximation $D_{ij}(\boldsymbol{x})$ is independent of $\boldsymbol{s}\big(\lambda(t)\big)$, and hence is not an explicit function of $\lambda(t)$. If we are interested in the dynamics on time scales much larger than a single generation, the probability $p(\boldsymbol{x},t)$ to observe a genotype state $\boldsymbol{x}$ at time t obeys a multivariate Fokker-Planck equation[55],

$$\begin{aligned} \partial_t p(\boldsymbol{x},t) &= -\partial_i \big(v_i\big(\boldsymbol{x};\lambda(t)\big)p(\boldsymbol{x},t)\big) + \partial_i\partial_j\big(D_{ij}(\boldsymbol{x})p(\boldsymbol{x},t)\big) \\ &\equiv \mathcal{L}\big(\lambda(t)\big)p(\boldsymbol{x},t), \end{aligned} \tag{5}$$

where $\partial_t \equiv \partial/\partial t$ and $\partial_i \equiv \partial/\partial x_i$. Note that we use Einstein summation notation, where repeated indices are summed over, and furthermore designate Greek indices to range from 1 to M while Roman indices range from 1 to $M-1$. So for example the term $\partial_i\partial_j\big(D_{ij}(\boldsymbol{x})p(\boldsymbol{x},t)\big) \equiv \sum_{i=1}^{M-1}\sum_{j=1}^{M-1}\partial_i\partial_j\big(D_{ij}(\boldsymbol{x})p(\boldsymbol{x},t)\big)$. The right-hand-side of Eq. (5) defines the Fokker-Planck differential operator $\mathcal{L}\big(\lambda(t)\big)$ in main text Eq. (1). In order to correspond to genotype fractions, the vectors $\boldsymbol{x}$ have to lie in the $(M-1)$-dimensional simplex Δ defined by the conditions $x_i \geq 0$ for all i and $\sum_{j=1}^{M-1} x_j \leq 1$. If $\boldsymbol{x} \in \Delta$, then the wild type fraction $x_M = 1 - \sum_{j=1}^{M-1} x_j$ lies between 0 and 1. Normalization of $p(\boldsymbol{x},t)$ takes the form $\int_\Delta d\boldsymbol{x}\, p(\boldsymbol{x},t) = 1$, where the integral is over the volume of the simplex Δ.

To complete the description of the model, we need expressions for the functions $v_i\big(\boldsymbol{x};\lambda(t)\big)$ and $D_{ij}(\boldsymbol{x})$. Given a Wright-Fisher evolutionary model, these take the following form (see SI for a detailed derivation):

$$v_i(\boldsymbol{x};\lambda(t)) = m_{i\mu}x_\mu + g_{ij}(\boldsymbol{x})s_j(\lambda(t)), \qquad D_{ij}(\boldsymbol{x},t) = \frac{g_{ij}(\boldsymbol{x})}{2N}, \tag{6}$$

where m is the $M\times M$ mutation rate matrix defined in the main text, and $g(\boldsymbol{x})$ is an $(M-1)\times(M-1)$ matrix with elements given by

$$g_{ij}(\boldsymbol{x}) \equiv \begin{cases} -x_i x_j & i \neq j \\ x_i(1-x_i) & i = j, \text{ no sum over } i \end{cases}. \tag{7}$$

4.2 Instantaneous equilibrium distributions

The instantaneous equilibrium (IE) distribution $\rho(\boldsymbol{x};\lambda(t))$ is defined through main text Eq. (2), $\mathcal{L}(\lambda(t))\rho(\boldsymbol{x};\lambda(t)) = 0$. Because we evaluate the effectiveness of our driving by comparing the actual distribution $p(\boldsymbol{x},t)$ to the IE distribution, it is useful to know the form of $\rho(\boldsymbol{x};\lambda(t))$. Unfortunately it is generally not possible to find an IE analytical expression, except in some specific cases[29,53]. The two genotype system ($M=2$) is one example where an exact solution is known. It has a form analogous to the Boltzmann distribution of statistical physics[29,53],

$$\rho(\boldsymbol{x};\lambda(t)) = \frac{e^{-\Phi(\boldsymbol{x};\lambda(t))}}{Z(\lambda(t))}, \tag{8}$$

where $\Phi(\boldsymbol{x};\lambda(t))$ is an effective "potential" given by

$$\Phi(\boldsymbol{x};\lambda(t)) = -2N\left(m_{12}\log x_1 + m_{21}\log(1-x_1) + s_1(\lambda(t))x_1\right) + \log\det g(\boldsymbol{x}) \tag{9}$$

and $Z(\lambda(t))$ is a normalization constant.

To estimate the IE distribution for general M, we take advantage of the large population, frequent mutation regime: $m_{\beta\alpha} \sim \mathcal{O}(\mu)$ for all nonzero matrix entries where $\alpha \neq \beta$, with $\mu N \gg 1$, $N \gg 1$. In this case we know that $\rho(\boldsymbol{x};\lambda(t))$ is approximately a multivariate normal distribution of the form

$$\rho(\boldsymbol{x};\lambda(t)) \approx \left((2\pi)^{M-1}\det\Sigma(\lambda(t))\right)^{-1/2}\exp\left(-\frac{1}{2}\big(x_i-\overline{x}_i(\lambda(t))\big)\Sigma^{-1}_{ij}(\lambda(t))\big(x_j-\overline{x}_j(\lambda(t))\big)\right). \tag{10}$$

Here $\overline{x}_i(\lambda) = \int_\Delta d\boldsymbol{x}\, x_i\rho(\boldsymbol{x};\lambda)$ is the ith mean genotype fraction for the IE distribution, and $\Sigma^{-1}(\lambda)$ is the inverse of the covariance matrix $\Sigma(\lambda)$ for this distribution. The latter has entries $\Sigma_{ij} \equiv \overline{x_i x_j} - \overline{x}_i\overline{x}_j$. In order to make practical use of Eq. (10), we need a method to estimate $\overline{x}_i(\lambda)$ and $\Sigma(\lambda)$. As shown in the SI, this can be done through an approximate numerical solution to a set of exact equations involving the moments of $\rho(\boldsymbol{x};\lambda)$.

4.3 Counterdiabatic driving protocol

To implement CD driving, we need to solve for the CD Fokker-Planck operator $\widetilde{\mathcal{L}}\big(\lambda(t),\dot{\lambda}(t)\big)$ that satisfies main text Eq. (3). We posit that $\widetilde{\mathcal{L}}$ should be in the Fokker-Planck form of Eq. (5), but with some CD version of the selection coefficient, $\tilde{\boldsymbol{s}}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)$, instead of the original $\boldsymbol{s}(\lambda(t))$. The necessary perturbation to the fitness seascape to achieve CD driving, $\delta\tilde{\boldsymbol{s}}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big) \equiv \tilde{\boldsymbol{s}}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big) - \boldsymbol{s}(\lambda(t))$, we take for now to be frequency-dependent for generality. Thus main text Eq. (3) takes the form

$$\begin{aligned}\partial_t\rho\big(\boldsymbol{x};\lambda(t)\big) &= \widetilde{\mathcal{L}}\big(\lambda(t),\dot{\lambda}(t)\big)\rho\big(\boldsymbol{x};\lambda(t)\big)\\ &= -\partial_i\Big(\tilde{v}_i\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)\rho\big(\boldsymbol{x};\lambda(t)\big)\Big) + \partial_i\partial_j\big(D_{ij}(\boldsymbol{x})\rho\big(\boldsymbol{x};\lambda(t)\big)\big)\end{aligned} \tag{11}$$

with a modified velocity function:

$$\begin{aligned}\tilde{v}_i\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big) &= m_{i\mu}x_\mu + g_{ij}(\boldsymbol{x})\tilde{s}_j\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)\\ &= v_i\big(\boldsymbol{x};\lambda(t)\big) + g_{ij}(\boldsymbol{x})\delta\tilde{s}_j\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big).\end{aligned} \tag{12}$$

Using the fact that $\mathcal{L}(\lambda(t))\rho\big(\boldsymbol{x};\lambda(t)\big) = 0$, since $\rho\big(\boldsymbol{x};\lambda(t)\big)$ is the IE distribution of the original operator $\mathcal{L}$, we can rewrite Eq. (11) as

$$\begin{aligned}\partial_t\rho\big(\boldsymbol{x};\lambda(t)\big) &= -\partial_i\Big(\rho\big(\boldsymbol{x};\lambda(t)\big)g_{ij}(\boldsymbol{x})\delta\tilde{s}_j\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)\Big)\\ &\equiv -\partial_i\mathcal{J}_i.\end{aligned} \tag{13}$$

where $\mathcal{J}_i \equiv \rho\big(\boldsymbol{x};\lambda(t)\big)g_{ij}(\boldsymbol{x})\delta\tilde{s}_j\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)$ is a probability current. In this form Eq. (13) looks like a continuity equation, describing the local transport of probability density due to the current field $\mathcal{J}$. In order for this equation to conserve total probability over the simplex Δ, we also require the condition that $\mathcal{J}_i n_i = 0$ at any point on the boundary of the simplex, where the vector $\boldsymbol{n}$ is normal to the boundary at the point. The perturbation $\delta\tilde{\boldsymbol{s}}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)$ that satisfies Eq. (13) and the boundary condition defines an exact CD protocol for the evolutionary system.

Given an arbitrary continuous time sequence of IE distributions $\rho\big(\boldsymbol{x};\lambda(t)\big)$, such a perturbation always exists. In fact, from a formal mathematical standpoint[56], any perturbation of the following form is a solution (note that for clarity we do not use Einstein summation in this case):

$$\delta\tilde{\boldsymbol{s}}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big) = \boldsymbol{g}^{-1}(\boldsymbol{x})\left[\frac{1}{\rho\big(\boldsymbol{x};\lambda(t)\big)}\left(-\sum_{i=1}^{M-1}\hat{\boldsymbol{x}}_i w_i\int_0^{x_i}dx_i'\,\partial_t\rho\big(x_1,\ldots,x_i',\ldots,x_{M-1};\lambda(t)\big) + \boldsymbol{f}\big(\boldsymbol{x};\lambda(t),\dot{\lambda}(t)\big)\right)\right].$$

(14)

Here $\boldsymbol{g}^{-1}(\boldsymbol{x})$ is the inverse of the matrix $\boldsymbol{g}(\boldsymbol{x})$, $\hat{\boldsymbol{x}}_i$ is the unit vector along the ith axis, and the integral in the ith term of the sum is carried out only over the ith genotype fraction, keeping all other components x_j, $j \neq i$, fixed. There are two quantities in Eq. (14) that make the solution potentially non-unique: the weights w_i can be any real numbers, so long as $\sum_{i=1}^{M-1} w_i = 1$; and $\boldsymbol{f}\left(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)\right)$ is an $(M-1)$-dimensional vector function which has zero divergence, $\partial_i f_i = 0$. However we have additional constraints on this function $\boldsymbol{f}$: it has to be compatible with the vanishing of the current orthogonal to boundary, $\mathcal{J}_i n_i = 0$. For $M = 2$, where necessarily $w_1 = 1$, these constraints mean that only $\boldsymbol{f} = 0$ is allowed, and we get a unique exact CD solution. For $M > 2$, the partial differential equation $\partial_i f_i = 0$ and the boundary condition do not specify $\boldsymbol{f}$ uniquely, and hence we get many possible allowable CD solutions all of which satisfy Eq. (13). This in turn means that we can always find CD Fokker-Planck operators $\widetilde{\mathcal{L}}\left(\lambda(t), \dot{\lambda}(t)\right)$ that satisfy main text Eq. (3).

However the formal existence of such perturbations $\delta\tilde{\boldsymbol{s}}$ is not the end of the story, because many of the solutions described by Eq. (14) may not be physically realizable. In order to get at a more practical (though approximate) CD solution, we proceed as follows. As discussed Sec. 4.2, in the regime of interest it is easier to work with moments of the IE distribution, so it is useful to convert Eq. (13) into a relation involving the IE first moment $\overline{x}_i(\lambda)$. Multiply both sides of Eq. (13) by x_k, and notice that $x_k \partial_i \mathcal{J}_i = \partial_i \left(x_k \mathcal{J}_i\right) - \delta_{ik} \mathcal{J}_i$, where δ_{ik} is the Kronecker delta function. Integrating over the entire simplex gives

$$\int_\Delta d\boldsymbol{x}\, x_k \partial_t \rho(\boldsymbol{x}; \lambda(t)) = -\int_\Delta d\boldsymbol{x}\, \partial_i \left(x_k \mathcal{J}_i\right) + \int_\Delta d\boldsymbol{x}\, \mathcal{J}_k. \tag{15}$$

By Gauss's theorem, $\int_\Delta d\boldsymbol{x}\, \partial_i \left(x_k \mathcal{J}_i\right) = \int_{\partial\Delta} d\sigma\, x_k \mathcal{J}_i n_i$, where the integral involves area elements $d\sigma$ of the simplex boundary $\partial\Delta$, and n_i are the components of the normal vector to this boundary. By conservation of probability, the component of $\mathcal{J}$ normal to $\partial\Delta$ vanishes, i.e. $\mathcal{J}_i n_i = 0$, so the first term in Eq. (15) is zero. Plugging the definition of $\mathcal{J}_k$ into the second term, we get

$$\int_\Delta d\boldsymbol{x}\, x_k \partial_t \rho(\boldsymbol{x}; \lambda(t)) = \int_\Delta d\boldsymbol{x}\, \rho(\boldsymbol{x}; \lambda(t)) g_{kj}(\boldsymbol{x}) \delta\tilde{s}_j(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)), \tag{16}$$

or equivalently

$$\partial_t \overline{\boldsymbol{x}}(\lambda(t)) = \left\langle \boldsymbol{g}(\boldsymbol{x}) \delta\tilde{\boldsymbol{s}}\left(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)\right) \right\rangle, \tag{17}$$

where the brackets $\langle\,\rangle$ denote an average over the simplex with respect to $\rho(\boldsymbol{x}; \lambda(t))$.

So far both Eq. (13) and (17) are exact relations satisfied by the CD perturbation $\delta\tilde{\boldsymbol{s}}$. However we can simplify the results in the large population, frequent mutation regime, where $\rho(\boldsymbol{x}; \lambda(t))$ has the approximate normal form of Eq. (10). As argued in the SI, in this case the leading contribution to $\delta\tilde{\boldsymbol{s}}$ is frequency-independent, $\delta\tilde{\boldsymbol{s}}(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)) \approx \delta\tilde{\boldsymbol{s}}(\lambda(t), \dot{\lambda}(t))$, with corrections that vanish in the large N limit. The leading contribution $\delta\tilde{\boldsymbol{s}}(\lambda(t), \dot{\lambda}(t))$ satisfies a version of Eq. (17) with $\boldsymbol{x}$ on the right-hand side replaced by the IE mean $\overline{\boldsymbol{x}}(\lambda(t))$,

$$\partial_t \overline{\boldsymbol{x}}(\lambda(t)) = \boldsymbol{g}\left(\overline{\boldsymbol{x}}(\lambda(t))\right) \delta\tilde{\boldsymbol{s}}\left(\lambda(t), \dot{\lambda}(t)\right). \tag{18}$$

This equation can be directly solved for $\delta\tilde{\boldsymbol{s}}(\lambda(t), \dot{\lambda}(t))$ in terms of $\overline{\boldsymbol{x}}(\lambda(t))$, yielding the approximate CD solution of main text Eq. (4). Thus knowing the IE first moment $\overline{\boldsymbol{x}}(\lambda(t))$ over the duration of the protocol (via the numerical procedure described in the SI) allows us to estimate a CD driving prescription.

4.4 CD driving for the two genotype example

For the $M = 2$ system, the exact IE distribution is given by Eqs. (8)-(9). In the large population, frequent mutation limit we can estimate the mean frequency $\overline{x}_1(\lambda(t))$ corresponding to this distribution as:

$$\overline{x}_1(\lambda(t)) \approx \frac{-m_{12} - m_{21} + s_1(\lambda(t)) + \sqrt{\left(m_{12} + m_{21} - s_1(\lambda(t))\right)^2 + 4 m_{12} s_1(\lambda(t))}}{2 s_1(\lambda(t))}. \tag{19}$$

This allows the CD prescription in Eq. (4) to be evaluated analytically, yielding

$$\tilde{s}_1\big(\lambda(t),\dot{\lambda}(t)\big) = s_1\big(\lambda(t)\big) + \frac{\partial_t s_1\big(\lambda(t)\big)}{\sqrt{\big(m_{12}+m_{21}-s_1\big(\lambda(t)\big)\big)^2 + 4m_{12}s_1\big(\lambda(t)\big)}}. \tag{20}$$

For the results in Fig. 2, we assume the following ramp for the selection coefficient: $s_1\big(\lambda(t)\big) = \sigma/(1+ae^{-kt}) - \sigma/(1+a)$, with $\sigma = 0.02$, $a = 817$, $k = 0.06$. The other model parameters are set to: $N = 10^4$, $m_{12} = m_{21} = 2.5\times10^{-3}$.

4.5 Agent based model

4.5.1 Model description

For the agent based model (ABM) simulations, we track a population of single-celled organisms that undergo birth (through binary division), death, and mutations. There are M genotypes, and the fitness of genotype $i < M$ relative to the Mth one (the wildtype) is $1+s_i(\lambda(\tau))$, which depends on the drug dosage $\lambda(\tau)$ at the current simulation time step τ. (The mapping between simulation time steps τ and Wright-Fisher generations t will be discussed below.) At each simulation time step, every cell in the population undergoes the following process: i) with probability d it dies; ii) if it survives, the cell divides with a genotype-dependent probability

$$b_i(\tau) = \begin{cases} \min\left(b_0(1+s_i(\lambda(\tau))\left(1-\frac{N_{\text{cell}}(\tau)}{K}\right),1\right) & N_{\text{cell}}(\tau) \le K \\ 0 & N_{\text{cell}}(\tau) > K, \end{cases} \tag{21}$$

where i is the cell's genotype, b_0 is a baseline birth rate, $N_{\text{cell}}(\tau)$ is the current number of cells in the population, and K is the carrying capacity. Upon division, the daughter cell mutates to another genotype j with probability $\hat{m}_{ji}$, $j \neq i$.

4.5.2 In silico implementation

The ABM was implemented for the $M = 2$ and $M = 16$ examples described in the main text using code written in the C++ programming language. Code, configuration files, and analysis scripts for these models can be found on https://github.com/Peyara/Evolution-Counterdiabatic-Driving. The code directly implements the model of the previous section, and is summarized in the flowchart of Fig. 4. The $M = 2$ selection coefficient $s_1(\lambda(t))$ and other model parameters are as described in the main text.

For $M = 16$, the simulations were run for 4.5×10^4 time steps with a death rate $d = 0.05$, a baseline birth rate $b_0 = 2$, and a carrying capacity $K = 5\times10^6$. The mutation probability $\hat{m}_{ji}$, $j \neq i$ is zero unless the Hamming distance between the binary string representation of i and j is 1. This gives the "tesseract" connectivity seen in Fig. 3A,D. Where nonzero, the probability $\hat{m}_{ji} = 2.5\times10^{-4}$, giving a total mutation probability $\sum_{j\neq i}\hat{m}_{ji} = 10^{-3}$ for all offspring. To give the population time to reach an initial equilibrium, the drug concentration $\lambda(\tau)$ is initially small, increases substantially around time step $\tau \sim 10^4$, and then plateaus at later times. The dosage follows the equation,

$$\lambda(\tau) = \frac{a}{1+\exp\big(-b(\tau-c)\big)}, \tag{22}$$

with parameters: $a = 1.5\times10^{-4}$ M, $b = 2\times10^{-3}$, $c = 10{,}110$. The selection coefficients $s_i(\lambda(\tau))$ were varied with concentration $\lambda(\tau)$ in accordance with the experimentally measured dose-fitness curves of 16 genotypes for the anti-malarial drug pyrimethamine[2,3]. To calculate distributions of genotype frequencies, every simulation is repeated 1000 times.

4.5.3 Mapping the ABM simulations to a Fokker-Planck equation

In order to implement the CD driving protocol, derived for Wright-Fisher Fokker-Planck dynamics, in the context of the ABM, we need a mapping between the ABM parameters and the corresponding Fokker-Planck parameters. As shown in the SI, this can be done by describing the ABM simulation population updating at each time step as an effective Langevin equation, and then using the connection between the Langevin and Fokker-Planck descriptions[55,57]. The resulting approximate correspondence is summarized as follows: i) a duration of τ ABM simulation time steps maps to $t \approx \tau d$ Wright-Fisher generations, where d is the ABM death rate. ii) The Fokker-Planck mutation matrix entries $m_{i\nu}$, $i \neq \nu$, are given by $m_{i\nu} \approx \hat{m}_{i\nu}(1+s_\nu)$, where $\hat{m}_{i\nu}$ are the ABM mutation probabilities. iii) The effective

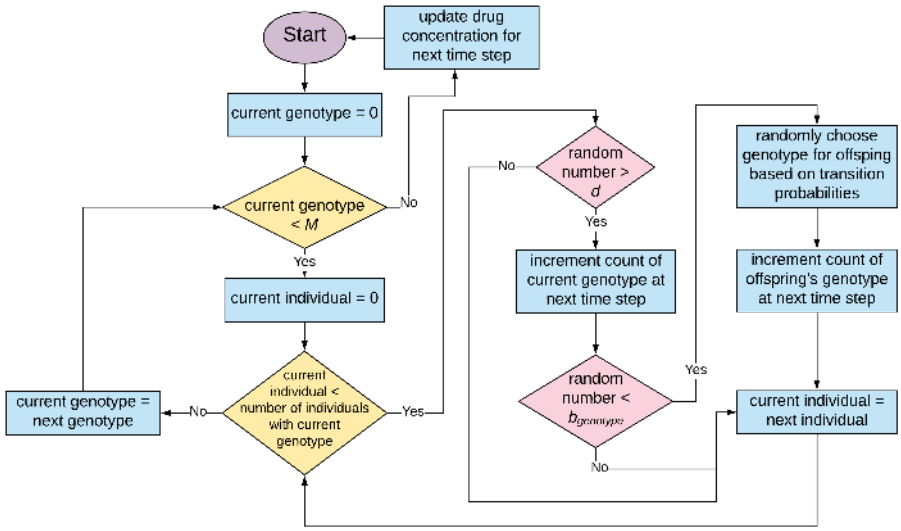

Figure 4. Flow chart describing the code executed at each time step by ABMs in this paper. M refers to the total number of genotypes in the model, d to the death rate, and $b_{genotype}$ to the birth rate of the current genotype. Random numbers in the chart are drawn uniformly from the range 0 to 1.

population N in the Fokker-Planck model is given by $N \approx \frac{1}{2}K(1-db_0^{-1}(1-d)^{-1})$, where K is the ABM carrying capacity and b_0 the baseline birth rate. The accuracy of this mapping is illustrated in Fig. 2A,C, where the distributions from ABM simulations for $M=2$ (red circles) are compared against numerical Fokker-Planck solutions with parameters calculated using the mapping (red curves).

4.5.4 Numerical estimation of the KL divergence and reduction in lag time

To quantify the effectiveness of the CD driving, we use the KL divergence between the actual distribution, $p(\boldsymbol{x},t)$ and the IE one, $\rho(\boldsymbol{x};\lambda(t))$, defined as $D_{\mathrm{KL}}(\rho||p)=\int d\boldsymbol{x}\,\rho(\boldsymbol{x};\lambda(t))\log_2\left(\rho(\boldsymbol{x};\lambda(t))/p(\boldsymbol{x},t)\right)$. For $M=2$, the Fokker-Planck equation can be solved numerically for $p(x_1,t)$, while $\rho(x_1;\lambda(t))$ is known analytically (Eqs. (8)-(9)). Hence the one-dimensional integral for $D_{\mathrm{KL}}(\rho||p)$ can be numerically evaluated. For $M=16$ the situation becomes more complicated. There is no analytical solution for $\rho(\boldsymbol{x};\lambda(t))$, but we do have a good approximation in terms of the multivariate normal distribution of Eq. (10), expressed in terms of the mean vector $\overline{\boldsymbol{x}}(\lambda(t))$ and covariance matrix $\Sigma(\lambda(t))$ that are calculated using the moment approach described in the SI. The ABM simulation results are also normally distributed in this parameter regime, and hence there is a corresponding simulation mean $\overline{\boldsymbol{x}}_{\mathrm{sim}}(t)$ and covariance $\Sigma_{\mathrm{sim}}(t)$ that can be calculated at each time t. These are calculated from the ensemble of 1000 simulations that are run for each parameter set. The integral for the KL divergence $D_{\mathrm{KL}}(\rho||p)$ between the simulation and IE multivariate normal distributions can then be evaluated directly, yielding

$$D_{\mathrm{KL}}(\rho||p)=\frac{1}{2\ln 2}\left[\ln\frac{\det\Sigma_{\mathrm{sim}}(t)}{\det\Sigma(\lambda(t))}-M+1+\mathrm{tr}\left(\Sigma_{\mathrm{sim}}^{-1}(t)\Sigma(\lambda(t))\right)+\left(\overline{\boldsymbol{x}}_{\mathrm{sim}}(t)-\overline{\boldsymbol{x}}(\lambda(t))\right)^T\Sigma_{\mathrm{sim}}^{-1}(t)\left(\overline{\boldsymbol{x}}_{\mathrm{sim}}(t)-\overline{\boldsymbol{x}}(\lambda(t))\right)\right] \tag{23}$$

Since $\Sigma_{\mathrm{sim}}(t)$ will have some degree of sampling errors due to the finite size of the simulation ensemble, it can in some cases be badly conditioned. In these scenarios the Moore-Penrose pseudo-inverse is used to estimate $\Sigma_{\mathrm{sim}}^{-1}(t)$.

We can use the curves of $D_{\mathrm{KL}}(\rho||p)$ as a function of time, for example those of Fig. 3F, to estimate how much lag time (Δt) is being eliminated using a given approximate CD protocol, relative to the original one. This lag time savings $\Delta t = t_{\mathrm{eq}}^{\mathrm{orig}} - t_{\mathrm{eq}}^{\mathrm{CD}}$, where $t_{\mathrm{eq}}^{\mathrm{orig}}$ and $t_{\mathrm{eq}}^{\mathrm{CD}}$ are respectively the times at which probability distributions in the original and CD protocols reach their final IE target values. In terms of $D_{\mathrm{KL}}(\rho||p)$, there is minimum value $D_{\mathrm{KL}}^{\mathrm{eq}}$ attained at long times when $p(\boldsymbol{x},t)$ has converged with $\rho(\boldsymbol{x};\lambda(t))$. Note this value is not precisely zero because of numerical noise associated with the estimation of the distribution $p(\boldsymbol{x},t)$ from a finite number of simulations. At long times when $D_{\mathrm{KL}}(\rho||p)$ approaches $D_{\mathrm{KL}}^{\mathrm{eq}}$, the final approach can be fit well by the following exponential decay

function,

$$D_{\mathrm{KL}}(\rho||p) \approx \begin{cases} D_{\mathrm{KL}}^{\mathrm{eq}} e^{-(t_{\mathrm{eq}}-t)/\tau} & t \leq t_{\mathrm{eq}} \\ D_{\mathrm{KL}}^{\mathrm{eq}} & t > t_{\mathrm{eq}} \end{cases} . \tag{24}$$

Since we know $D_{\mathrm{KL}}^{\mathrm{eq}}$ from the long-time behavior of the KL divergence curves, we then can estimate τ and t_{eq} by fitting Eq. (24) to the final decay portion of each $D_{\mathrm{KL}}(\rho||p)$ curve (the time range where $D_{\mathrm{KL}}(\rho||p)$ is within two orders of magnitude of $D_{\mathrm{KL}}^{\mathrm{eq}}$). After finding t_{eq} for the original and CD protocols, the difference gives us the Δt values quoted in the main text and SI.

References

1. Mira, P. M. *et al.* Rational design of antibiotic treatment plans: a treatment strategy for managing evolution and reversing resistance. *PloS One* **10**, e0122283 (2015).
2. Ogbunugafor, C. B., Wylie, C. S., Diakite, I., Weinreich, D. M. & Hartl, D. L. Adaptive landscape by environment interactions dictate evolutionary dynamics in models of drug resistance. *PLoS Comp. Biol.* **12**, e1004710 (2016).
3. Brown, K. M. *et al.* Compensatory mutations restore fitness during the evolution of dihydrofolate reductase. *Mol. Biol. Evol.* **27**, 2682–2690 (2010).
4. World Health Organization. Antimicrobial resistance: global report on surveillance. (2014). ISBN: 978 92 4 156474 8.
5. Holohan, C., Van Schaeybroeck, S., Longley, D. B. & Johnston, P. G. Cancer drug resistance: an evolving paradigm. *Nat. Rev. Cancer* **13**, 714–726 (2013).
6. World Health Organization. HIV drug resistance report 2019 (2019). Licence: CC BY-NC-SA 3.0 IGO.
7. Nichol, D. *et al.* Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS Comp. Biol.* **11**, e1004493 (2015).
8. Maltas, J. & Wood, K. B. Pervasive and diverse collateral sensitivity profiles inform optimal strategies to limit antibiotic resistance. *bioRxiv* 241075 (2018).
9. Bason, M. G. *et al.* High-fidelity quantum driving. *Nat. Phys.* **8**, 147 (2012).
10. Zhou, B. B. *et al.* Accelerated quantum control using superadiabatic dynamics in a solid-state lambda system. *Nat. Phys.* **13**, 330 (2017).
11. Walther, A. *et al.* Controlling fast transport of cold trapped ions. *Phys. Rev. Lett.* **109**, 080501 (2012).
12. Farhi, E. *et al.* A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem. *Science* **292**, 472–475 (2001).
13. Torrontegui, E. *et al.* Shortcuts to adiabaticity. In *Adv. At. Mol. Opt. Phys.*, vol. 62, 117–169 (Elsevier, 2013).
14. Deffner, S., Jarzynski, C. & del Campo, A. Classical and quantum shortcuts to adiabaticity for scale-invariant driving. *Phys. Rev. X* **4**, 021013 (2014).
15. Deffner, S. Shortcuts to adiabaticity: suppression of pair production in driven dirac dynamics. *New J. Phys.* **18**, 012001 (2015).
16. Acconcia, T. V., Bonança, M. V. S. & Deffner, S. Shortcuts to adiabaticity from linear response theory. *Phys. Rev. E* **92**, 042148 (2015).
17. Campbell, S. & Deffner, S. Trade-off between speed and cost in shortcuts to adiabaticity. *Phys. Rev. Lett.* **118**, 100601 (2017).
18. Guéry-Odelin, D. *et al.* Shortcuts to adiabaticity: Concepts, methods, and applications. *Rev. Mod. Phys.* **91**, 045001 (2019).
19. Demirplak, M. & Rice, S. A. Adiabatic population transfer with control fields. *J. Phys. Chem. A* **107**, 9937–9945 (2003).
20. Demirplak, M. & Rice, S. A. Assisted adiabatic passage revisited. *J. Phys. Chem. B* **109**, 6838–6844 (2005).
21. Berry, M. V. Transitionless quantum driving. *J. Phys. A: Math. Theor.* **42**, 365303 (2009).
22. Patra, A. & Jarzynski, C. Shortcuts to adiabaticity using flow fields. *New J. Phys.* **19**, 125009 (2017).
23. Li, G., Quan, H. & Tu, Z. Shortcuts to isothermality and nonequilibrium work relations. *Phys. Rev. E* **96**, 012144 (2017).
24. Martínez, I. A., Petrosyan, A., Guéry-Odelin, D., Trizac, E. & Ciliberto, S. Engineered swift equilibration of a brownian particle. *Nat. Phys.* **12**, 843 (2016).
25. Le Cunuder, A. *et al.* Fast equilibrium switch of a micro mechanical oscillator. *Appl. Phys. Lett.* **109**, 113502 (2016).
26. Schmiedl, T. & Seifert, U. Optimal finite-time processes in stochastic thermodynamics. *Phys. Rev. Lett.* **98**, 108301 (2007).

27. Aurell, E., Gawędzki, K., Mejía-Monasterio, C., Mohayaee, R. & Muratore-Ginanneschi, P. Refined second law of thermodynamics for fast random processes. *J. Stat. Phys.* **147**, 487–505 (2012).

28. Wright, S. *The roles of mutation, inbreeding, crossbreeding, and selection in evolution*, vol. 1 (na, 1932).

29. Mustonen, V. & Lässig, M. Fitness flux and ubiquity of adaptive evolution. *Proc. Natl. Acad. Sci.* **107**, 4248–4253 (2010).

30. Grabert, H., Hänggi, P. & Talkner, P. Is quantum mechanics equivalent to a classical stochastic process? *Phys. Rev. A* **19**, 2440–2445 (1979).

31. Van Kampen, N. G. *Stochastic processes in physics and chemistry* (Elsevier, 1992).

32. Risken, H. *The Fokker-Planck Equation* (Springer, 1996).

33. Born, M. & Fock, V. Beweis des adiabatensatzes. *Z. Phys.* **51**, 165–180 (1928).

34. Nichol, D. *et al.* Antibiotic collateral sensitivity is contingent on the repeatability of evolution. *Nat. Commun.* **10**, 334 (2019).

35. Li, Y., Petrov, D. A. & Sherlock, G. Single nucleotide mapping of trait space reveals pareto fronts that constrain adaptation. *Nat. Ecol. Evol.* 1–13 (2019).

36. Vaikuntanathan, S. & Jarzynski, C. Dissipation and lag in irreversible processes. *EPL (Europhysics Lett.* **87**, 60005 (2009).

37. Gillespie, J. H. A simple stochastic gene substitution model. *Theor. Popul. Biol.* **23**, 202–215 (1983).

38. Gerrish, P. J. & Lenski, R. E. The fate of competing beneficial mutations in an asexual population. *Genetica* **102**, 127 (1998).

39. Desai, M. M. & Fisher, D. S. Beneficial mutation–selection balance and the effect of linkage on positive selection. *Genetics* **176**, 1759–1798 (2007).

40. Sniegowski, P. D. & Gerrish, P. J. Beneficial mutations and the dynamics of adaptation in asexual populations. *Philos. Transactions Royal Soc. B: Biol. Sci.* **365**, 1255–1263 (2010).

41. Martens, E. A. & Hallatschek, O. Interfering waves of adaptation promote spatial mixing. *Genetics* **189**, 1045–1060 (2011).

42. Magdanova, L. & Golyasnaya, N. Heterogeneity as an adaptive trait of microbial populations. *Microbiology* **82**, 1–10 (2013).

43. Krishnan, N. & Scott, J. G. Range expansion shifts clonal interference patterns in evolving populations. *bioRxiv* 794867 (2019).

44. Sella, G. & Hirsh, A. E. The application of statistical physics to evolutionary biology. *Proc. Natl. Acad. Sci.* **102**, 9541–9546 (2005).

45. Chiel, J. *et al.* (2020). In preparation.

46. Kullback, S. & Leibler, R. A. On Information and Sufficiency. *Ann. Math. Stat.* **22**, 79–86 (1951).

47. Kaznatcheev, A. Computational complexity as an ultimate constraint on evolution. *Genetics* **212**, 245–265 (2019).

48. Dobzhansky, T. Nothing in biology makes sense except in the light of evolution. *Am. Biol. Teach.* **35**, 125–129 (1973).

49. Romero, P. A. & Arnold, F. H. Exploring protein fitness landscapes by directed evolution. *Nat. Rev. Mol. Cell Biol* **10**, 866 (2009).

50. Tuerk, C. & Gold, L. Systematic evolution of ligands by exponential enrichment: Rna ligands to bacteriophage t4 dna polymerase. *Science* **249**, 505–510 (1990).

51. Ellington, A. D. & Szostak, J. W. In vitro selection of rna molecules that bind specific ligands. *Nature* **346**, 818 (1990).

52. Bita, C. & Gerats, T. Plant tolerance to high temperature in a changing environment: scientific fundamentals and production of heat stress-tolerant crops. *Front. Plant Sci.* **4**, 273 (2013).

53. Baxter, G. J., Blythe, R. A. & McKane, A. J. Exact solution of the multi-allelic diffusion model. *Math. Biosci.* **209**, 124–170 (2007).

54. Kimura, M. Stochastic processes and distribution of gene frequencies under natural selection. *Cold Spring Harb. Symp. Quant. Biol.* **20**, 33–53 (1955).

55. Gillespie, D. T. The multivariate Langevin and Fokker-Planck equations. *Am. J. Phys.* **64**, 1246–1257 (1996).

56. Sahoo, S. Inverse vector operators. *arXiv* 0804.2239 (2008).

57. Gillespie, D. T. The chemical Langevin equation. *J. Chem. Phys.* **113**, 297–306 (2000).

Acknowledgements

MH would like to thank the U.S. National Science Foundation for support through the CAREER grant (BIO/MCB 1651560). JGS would like to thank the NIH Loan Repayment Program for their generous support and the Paul Calabresi Career Development Award for Clinical Oncology (NIH K12CA076917). SD acknowledges support from the U.S. National Science Foundation under Grant No. CHE-1648973.

Author contributions statement

SI and JP performed mathematical analysis, wrote the two-allele code, peformed simulations, analyzed the data and wrote the manuscript. JC, EI, OG, BK performed mathematical analysis, analyzed the data and wrote the manuscript. ED and NK wrote the multidimensional ABM code, performed the simulations, analyzed data and wrote the manuscript. JGS analyzed the data and wrote the manuscript. MH performed the mathematical analysis and simulations, wrote code, analyzed the data and wrote the manuscript. SD wrote the manuscript, and SD, EI, JGS, MH contributed to developing the overall theoretical framework. These contributions are graphically illustrated in Figure 5.

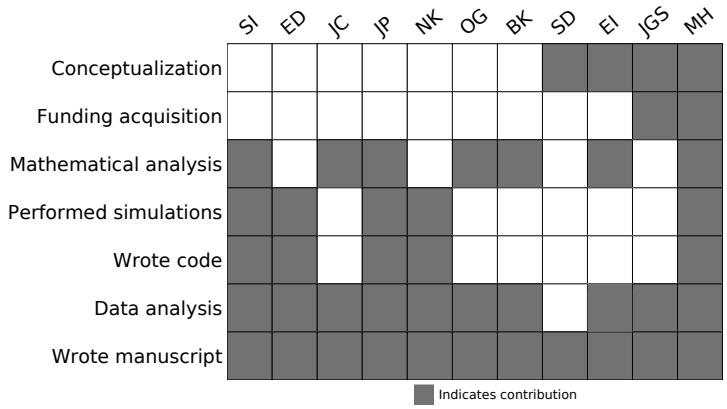

Figure 5. Author Contributions

Data availability

The raw numerical data for the figures in the main text and SI, as well as the code to generate the figures, is available via github at https://github.com/Peyara/Evolution-Counterdiabatic-Driving.

Code availability

The code to perform the numerical simulations and the specific driving protocols is available via github at https://github.com/Peyara/Evolution-Counterdiabatic-Driving.

Supplementary Information: Controlling the speed and trajectory of evolution with counterdiabatic driving

Contents

1 Fokker-Planck analogue of quantum adiabatic theorem

The Fokker-Planck dynamics of our model obeys a classical, stochastic analogue of the quantum adiabatic theorem[1,2]. As described in the main text, this means that if the system starts at $t = t_0$ in the equilibrium distribution $\rho(\boldsymbol{x};\lambda(0))$ corresponding to control parameter $\lambda(0)$, it will remain in the corresponding instantaneous equilibrium (IE) distribution $\rho(\boldsymbol{x};\lambda(t))$ at all $t > t_0$ if $\lambda(t)$ is varied infinitesimally slowly. Note that for the quantum case, if a system starts in *any* eigenstate of the Hamiltonian, it will remain in that eigenstate if the Hamiltonian is varied adiabatically, assuming the eigenvalues never become degenerate. The classical version that we demonstrate here only applies to one eigenstate, the IE distribution, which corresponds to the quantum ground state.

1.1 Fokker-Planck eigenfunction expansion, adjoint operator

Before considering adiabatic driving, we start with some preliminaries for a Fokker-Planck (FP) system with a time-independent control parameter $\lambda(t) = \lambda$. The FP operator for a given λ is defined in Eq. (5) of the Methods,

$$\mathcal{L}(\lambda)p(\boldsymbol{x},t) \equiv -\partial_i\left(v_i(\boldsymbol{x};\lambda)p(\boldsymbol{x},t)\right) + \partial_i\partial_j\left(D_{ij}(\boldsymbol{x})p(\boldsymbol{x},t)\right). \tag{S1}$$

Throughout the Supplementary Information (SI) we will use the same Einstein summation notation that we described in the Methods, with repeated Roman indices summed from 1 to $M-1$ and repeated Greek indices summed from 1 to M. The two exceptions for clarity will be: i) the eigenfunction indices n and m used in this section, where the summation convention will not apply and explicit sums will be always be indicated; ii) the final section describing the mapping between the agent-based model and the Fokker-Planck equation, where it will be more convenient to write out all sums explicitly.

The operator $\mathcal{L}(\lambda)$ has an associated set of eigenfunctions $\psi_n(\boldsymbol{x};\lambda)$ and eigenvalues $-\kappa_n(\lambda)$, satisfying[3]

$$\mathcal{L}(\lambda)\psi_n(\boldsymbol{x};\lambda) = -\kappa_n(\lambda)\psi_n(\boldsymbol{x};\lambda) \tag{S2}$$

for $n = 0, 1, 2, \ldots$. We assume an equilibrium distribution exists for every value of the control parameter λ, in which case we know one eigenvalue is zero. By convention we choose this to be $n = 0$, so $\kappa_0(\lambda) = 0$, and the corresponding eigenfunction $\psi_0(\boldsymbol{x};\lambda) \equiv \rho(\boldsymbol{x};\lambda)$. Eigenvalues for $n > 0$ can be in general complex, but have positive real parts, $\mathrm{Re}(\kappa_n(\lambda)) > 0$, which guarantees that the system eventually equilibrates, as discussed below[3,4].

The Fokker-Planck equation for constant λ,

$$\partial_t p(\boldsymbol{x},t) = \mathcal{L}(\lambda) p(\boldsymbol{x},t) \tag{S3}$$

has a solution that can be expressed as a linear combination of the eigenfunctions,

$$p(\boldsymbol{x},t) = \sum_{n=0}^{\infty} c_n \psi_n(\boldsymbol{x};\lambda) e^{-\kappa_n(\lambda)t}, \tag{S4}$$

for some constants c_n where $c_0 = 1$. The fact that $\mathrm{Re}(\kappa_n(\lambda)) > 0$ for $n > 0$ ensures eventual equilibration: $p(\boldsymbol{x},t) \to \psi_0(\boldsymbol{x};\lambda) = \rho(\boldsymbol{x};\lambda)$ as $t \to \infty$.

It is convenient to introduce the adjoint $\mathcal{L}^\dagger(\lambda)$ of the FP operator[3,4],

$$\mathcal{L}^\dagger(\lambda) p(\boldsymbol{x},t) \equiv v_i(\boldsymbol{x};\lambda)\partial_i p(\boldsymbol{x},t) + D_{ij}(\boldsymbol{x})\partial_i\partial_j p(\boldsymbol{x},t), \tag{S5}$$

with corresponding eigenfunctions $\xi_n(\boldsymbol{x};\lambda)$,

$$\mathcal{L}^\dagger(\lambda)\xi_n(\boldsymbol{x};\lambda) = -\kappa_n^*(\lambda)\xi_n(\boldsymbol{x};\lambda). \tag{S6}$$

where the asterisk denotes complex conjugation. Let us define the scalar product of two functions $f(\boldsymbol{x};\lambda)$ and $h(\boldsymbol{x};\lambda)$ through

$$\langle f, h\rangle_\lambda \equiv \int_\Delta d\boldsymbol{x}\, f(\boldsymbol{x};\lambda) h^*(\boldsymbol{x};\lambda) \tag{S7}$$

where the integral is over the $M-1$ dimensional simplex Δ defined in the Methods, and the λ subscript denotes the dependence on λ. Then $\mathcal{L}^\dagger(\lambda)$ has the conventional property of an adjoint:

$$\langle \mathcal{L} f, h\rangle_\lambda = \langle f, \mathcal{L}^\dagger h\rangle_\lambda. \tag{S8}$$

A consequence of this property, using $f = \psi_n$ and $h = \xi_m$, is that the eigenfunctions can be chosen to ensure biorthonormality of the form

$$\langle \psi_n, \xi_m\rangle_\lambda = \delta_{nm}, \tag{S9}$$

where δ_{nm} is the Kronecker delta. By inspection of Eq. (S5) one can see that the $n = 0$ adjoint eigenfunction, with eigenvalue $\kappa_0^*(\lambda) = 0$, is $\xi_0(\boldsymbol{x};\lambda) = 1$. The biorthonormality relation in this case,

$$1 = \langle \psi_0, \xi_0\rangle_\lambda = \int_\Delta d\boldsymbol{x}\, \rho(\boldsymbol{x};\lambda) \tag{S10}$$

corresponds to the normalization of the equilibrium distribution $\rho(\boldsymbol{x};\lambda)$. We also know that for $n > 0$,

$$0 = \langle \psi_n, \xi_0\rangle_\lambda = \int_\Delta d\boldsymbol{x}\, \psi_n(\boldsymbol{x};\lambda). \tag{S11}$$

This property ensures that $p(\boldsymbol{x},t)$ from Eq. (S4), with $c_0 = 1$, is also properly normalized, $\int_\Delta d\boldsymbol{x}\, p(\boldsymbol{x},t) = 1$.

1.2 Fokker-Planck adiabatic driving

We now allow the control parameter $\lambda(t)$ to vary with time. At any given time t, the definitions of the previous section generalize to give instantaneous operators $\mathcal{L}(\lambda(t))$, $\mathcal{L}^\dagger(\lambda(t))$, and corresponding instantaneous eigenfunctions/eigenvalues $\psi_n(\boldsymbol{x};\lambda(t))$, $\xi_n(\boldsymbol{x};\lambda(t))$ and $\kappa_n(\lambda(t))$. The dynamics of the system are now described by the FP equation

$$\partial_t p(\boldsymbol{x},t) = \mathcal{L}(\lambda(t)) p(\boldsymbol{x},t). \tag{S12}$$

Working by analogy with the standard proof of the quantum adiabatic theorem[2], let us posit a solution to this equation of the form

$$p(\boldsymbol{x},t) = \sum_{n=0}^{\infty} c_n(t)\psi_n(\boldsymbol{x};\lambda(t)) e^{-\Theta_n(t)}, \tag{S13}$$

where $c_n(t)$ are some functions to be determined and $\Theta_n(t) \equiv \int_{t_0}^{t} dt' \, \kappa_n(\lambda(t'))$. Plugging Eq. (S13) into Eq. (S12), and using the fact that $\mathcal{L}(\lambda(t))\psi_n(\boldsymbol{x};\lambda(t)) = -\kappa_n(\lambda(t))\psi_n(\boldsymbol{x};\lambda(t))$, we see the $p(\boldsymbol{x},t)$ form satisfies the FP equation assuming the following relation is true:

$$\sum_{n=0}^{\infty} \dot{c}_n(t)\psi_n(\boldsymbol{x};\lambda(t))e^{-\Theta_n(t)} = -\sum_{n=0}^{\infty} c_n(t)\dot{\lambda}(t)\partial_\lambda \psi_n(\boldsymbol{x};\lambda(t))e^{-\Theta_n(t)}, \tag{S14}$$

where $\dot{f}(t) \equiv df/dt$ for a function $f(t)$. Taking the scalar product of both sides with respect to $\xi_m(\boldsymbol{x};\lambda(t))$, and using the biorthonormality relations we find a set of coupled differential equations for $m = 0, 1, 2, \ldots$ that can in principle be used to solve for the functions $c_n(t)$ if we knew the eigenfunctions/eigenvalues:

$$\dot{c}_m(t) = -\sum_{n=0}^{\infty} c_n(t)\dot{\lambda}(t)\langle \partial_\lambda \psi_n, \xi_m \rangle_{\lambda(t)} e^{-(\Theta_n(t)-\Theta_m(t))}. \tag{S15}$$

For $m = 0$ the relations in Eqs. (S10)-(S11) allow us to simplify the above to $\dot{c}_0(t) = 0$, which means $c_0(t)$ is time-independent (and has to be equal to 1 to ensure normalization). Thus we can write Eq. (S13) as

$$p(\boldsymbol{x},t) = \rho(\boldsymbol{x};\lambda(t)) + \sum_{n=1}^{\infty} c_n(t)\psi_n(\boldsymbol{x};\lambda(t))e^{-\Theta_n(t)}. \tag{S16}$$

Let us imagine that we start at t_0 in equilibrium, so $c_n(0) = 0$ for $n > 0$. For a general control protocol $\lambda(t)$, Eq. (S15) implies that $c_n(t)$, $n > 0$, would not necessarily stay zero at later times $t > t_0$. Hence $p(\boldsymbol{x},t)$ in Eq. (S16) would gain contributions from higher eigenfunctions in the second term, and no longer remain in the IE distribution. This is the classical analogue of the observation that for a general time-dependent Hamiltonian driven at a finite rate, a quantum system that started in a ground state will evolve into a superposition of instantaneous ground and excited states.

However if the driving was infinitesimally slow, $\dot{\lambda}(t) \to 0$, then the right-hand side of Eq. (S15) becomes negligible, and hence $c_n(t) \approx 0$ for $n > 0$. Thus for adiabatically slow driving, $p(x,t) \approx \rho(\boldsymbol{x};\lambda(t))$ at all times $t > t_0$. The system remains in the IE distribution, just like the corresponding quantum system remains in the instantaneous ground state.

2 Derivation of v_i and D_{ij} for Fokker-Planck Wright-Fisher model with mutation and selection

To derive the expressions for v_i and D_{ij} in Eqs. (6)-(7) of the Methods, we start with haploid WF evolutionary dynamics defined as follows: each time step corresponds to a generation, and every new generation is created by each child randomly "choosing" a parent in the previous generation and copying the parental genotype. Let N be the total population, assumed fixed between generations (i.e. at carrying capacity). Let $\phi_i(\boldsymbol{x};\lambda(t))$ be the probability of choosing a parent of genotype i, which may depend on the control parameter $\lambda(t)$ through its influence on selection coefficients. In the absence of mutation and selection, $\phi_i(\boldsymbol{x};\lambda(t)) = x_i$, the fraction of that genotype in the parental generation. However we will keep ϕ_i general, in order to incorporate mutation / selection effects later on. The probability of the new generation having a set of genotype populations $\boldsymbol{n} \equiv (n_1, n_2, \ldots, n_{M-1})$, where n_i is the number of type i individuals, is given by the multinomial distribution,

$$\mathcal{P}(\boldsymbol{n};\boldsymbol{x},\lambda(t)) = \frac{N!}{n_1!n_2!\cdots n_{M-1}!n_M!}\phi_1^{n_1}(\boldsymbol{x};\lambda(t))\phi_2^{n_2}(\boldsymbol{x};\lambda(t))\cdots\phi_{M-1}^{n_{M-1}}(\boldsymbol{x};\lambda(t))\phi_M^{n_M}(\boldsymbol{x};\lambda(t)) \tag{S17}$$

where $n_M = N - \sum_{j=1}^{M-1} n_j$ is the number of type M individuals. Note that Eq. (S17) is defined for all allowable configurations of types, or in other words for any $\boldsymbol{n}$ where $N - \sum_{j=1}^{M-1} n_j \geq 0$. If we denote $\sum_{\boldsymbol{n}}$ as a sum over these allowable configurations, then $\sum_{\boldsymbol{n}} \mathcal{P}(\boldsymbol{n};\boldsymbol{x};\lambda(t)) = 1$. The genotype fraction x_i' in the new generation is just $x_i' = n_i/N$, and hence the mean difference $v_i(\boldsymbol{x};\lambda(t))$ in genotype fractions per generation can be expressed as

$$v_i(\boldsymbol{x};\lambda(t)) = \langle \delta x_i \rangle = \langle x_i' - x_i \rangle = \sum_{\boldsymbol{n}} \left(\frac{n_i}{N} - x_i \right) \mathcal{P}(\boldsymbol{n};\boldsymbol{x},\lambda(t)) = \phi_i(\boldsymbol{x};\lambda(t)) - x_i. \tag{S18}$$

Similarly $D_{ij}(\boldsymbol{x};\lambda(t))$ can be calculated through the covariance as:

$$\begin{aligned} 2D_{ij}(\boldsymbol{x};\lambda(t)) &= \langle \delta x_i \delta x_j \rangle - \langle \delta x_i \rangle \langle \delta x_j \rangle \\ &= \left\langle \left(\frac{n_i}{N} - x_i\right)\left(\frac{n_j}{N} - x_j\right)\right\rangle - (\phi_i(\boldsymbol{x};\lambda(t)) - x_i)\,(\phi_j(\boldsymbol{x};\lambda(t)) - x_j) \\ &= \frac{g_{ij}(\boldsymbol{\phi}(\boldsymbol{x};\lambda(t)))}{N}, \end{aligned} \tag{S19}$$

where $g_{ij}(\boldsymbol{\phi}(\boldsymbol{x};\lambda(t)))$ is given by Methods Eq. (7) with $\boldsymbol{\phi}(\boldsymbol{x};\lambda(t)) = (\phi_1(\boldsymbol{x};\lambda(t)),\ldots,\phi_{M-1}(\boldsymbol{x};\lambda(t)))$ replacing $\boldsymbol{x}$. In order to derive Eqs. (S18)-(S19), we have used the mean and covariance properties of the multinomial distribution of Eq. (S17):

$$\begin{aligned} \langle n_i \rangle &= N\phi_i(\boldsymbol{x};\lambda(t)), \\ \langle n_i n_j \rangle - \langle n_i \rangle \langle n_j \rangle &= \begin{cases} N\phi_i(\boldsymbol{x};\lambda(t))(1-\phi_i(\boldsymbol{x};\lambda(t))) & j = i \\ -N\phi_i(\boldsymbol{x};\lambda(t))\phi_j(\boldsymbol{x};\lambda(t)) & j \neq i \end{cases}. \end{aligned} \tag{S20}$$

To complete the derivation, we need an expression for $\phi_i(\boldsymbol{x};\lambda(t))$ when mutation and selection are included in the model. Consider first the probability $\phi_i^0(\boldsymbol{x})$ of picking a parent of type i, assuming only mutation was allowed, but no selection. Accounting for the gain and loss of possible type i parents through mutation, we have $\phi_i^0(\boldsymbol{x}) = x_i + m_{i\mu}x_\mu$, for $1 \leq i \leq M-1$. We can express the probability to choose a parent of type M as $\phi_M^0(\boldsymbol{x}) = 1 - \sum_{j=1}^{M-1}\phi_j^0(\boldsymbol{x})$ by normalization. To include selection and obtain the final expressions for $\phi_i(\boldsymbol{x};\lambda(t))$, we note that we can write

$$\frac{\phi_i(\boldsymbol{x};\lambda(t))}{\phi_M(\boldsymbol{x};\lambda(t))} = (1+s_i(\lambda(t)))\frac{\phi_i^0(\boldsymbol{x})}{\phi_M^0(\boldsymbol{x})} \qquad [\text{no summation over } i], \tag{S21}$$

Eq. (S21), for $i = 1,\ldots,M-1$, states that the ratio of $\phi_i(\boldsymbol{x};\lambda(t))$, the chance of picking a type i parent, to $\phi_M(\boldsymbol{x};\lambda(t))$, the chance of picking a wild type parent, is modified by a factor $1+s_i(t)$ due to selection, compared to the case without selection. In other words, if $s_i(t)$ is positive because type i has a greater fitness than the wild type, the chance of getting a type of i parent relative to a wild type parent increases by $1+s_i(t)$. Recalling the definition of $\phi_M(\boldsymbol{x};\lambda(t))$ in terms of $\boldsymbol{\phi}(\boldsymbol{x};\lambda(t))$'s components, we can solve the systems of equations in Eq. (S21) to find

$$\phi_i(\boldsymbol{x};\lambda(t)) = \frac{(1+s_i(\lambda(t)))\phi_i^0(\boldsymbol{x})}{1+s_j(\lambda(t))\phi_j^0(\boldsymbol{x})} \quad [\text{no summation over } i]. \tag{S22}$$

Let us substitute in the expression for $\phi_i^0(\boldsymbol{x})$, and make the typical assumption that $|s_i(\lambda(t))|$, $|m_{ij}| \ll 1$. After Taylor expanding to first order in these quantities, Eq. (S22) becomes

$$\phi_i(\boldsymbol{x};\lambda(t)) \approx x_i + m_{i\mu}x_\mu + g_{ij}(\boldsymbol{x})s_j(\lambda(t)). \tag{S23}$$

Plugging this into Eq. (S18) gives the expression in Methods Eq. (6) for v_i. Similarly, if we plug Eq. (S23) into Eq. (S19), and keep only the leading order contribution, we get $D_{ij}(\boldsymbol{x}) \approx g_{ij}(\boldsymbol{x})/(2N)$, which is the expression in Methods Eq. (6) for D_{ij}.

3 Estimating the mean and covariance of the instantaneous equilibrium distribution

As discussed in the Methods, for general M we do not have an analytical expression for IE distribution. However in the large population, frequent mutation regime we can use the multivariate normal approximation of Methods Eq. (10). This requires us to be able to calculate the mean genotype frequencies $\overline{x}_i(\lambda)$ and covariance matrix $\Sigma(\lambda)$ associated with the IE distribution $\rho(\boldsymbol{x};\lambda)$ at a given value of the control parameter λ. In this section we outline how to do this calculation. The first step is deriving a set of coupled equations for the first and second moments of the IE distribution (part of a larger hierarchy of moment equations). The second step will be to approximately solve these equations using a moment closure technique. We consider each step in turn.

3.1 Deriving equations for the first and second moments of the IE distribution

In order to characterize the moments of the IE distribution $\rho(\boldsymbol{x};\lambda)$, let us consider an auxiliary problem: imagine a system described by a genotype frequency probability $\tilde{p}(\boldsymbol{x},\tau)$ evolving under a Fokker-Planck equation with constant

λ:

$$\begin{aligned}\partial_\tau \tilde{p}(\boldsymbol{x},\tau) &= \mathcal{L}(\lambda)\tilde{p}(\boldsymbol{x},\tau)\\ &= -\partial_i\left[v_i(\boldsymbol{x};\lambda)\tilde{p}(\boldsymbol{x},\tau)+\partial_j\left(D_{ij}(\boldsymbol{x})\tilde{p}(\boldsymbol{x},\tau)\right)\right]\\ &\equiv -\partial_i J_i(\boldsymbol{x},\tau),\end{aligned} \tag{S24}$$

where for later convenience we have introduced the probability current $J_i(\boldsymbol{x},\tau)$. We know that in the limit $\tau\to\infty$ the system equilibrates, $\tilde{p}(\boldsymbol{x},\tau)\to\rho(\boldsymbol{x};\lambda)$. We use the auxiliary time τ here for clarity, since it is distinct from the actual time t in the original system. Let us define the first and second moments of the genotype frequencies with respect to $\tilde{p}(\boldsymbol{x},\tau)$:

$$\langle x_i\rangle_\tau \equiv \int_\Delta d\boldsymbol{x}\, x_i\tilde{p}(\boldsymbol{x},\tau), \qquad \langle x_ix_j\rangle_\tau \equiv \int_\Delta d\boldsymbol{x}\, x_ix_j\tilde{p}(\boldsymbol{x},\tau), \tag{S25}$$

where the τ subscript denotes the dependence of the moments on τ. Since $\tilde{p}(\boldsymbol{x},\tau)\to\rho(\boldsymbol{x};\lambda)$ as $\tau\to\infty$, the IE distribution quantities we are interested in are just the following limiting values:

$$\overline{x}_i(\lambda) = \lim_{\tau\to\infty}\langle x_i\rangle_\tau, \qquad \Sigma_{ij}(\lambda) = \lim_{\tau\to\infty}\left(\langle x_ix_j\rangle_\tau - \langle x_i\rangle_\tau\langle x_j\rangle_\tau\right). \tag{S26}$$

We will derive the following exact moment relationships:

$$0 = m_{i\mu}\overline{x}_\mu(\lambda) + \left[\overline{x}_i(\lambda)\delta_{ik} - \Sigma_{ik}(\lambda) - \overline{x}_i(\lambda)\overline{x}_k(\lambda)\right]s_k(\lambda), \tag{S27}$$

$$\begin{aligned}0 &= m_{i\mu}\Sigma_{j\mu}(\lambda) + m_{j\mu}\Sigma_{i\mu}(\lambda) + \Sigma_{ij}(\lambda)\left(s_i(\lambda)+s_j(\lambda)\right)\\ &\quad -\left[2T_{ijk}(\lambda) - \overline{x}_j(\lambda)\Sigma_{ik}(\lambda) - \overline{x}_i(\lambda)\Sigma_{jk}(\lambda) - 2\overline{x}_i(\lambda)\overline{x}_j(\lambda)\overline{x}_k(\lambda)\right]s_k(\lambda)\\ &\quad +\frac{1}{N}\left[\overline{x}_i(\lambda)\delta_{ij} - \Sigma_{ij}(\lambda) - \overline{x}_i(\lambda)\overline{x}_j(\lambda)\right],\end{aligned} \tag{S28}$$

where $i,j=1,\ldots,M-1$ (the indices i, j are not summed over). Here $T_{ijk}(\lambda)\equiv\lim_{\tau\to\infty}\langle x_ix_jx_k\rangle_\tau$ is a third moment of the IE distribution. The derivation of these equations is shown below. For readers not interested in the details, they can skip ahead to Sec. 3.2.

To find Eq. (S27), let us start with the first moment $\langle x_i\rangle_\tau$ from Eq. (S25), take the derivative with respect to τ, and plug in the Fokker-Planck equation from Eq. (S24):

$$\frac{d}{d\tau}\langle x_i\rangle_\tau = \int_\Delta d\boldsymbol{x}\, x_i\partial_\tau\tilde{p}(\boldsymbol{x},\tau) = -\int_\Delta d\boldsymbol{x}\, x_i\partial_j J_j(\boldsymbol{x},\tau), \tag{S29}$$

Notice that $x_i\partial_jJ_j = \partial_j(x_iJ_j) - \delta_{ij}J_j$. By Gauss's theorem, the integral over the first term is $\int_\Delta d\boldsymbol{x}\,\partial_j(x_iJ_j) = \int_{\partial\Delta} d\sigma\, x_iJ_jn_j$. The latter integral is expressed in terms of area elements $d\sigma$ of the simplex boundary $\partial\Delta$, and n_j is the jth component of the normal vector to this boundary. Since probability is conserved within the simplex, $J_jn_j=0$, and hence the integral vanishes. Thus only the integral over the second term contributes, and Eq. (S29) can be rewritten:

$$\frac{d}{d\tau}\langle x_i\rangle_\tau = \int d\boldsymbol{x}\, J_i(\boldsymbol{x},\tau) = \int d\boldsymbol{x}\left[v_i(\boldsymbol{x};\lambda)\tilde{p}(\boldsymbol{x},\tau) - \partial_j\left(D_{ij}(\boldsymbol{x})\tilde{p}(\boldsymbol{x},\tau)\right)\right]. \tag{S30}$$

Focusing on the second term of the integrand in Eq. (S30), we can rewrite this term using Gauss's law as:

$$-\int_\Delta d\boldsymbol{x}\,\partial_j\left(D_{ij}(\boldsymbol{x})\tilde{p}(\boldsymbol{x},\tau)\right) = -\int_{\partial\Delta} d\sigma\,\tilde{p}(\boldsymbol{x},\tau)D_{ij}(\boldsymbol{x})n_j = -\frac{1}{2N}\int_{\partial\Delta} d\sigma\,\tilde{p}(\boldsymbol{x},\tau)g_{ij}(\boldsymbol{x})n_j. \tag{S31}$$

The term $g_{ij}(\boldsymbol{x})n_j=0$ for $\boldsymbol{x}\in\partial\Delta$, which makes the integral vanish. We prove this in the following lemma.

Lemma: For $g_{ij}(\boldsymbol{x})$ as defined in Methods Eq. (7) and for any normal vector $\boldsymbol{n}$ to the simplex boundary $\partial\Delta$, $g_{ij}(\boldsymbol{x})n_j=0$ for all $\boldsymbol{x}\in\partial V$.

Proof: A simplex is defined by $\Delta=\{\boldsymbol{x}\,|\,x_i\geq 0, x_ie_i\leqslant 1\}$, where $\boldsymbol{e}$ is an $M-1$ dimensional vector with all components equal to 1. There are two classes of hypersurfaces on the simplex boundary (where $x_ie_i=1$) to consider:

Case 1: One type of hypersurface on the simplex boundary is S_k, defined by the conditions $x_k = 0$ and $x_j e_j = 1$. There are $M-1$ such hypersurfaces, one for each $k = 1, \ldots, M-1$. The components of the normal vector $\boldsymbol{n}_k$ to S_k are given by $n_{kj} = -\delta_{kj}$ (the minus sign ensures $\boldsymbol{n}_k$ faces away from the simplex volume). Then, for all $\boldsymbol{x} \in S_k$, $g_{ij}(\boldsymbol{x}) n_{kj} = -g_{ij}(\boldsymbol{x})\delta_{kj} = -g_{ik}(\boldsymbol{x})$, and we know from the definition of g that $g_{ik}(\boldsymbol{x}) \propto x_k = 0$. The last result follows because $x_k = 0$ for $\boldsymbol{x} \in S_k$.

Case 2: The only other type of hypersurface on the simplex boundary is S', defined by the conditions $x_i > 0$ for all i, and $x_i e_i = 1$. We find the normal to this surface as follows. Define $F(\boldsymbol{x}) = x_i e_i - 1$. Then $\boldsymbol{n} = \nabla F = \boldsymbol{e}$. For all $\boldsymbol{x} \in S'$, we have $g_{ij}(\boldsymbol{x}) n_j = g_{ij}(\boldsymbol{x}) e_j = x_i(1 - x_j e_j)$. Since $x_j e_j = 1$ for all $\boldsymbol{x} \in S'$, we see that $g_{ij}(\boldsymbol{x}) n_j = 0$.

Using the lemma, and the definition of $v_i(\boldsymbol{x};\lambda)$ from Methods Eq. (6), we can rewrite Eq. (S30) as:

$$\begin{aligned} \frac{d}{d\tau}\langle x_i \rangle_\tau = \int d\boldsymbol{x}\, v_i(\boldsymbol{x};\lambda)\tilde{p}(\boldsymbol{x},\tau) &= \langle v_i(\boldsymbol{x};\lambda)\rangle_\tau \\ &= m_{i\mu}\langle x_\mu \rangle_\tau + \langle g_{ij}(\boldsymbol{x})\rangle_\tau s_j(\lambda). \end{aligned} \tag{S32}$$

Taking the $\tau \to \infty$ limit on sides of the equation, we note that the left-hand side vanishes because $\langle x_i \rangle_\tau \to \overline{x}_i(\lambda)$, a constant independent of τ. On the right-hand side we can subsitute in Eq. (S25) and use the definition of g in Methods Eq. (7). The end result is Eq. (S27). This is the first moment equation we will be interested in.

To derive Eq. (S28), we start analogously to Eq. (S29), but now with the second moment:

$$\frac{d}{d\tau}\langle x_i x_j \rangle_\tau = \int_\Delta d\boldsymbol{x}\, x_i x_j \partial_\tau \tilde{p}(\boldsymbol{x},\tau) = -\int_\Delta d\boldsymbol{x}\, x_i x_j \partial_k J_k(\boldsymbol{x},\tau). \tag{S33}$$

Notice, by the product rule, that $x_i x_j \partial_k J_k = \partial_k(x_i x_j J_k) - J_k \partial_k(x_i x_j) = \partial_k(x_i x_j J_k) - J_i x_j - J_j x_i$. By Gauss' law, we have $\int_\Delta d\boldsymbol{x}\, \partial_k(x_i x_j J_k) = \int_{\partial\Delta} d\sigma\, x_i x_j J_k n_k$. Since probability is conserved, $J_k n_k = 0$, and we can thus rewrite Eq. (S33) as:

$$\begin{aligned} \frac{d}{d\tau}\langle x_i x_j \rangle_\tau &= \int d\boldsymbol{x}\,(J_i x_j + J_j x_i) \\ &= \int d\boldsymbol{x}\Big[\big(v_i(\boldsymbol{x};\lambda)\tilde{p}(\boldsymbol{x},\tau) - \partial_k\big(D_{ik}(\boldsymbol{x})\tilde{p}(\boldsymbol{x},\tau)\big)\big)\, x_j + \big(v_j(\boldsymbol{x};\lambda)\tilde{p}(\boldsymbol{x},\tau) - \partial_k\big(D_{jk}(\boldsymbol{x})\tilde{p}(\boldsymbol{x};\tau)\big)\big)\, x_i\Big]. \end{aligned} \tag{S34}$$

Note that $\partial_k(D_{ik}\tilde{p})x_j = \partial_k(D_{ik}\tilde{p}x_j) - D_{ik}\tilde{p}\partial_k x_j = \partial_k(D_{ik}\tilde{p}x_j) - D_{ik}\tilde{p}\delta_{jk} = \partial_k(D_{ik}\tilde{p}x_j) - D_{ij}\tilde{p}$. By Gauss's theorem,

$$-\int_\Delta d\boldsymbol{x}\,\partial_k\left(D_{ik}(\boldsymbol{x})\tilde{p}(\boldsymbol{x},\tau)x_j\right) = -\int_{\partial\Delta} d\sigma\,\tilde{p}(\boldsymbol{x},\tau)x_j D_{ik}(\boldsymbol{x})n_k = -\frac{1}{2N}\int_{\partial\Delta} d\sigma\,\tilde{p}(\boldsymbol{x},\tau)x_j g_{ik}(\boldsymbol{x})n_k. \tag{S35}$$

By the lemma, $g_{ik}(\boldsymbol{x})n_k = 0$ for $\boldsymbol{x} \in \partial\Delta$, so the integral vanishes. Using this and the fact that $D_{ij} = D_{ji}$, Eq. (S34) becomes

$$\begin{aligned} \frac{d}{d\tau}\langle x_i x_j \rangle_\tau &= \int_\Delta d\boldsymbol{x}\left[x_j v_i(\boldsymbol{x};\lambda)\tilde{p}(\boldsymbol{x},\tau) + x_i v_j(\boldsymbol{x};\lambda)\tilde{p}(\boldsymbol{x},\tau) + 2D_{ij}(\boldsymbol{x})\tilde{p}(\boldsymbol{x},\tau)\right] \\ &= \langle x_j v_i(\boldsymbol{x};\lambda)\rangle_\tau + \langle x_i v_j(\boldsymbol{x};\lambda)\rangle_\tau + 2\langle D_{ij}(\boldsymbol{x})\rangle_\tau. \end{aligned} \tag{S36}$$

In the $\tau \to \infty$ limit this equation can be written in the form of Eq. (S28).

3.2 Approximate solution of moment equations

Eqs. (S27) and (S28) constitute the first two of a hierarchy of coupled moment equations, with each set of equations involving moments of one higher order (i.e. Eq. (S27) involves Σ_{ij}, Eq. (S28) involves T_{ijk}). In the limit $N \to \infty$ Eq. (S28) for the second moments can be trivially satisfied, because the IE distribution becomes a delta function with zero spread. In this case $\Sigma_{ij}(\lambda) = 0$, $T_{ijk}(\lambda) = \overline{x}_i(\lambda)\overline{x}_j(\lambda)\overline{x}_k(\lambda)$ is a solution to Eq. (S28). For large but finite N (assuming $\mu N \gg 1$ is still satisfied, where μ is the order of magnitude of the nonzero mutation rates) the distribution becomes spread out by a small amount, and we will approximate it by a multivariate Gaussian. As a result we assume third and higher order cumulants are negligible, which will allow us to approximately solve Eqs. (S27) and (S28) for $\overline{x}_i(\lambda)$ and $\Sigma_{ij}(\lambda)$. This approach, known as moment closure, effectively truncates the moment hierarchy after the second order. It is justified by the fact that while $\Sigma_{ij}(\lambda)$ scales like N^{-1}, the third cumulant scales like

N^{-2}, etc., which allows the higher order cumulants to be neglected for large N. Letting the third cumulant be zero means the third moment can be approximated as follows:

$$T_{ijk}(\lambda) \approx \overline{x}_i(\lambda)\Sigma_{jk}(\lambda) + \overline{x}_j(\lambda)\Sigma_{ik}(\lambda) + \overline{x}_k(\lambda)\Sigma_{ij}(\lambda) + \overline{x}_i(\lambda)\overline{x}_k(\lambda)\overline{x}_k(\lambda). \tag{S37}$$

Plugging this into Eq. (S28) gives:

$$\begin{aligned} 0 \approx\ & m_{i\mu}\Sigma_{j\mu}(\lambda) + m_{j\mu}\Sigma_{i\mu}(\lambda) + \Sigma_{ij}(\lambda)\big(s_i(\lambda) + s_j(\lambda)\big) \\ & - \big[2\overline{x}_k(\lambda)\Sigma_{ij}(\lambda) + \overline{x}_i(\lambda)\Sigma_{jk}(\lambda) + \overline{x}_j(\lambda)\Sigma_{ik}(\lambda)\big]s_k(\lambda) \\ & + \frac{1}{N}\big[\overline{x}_i(\lambda)\delta_{ij} - \Sigma_{ij}(\lambda) - \overline{x}_i(\lambda)\overline{x}_j(\lambda)\big]. \end{aligned} \tag{S38}$$

For Eq. (S27), since the term $\Sigma_{ik}(\lambda) \sim \mathcal{O}(N^{-1})$ becomes negligible relative to the other terms for large N, we can approximate the equation as

$$0 \approx m_{i\mu}\overline{x}_\mu(\lambda) + \big[\overline{x}_i(\lambda)\delta_{ik} - \overline{x}_i(\lambda)\overline{x}_k(\lambda)\big]s_k(\lambda). \tag{S39}$$

The following procedure can then be used to solve Eqs. (S38)-(S39): i) numerically solve the nonlinear set of equations in Eq. (S39) for $\overline{x}_i(\lambda)$, $i = 1, \ldots, M-1$. ii) Plug this solution into Eq. (S38), and then numerically solve the resulting set of linear equations for $\Sigma_{ij}(\lambda)$, $i, j = 1, \ldots, M-1$. Because the Σ indices in Eq. (S38) run up to M, we use the following identities to express those elements in terms of lower indices: $\Sigma_{MM} = \Sigma_{k\ell}e_k e_\ell$, $\Sigma_{iM} = -\Sigma_{ik}e_k$. Both of these identities follow from the fact that $x_M = 1 - x_k e_k$.

In principle we can iterate this procedure to progressively add small corrections to the solution, converging to self-consistency between Eq. (S27) and Eq. (S38): plug the $\Sigma_{ij}(\lambda)$ values obtained from the first iteration into Eq. (S27), solve for the updated $\overline{x}_i(\lambda)$, plug these into Eq. (S38), and so on. However for all the cases we examined in the main text the corrections resulting from multiple iterations are negligible, so we use one iteration only.

4 Approximating the counterdiabatic driving protocol in the large population, frequent mutation regime

As derived in Methods Sec. 4, the selection coefficient perturbation $\delta\tilde{\boldsymbol{s}}\big(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)\big) = \tilde{\boldsymbol{s}}\big(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)\big) - \boldsymbol{s}(\lambda(t))$ needed to implement the CD protocol satisfies Methods Eq. (13):

$$\partial_t \rho\big(\boldsymbol{x}; \lambda(t)\big) = -\partial_i\Big(\rho\big(\boldsymbol{x}; \lambda(t)\big) g_{ij}(\boldsymbol{x}) \delta\tilde{s}_j\big(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)\big)\Big), \tag{S40}$$

which in turn implies the relation in Methods Eq. (17),

$$\partial_t \overline{\boldsymbol{x}}(\lambda(t)) = \Big\langle \boldsymbol{g}(\boldsymbol{x}) \delta\tilde{\boldsymbol{s}}\big(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)\big)\Big\rangle. \tag{S41}$$

Here the brackets $\langle\,\rangle$ denote an average over the simplex with respect to the IE distribution $\rho(\boldsymbol{x}; \lambda(t))$. Let us define a function $\boldsymbol{F}(\boldsymbol{x}) \equiv \boldsymbol{g}(\boldsymbol{x})\delta\tilde{\boldsymbol{s}}\big(\boldsymbol{x}; \lambda(t), \dot{\lambda}(t)\big)$. For simplicity of notation we do not explicitly show the $\lambda(t)$, $\dot{\lambda}(t)$ dependence in $\boldsymbol{F}(\boldsymbol{x})$. We can then Taylor expand the right-hand side of Eq. (S41) around $\overline{\boldsymbol{x}}(\lambda(t))$ up to second order. In component form, this looks like

$$\begin{aligned} \partial_t \overline{x}_i(\lambda(t)) &= \langle F_i(\boldsymbol{x})\rangle \\ &= F_i\big(\overline{\boldsymbol{x}}(\lambda(t))\big) + \partial_j F_i\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\langle x_j - \overline{x}_j(\lambda(t))\rangle + \frac{1}{2}\partial_j\partial_k F_i\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\big\langle\big(x_j - \overline{x}_j(\lambda(t))\big)\big(x_k - \overline{x}_k(\lambda(t))\big)\big\rangle + \cdots \\ &= F_i\big(\overline{\boldsymbol{x}}(\lambda(t))\big) + \frac{1}{2}\partial_j\partial_k F_i\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\Sigma_{jk}(\lambda(t)) + \cdots \end{aligned} \tag{S42}$$

In the last line we have used the definition of the IE covariance matrix $\Sigma_{jk}(\lambda(t))$, and the fact that $\langle x_j - \overline{x}_j(\lambda(t))\rangle = 0$ since $\overline{x}_j(\lambda(t))$ is the mean of the IE distribution. From the discussion in the previous section we know that $\Sigma_{jk}(\lambda(t))$ scales like N^{-1} when N is large (and $\mu N \gg 1$). This means that the Σ term in Eq. (S42) becomes small compared to the leading term for large N. Keeping only the leading term, and substituting in the definition of $\boldsymbol{F}(\boldsymbol{x})$, we find the approximate relation

$$\partial_t \overline{\boldsymbol{x}}(\lambda(t)) \approx \boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\delta\tilde{\boldsymbol{s}}\big(\lambda(t), \dot{\lambda}(t)\big), \tag{S43}$$

where $\delta\tilde{\boldsymbol{s}}\big(\lambda(t), \dot{\lambda}(t)\big) \equiv \delta\tilde{\boldsymbol{s}}\big(\overline{\boldsymbol{x}}(\lambda(t)); \lambda(t), \dot{\lambda}(t)\big)$. This is what is shown in Methods Eq. (18).

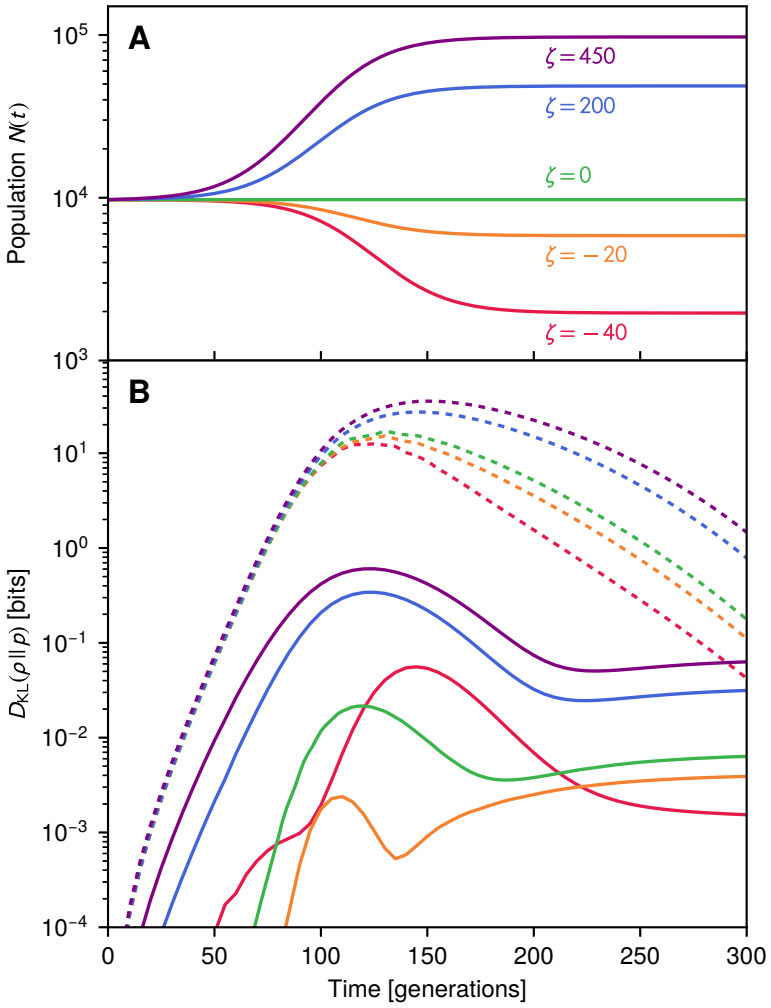

Figure S1. **(A)** Five examples of time-varying total populations $N(t) = N_0(1+\zeta s_i(\lambda(t))$ for the two genotype system described in Methods Sec. 4.4. Here $N_0 = 10^4$ and $\zeta = -40$, -20, 0, 200, and 450. **(B)** For each case in panel **A**, the Kullback-Leibler divergence between the IE and actual genotype frequency distributions while driving the system according to the original protocol (dashed curves) or the approximate CD protocol (solid curves). Colors are matched to the ζ values indicated in panel **A**. The KL divergences are calculated using numerical solutions to the associated Fokker-Planck equations.

5 Counterdiabatic driving under time-varying total populations

The derivation of the CD protocol discussed in Methods Sec. 4 and the previous section of the SI remains valid even when the total population $N(t)$ varies in time as a result of the control protocol, i.e. the carrying capacity of the system changes along with the control parameters. In this case we can effectively absorb $N(t)$ into the set of control parameters $\lambda(t)$ during the derivation, and we end up with the same approximate CD protocol defined through Eq. (S43), and whose explicit solution is shown in main text Eq. (4). This assumes the conditions for the validity of the approximation hold at all times of interest, $N(t) \gg 1$ and $\mu N(t) \gg 1$. It is interesting to note that main text Eq. (4) depends only on the selection coefficients under the protocol $s_i(\lambda(t))$ and the corresponding instantaneous equilibrium mean genotype frequencies $\overline{x}_i(\lambda(t))$. As can be seen from Eq. (S39), to leading order for large $N(t)$ the means $\overline{x}_i(\lambda(t))$ are independent of $N(t)$. If the selection coefficients are also independent of $N(t)$ then the entire CD driving protocol becomes (at least to leading order) independent of $N(t)$. Thus the same CD protocol should work for a variety of $N(t)$ behaviors.

To illustrate this, we have redone the two genotype CD driving results from main text Fig. 2 using time-varying $N(t)$. All parameters are as described in Methods Sec. 4.4, except that $N(t)$ varies with the control parameter

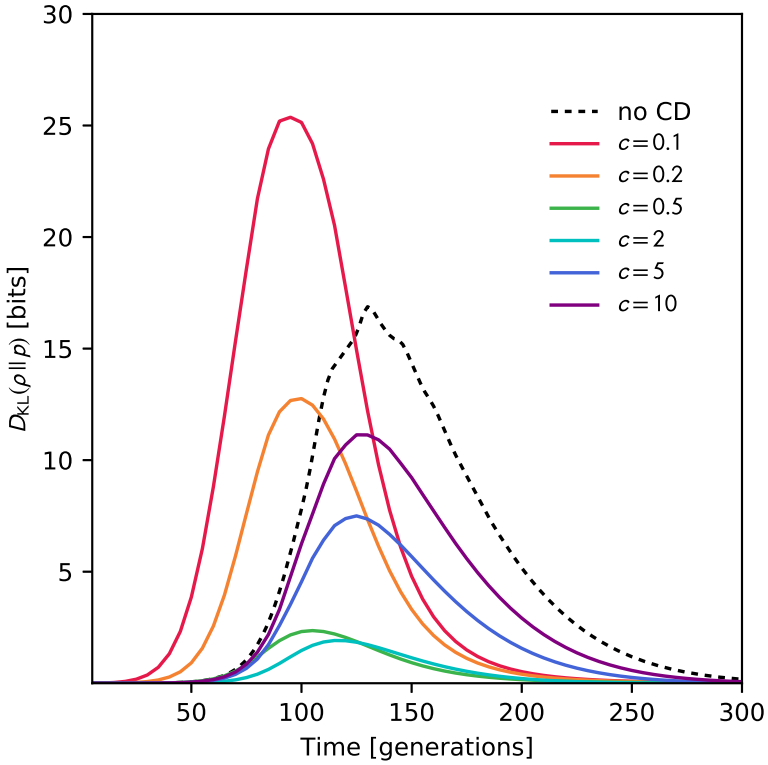

Figure S2. The KL divergence between the actual and IE genotype frequency distributions for the $M = 2$ system described in Methods Sec. 4.4, calculated using numerical solutions to the Fokker-Planck equation. The dashed curve shows the results of the original protocol, while the colored curves show deliberately inaccurate CD protocols: the mutation rates that enter into the CD solution of Methods Eq. (20) are scaled by a factor c to mimic experimental measurement errors, where $c = 1$ corresponds to the true CD protocol. Six different inaccurate protocols are shown, with c ranging from 0.1 to 10.

according to: $N(t) = N_0(1+\zeta s_1(\lambda(t)))$, where $N_0 = 10^4$ and ζ is a constant. The original two genotype results for constant N are recovered when $\zeta = 0$. Fig. S1A shows five forms for $N(t)$ for different ζ, and Fig. S1B shows the corresponding driving results, calculated using numerical solutions of the Fokker-Planck dynamics (main text Eq. (1)). The dashed curves show the Kullback-Leibler divergence between the IE and actual genotype frequency distributions using the original protocol, and the solid curves using the approximate CD protocol (which is independent of $N(t)$). In all cases CD driving dramatically reduces lag, improving the agreement between IE and actual distributions: the KL divergences under CD remain below 1 bit throughout the whole protocol, in contrast to the original cases where the divergence peaks above 10 bits.

6 Robustness of counterdiabatic driving to errors in the protocol

In the main text we explored one way in which counterdiabatic driving may still be effective even in scenarios where precise implementation of the protocol described by main text Eq. (4) is hampered by practical constraints. In the $M = 16$ example direct realization of the full $\tilde{\boldsymbol{s}}(\lambda(t), \dot{\lambda}(t))$ solution is constrained by the fact that there is one control variable (drug concentration) and this control variable may have a limited range (no larger than a certain maximum allowable dosage). However we showed protocols that approximate the CD solution could still produce excellent results in terms of driving the system near the desired trajectory of genotype distributions over finite times.

Here we look at another potential hindrance: what if there were errors in the quantities we use to calculate CD driving via main text Eq. (4)? The latter involves the mean IE genotypes frequencies $\overline{\boldsymbol{x}}(\lambda(t))$ and selection coefficients $\boldsymbol{s}(\lambda(t))$ in the original protocol as a function of the control parameter. Using the moment relationship of Eq. (S39), $\overline{\boldsymbol{x}}(\lambda(t))$ can in turn be calculated with knowledge of the mutation rate matrix m and $\boldsymbol{s}(\lambda(t))$. We can imagine different potential sources of error: i) One possibility is that our estimate for $\boldsymbol{s}(\lambda(t))$ of the original

protocol contains some inaccuracies. For example, in the case where the control parameter is a drug concentration, the genotype fitness versus drug dosage curve from earlier experiments might have measurement artifacts. Assuming that we had a good estimate of the mutation rates $m_{\mu\nu}$, we could still calculate the correct $\overline{\boldsymbol{x}}(\lambda(t))$ associated with each value of our $\boldsymbol{s}(\lambda(t))$ trajectory through Eq. (S39). Hence main text Eq. (4) would still yield a valid CD protocol, in the sense that the system would be approximately guided through a series of IE distributions corresponding to $\boldsymbol{s}(\lambda(t))$, assuming we could implement the calculated $\tilde{\boldsymbol{s}}(\lambda(t),\dot{\lambda}(t))$. It would not be precisely the series of IE distributions actually associated with $\lambda(t)$, but nevertheless Eq. (4) would still provide a recipe for following a path in the space of IE distributions. ii) Another possibility is that our estimate for the mutation rates $m_{\mu\nu}$ is flawed. This would propagate into errors in our calculation of $\overline{\boldsymbol{x}}(\lambda(t))$ associated with $\boldsymbol{s}(\lambda(t))$. Because of these errors the results of Eq. (4) would no longer necessarily be a CD protocol for $\boldsymbol{s}(\lambda(t))$.

To understand the effects of errors in the mutation rates, we investigated the $M=2$ example with deliberately inaccurate protocols. For $M=2$ the expression for the mean genotype frequency $\overline{x}_1(\lambda(t))$ is given by Methods Eq. (19), which in turn yields the CD protocol solution in Methods Eq. (20). We modified the protocol by scaling the mutation rates in Eq. (20) by a factor of c, so $m_{12}=m_{21}=c\mu_0$, where $\mu_0=2.5\times10^{-3}$ is the actual mutation rate of the $M=2$ system. All other parameters as described in Methods Sec. 4.4. The value $c=1$ corresponds to the true CD protocol, while other values correspond to inaccurate protocols. As seen in main text Fig. 2D, the true CD protocol works effectively at all times, with the KL divergence between the IE and actual genotype distributions getting no larger than about 0.02 bits. In contrast, the original protocol leads to severe lag at intermediate times, with the KL divergence peaking around 17 bits. Fig. S2 shows KL divergence results for inaccurate protocols, with c ranging from 0.1 to 10, calculated using a numerical Fokker-Planck approach. As expected, we see a breakdown of CD driving, with the IE and actual distributions differing dramatically at intermediate times. For the most inaccurate protocols, at $c=0.1$ and $c=10$, the peaks in the KL divergence are comparable to or greater than those without CD. However milder errors, like $c=0.5$ and $c=2$, perform much better than the original protocol, with the divergence peaking around 2 bits.

There is one silver lining in the error analysis: even though we lose a degree of control over the genotype distributions at intermediate times (they no longer follow precisely the desired series of IE distributions over time), we still retain some benefits in arriving at our destination faster. All the inaccurate protocols analyzed, particularly those where the mutation rates were underestimated ($c<1$) showed faster convergence to the final equilibrium distribution at long times (i.e. faster decay of the KL divergence) than the original protocol (dashed curve in Fig. S2). Thus some practical benefit of the calculated protocol remains for errors up to an order of magnitude in the mutation rates. Clearly having as good an estimate as possible of the mutation rates is beneficial for complete control, but the approach has some tolerance for the inevitable inaccuracies that will enter into experimentally estimated system parameters.

7 Mapping the agent-based model to a Fokker-Planck equation

In order to derive the mapping summarized in Methods Sec. 4.5.3, the first step is describing the dynamics of the agent-based model (ABM) over each time step in terms of a chemical Langevin equation[5]. This in turns allows us to map the ABM to an effective Fokker-Planck equation, using the standard relationship between Langevin and Fokker-Planck dynamics[6]. At simulation time step τ we have $n_\mu(\tau)$ organisms with genotype μ, and the ABM code will update this to $n_\mu(\tau+1)$ at the next time step. So long as we are interested in time scales much larger than a single simulation time step, and the updates per step are small compared to $n_\mu(\tau)\gg1$, we can treat $n_\mu(\tau)$ as a continuous population variable and τ as a continuous time variable, so that $n_\mu(\tau+1)-n_\mu(\tau)\approx dn_\mu(\tau)/d\tau$. In this continuum description, the rate of change $dn_\mu(\tau)/d\tau$ will be related to a series of stochastic "reactions", which we will label with an index $k=1,\ldots,R$, where R is the total number of possible reactions. Each reaction represents an aspect of the code that contributes to changes in the genotype populations. The amount by which the population of genotype μ changes in the kth reaction is denoted as $\Delta_{k\mu}$, and the probability of this reaction occurring at time step τ is $\pi_k(\tau)$. The corresponding chemical Langevin approximation describing the change in the population of genotype μ is given by[5]:

$$\frac{dn_\mu}{d\tau}=\sum_{k=1}^{R}\Delta_{k\mu}\pi_k(\tau)+\sum_{k=1}^{R}\Delta_{k\mu}\sqrt{\pi_k(\tau)}\Gamma_k(\tau). \tag{S44}$$

Note that throughout this section we will not be using the Einstein summation convention, so all sums will be explicitly indicated. The first term represents the deterministic contribution of all the reactions, while the second term is the corresponding noise introduced by the stochasticity of the reactions. The $\Gamma_k(\tau)$ functions are independent

Gaussian white noise functions with zero mean that satisfy $\langle \Gamma_k(\tau)\Gamma_{k'}(\tau')\rangle = \delta_{kk'}\delta(\tau-\tau')$. The average $\langle\,\rangle$ is taken over an ensemble of realizations of the system (i.e. an ensemble of simulation trajectories).

To complete the Langevin description, we have to identify $\Delta_{k\mu}$ and $\pi_k(\tau)$ for each reaction. Based on the ABM procedure described in Methods Sec. 4.5.3, we can enumerate the different types of reactions as follows:

- There are M possible cell death reactions, one for each genotype. If reaction k corresponds to the death of a genotype μ organism, then $\Delta_{k\mu} = -1$, $\Delta_{k\nu} = 0$ for all $\nu \neq \mu$, and $\pi_k(\tau) = n_\mu(\tau)d$, where d is the death probability at each simulation time step.
- There are M possible cell division reactions, one for each genotype. If reaction k corresponds to the division of a genotype μ organism, then $\Delta_{k\mu} = 1$, $\Delta_{k\nu} = 0$ for all $\nu \neq \mu$, and $\pi_k(\tau) = b_\mu(\tau)(1-d)n_\mu(\tau)$. Here $b_\mu(\tau)$ is the cell division probability given by Methods Eq. (21), which due to the carrying capacity in the system depends on the current total cell population $N_{\text{cell}}(\tau) = \sum_{\mu=1}^{M} n_\mu(\tau)$. The factor $(1-d)$ accounts for the fact that cell division can only occur if the cell survives the initial culling at the death step of the code.
- There are $M(M-1)$ possible mutation reactions (assumed to occur instantaneously after the cell division reactions) where newly born cells of one genotype mutate into a different genotype. If reaction k corresponds a new born cell of type ν mutating into a different type μ, then $\Delta_{k\mu} = 1$, $\Delta_{k\nu} = -1$, $\Delta_{k\sigma} = 0$ for all $\sigma \neq \mu, \nu$, and $\pi_k(\tau) = \hat{m}_{\mu\nu} b_\nu(\tau)(1-d)n_\nu(\tau)$ where $\hat{m}_{\mu\nu}$ is the mutation probability per time step in the ABM simulation.

Thus there are altogether $R = M(M+1)$ possible reactions. Before going further, we will narrow our focus to the system after it reaches equilibrium at some constant set of selection coefficients $s_i(\lambda)$. The total population will then fluctuate around some steady state mean value $\overline{N}_{\text{cell}} = \sum_{\mu=1}^{M}\langle n_\mu(\tau)\rangle$. The latter can be calculated by summing both sides of Eq. (S44) over all μ, and then taking the mean. The result is:

$$0 = \sum_{\mu=1}^{M}\sum_{k=1}^{R}\Delta_{k\mu}\langle\pi_k(\tau)\rangle = -d\overline{N}_{\text{cell}} + (1-d)\sum_{\mu=1}^{M}\langle b_\mu(\tau)n_\mu(\tau)\rangle. \tag{S45}$$

In the second equality we have explicitly plugged in all the different possible reactions and their probabilities. From Methods Eq. (21) we know that for $N_{\text{cell}}(\tau) < K$ the division probability is given by $b_\mu(\tau) = b_0(1+s_i(\lambda))(1 - N_{\text{cell}}(\tau)/K)$. For simplicity, let us ignore $N_{\text{cell}}(\tau)$ fluctuations in this expression and write

$$b_\mu(\tau) \approx b_0(1+s_\mu(\lambda))\left(1 - \frac{\overline{N}_{\text{cell}}}{K}\right). \tag{S46}$$

Substituting this into Eq. (S45) we get:

$$0 \approx -d\overline{N}_{\text{cell}} + (1-d)b_0\left(1 - \frac{\overline{N}_{\text{cell}}}{K}\right)\sum_{\mu=1}^{M}(1+s_\mu(\lambda))\langle n_\mu(\tau)\rangle. \tag{S47}$$

If we make the typical assumption that $|s_\mu(\lambda)| \ll 1$, then the sum on the right is approximately equal to $\overline{N}_{\text{cell}}$. This approximation still remains valid even if a subset of genotypes satisfies $|s_\mu(\lambda)| \ll 1$, while the others have substantial negative selection coefficients (as is the case for the fitness landscapes in our simulations). Then the latter group of less fit genotypes will have negligible populations relative to the former group, and the sum will still approximately evaluate to $\overline{N}_{\text{cell}}$. Using this simplification, we can solve Eq. (S47) for $\overline{N}_{\text{cell}}$:

$$\overline{N}_{\text{cell}} \approx K\left(1 - \frac{d}{b_0(1-d)}\right). \tag{S48}$$

Plugging this into Eq. (S46) we find

$$b_\mu(\tau) \approx (1+s_\mu(\lambda))\frac{d}{1-d}. \tag{S49}$$

This approximate form for $b_\mu(\tau)$ will prove useful below.

In order to connect the ABM dynamics described through Eq. (S44) to our Fokker-Planck formalism, we need to express these dynamics in terms of genotype frequencies, $x_\mu(\tau) = n_\mu(\tau)/N_{\rm cell}(\tau)$. Using the chain rule, the derivative of $x_\mu(\tau)$ with respect to τ can be written as:

$$\begin{aligned}\frac{dx_\mu(\tau)}{d\tau} &= \frac{1}{N_{\rm cell}(\tau)}\frac{dn_\mu(\tau)}{d\tau} - \frac{n_\mu(\tau)}{N_{\rm cell}^2(\tau)}\sum_{\nu=1}^{M}\frac{dn_\nu(\tau)}{d\tau} \\ &= \frac{(1-x_\mu(\tau))}{N_{\rm cell}(\tau)}\frac{dn_\mu(\tau)}{d\tau} - \frac{x_\mu(\tau)}{N_{\rm cell}(\tau)}\sum_{\nu\neq\mu}^{M}\frac{dn_\nu(\tau)}{d\tau}.\end{aligned} \tag{S50}$$

Using Eq. (S44) we can rewrite this as follows:

$$\frac{dx_\mu(\tau)}{d\tau} = \frac{1}{N_{\rm cell}(\tau)}\sum_{k=1}^{R}\pi_k(\tau)\left[(1-x_\mu(\tau))\Delta_{k\mu} - x_\mu(\tau)\sum_{\nu\neq\mu}^{M}\Delta_{k\nu}\right] + \eta_\mu(\tau)d, \tag{S51}$$

where $\eta_\mu(\tau)$ incorporates all the noise contributions,

$$\eta_\mu(\tau) \equiv \frac{1}{N_{\rm cell}(\tau)d}\sum_{k=1}^{R}\Gamma_k(\tau)\sqrt{\pi_k(\tau)}\left[(1-x_\mu(\tau))\Delta_{k\mu} - x_\mu(\tau)\sum_{\nu\neq\mu}^{M}\Delta_{k\nu}\right]. \tag{S52}$$

Let us plug in the details of all the reaction types into Eq. (S51), and use Eq. (S49) for $b_\mu(\tau)$ to simplify. For $\mu = i = 1,\ldots,M-1$, we find:

$$\frac{dx_i(\tau)}{d\tau} = \sum_{\nu\neq i}^{M}\hat{m}_{i\nu}(1+s_\nu(\lambda))x_\nu(\tau)d - \sum_{\nu\neq i}^{M}\hat{m}_{\nu i}(1+s_i(\lambda))x_i(\tau)d + \sum_{j=1}^{M-1}g_{ij}(\boldsymbol{x}(\tau))s_j(\lambda)d + \eta_i(\tau)d, \tag{S53}$$

where g is the matrix defined in Methods Eq. (7) and we note that $s_M(\lambda) = 0$ from the definition of the selection coefficients (since genotype M is the wild type). Let us define a rescaled time $t \equiv \tau d$, which allows Eq. (S53) to be written as:

$$\frac{dx_i(t)}{dt} = v_i(\boldsymbol{x}(t);\lambda) + \eta_i(t), \tag{S54}$$

where

$$v_i(\boldsymbol{x}(t);\lambda) = \sum_{\nu=1}^{M}\hat{m}_{i\nu}(1+s_\nu(\lambda))x_\nu(t) - \sum_{\nu\neq i}^{M}\hat{m}_{\nu i}(1+s_i(\lambda))x_i(t) + \sum_{j=1}^{M-1}g_{ij}(\boldsymbol{x}(t))s_j(\lambda). \tag{S55}$$

Using Eqs. (S52), (S49), the details of the reaction types, and the relation $\langle\Gamma_k(\tau)\Gamma_{k'}(\tau')\rangle = \delta_{kk'}\delta(\tau-\tau')$, we can write the correlations of the noise terms for $i,j = 1,\ldots,M-1$ as:

$$\langle\eta_i(t)\eta_j(t')\rangle = \frac{2\delta(t-t')}{N_{\rm cell}(t)}\left[g_{ij}(\boldsymbol{x}(t)) + \mathcal{O}(\hat{m},\boldsymbol{s}(\lambda))\right]. \tag{S56}$$

Note that we have used the fact that Dirac delta functions change under a rescaling of variables as $\delta(\tau-\tau') = \delta(t-t')d$. The correction terms not explicitly shown in the bracket of Eq. (S56) involve elements of the matrix $\hat{m}$ and selection coefficient vector $\boldsymbol{s}(\lambda)$. Assuming $\hat{m}_{i\mu}, |s_i(\lambda)| \ll 1$, these can be ignored relative to the leading contribution $g_{ij}(\boldsymbol{x}(\tau))$, and we will also approximate $N_{\rm cell}(t) \approx \overline{N}_{\rm cell}$. Finally, since the noise functions $\eta_i(t)$ have zero mean, Eq. (S56) represents the covariance of the noise, and hence can be used to define a diffusivity matrix:

$$\langle\eta_i(t)\eta_j(t')\rangle = 2D_{ij}(\boldsymbol{x}(t))\delta(t-t'), \qquad \text{where } D_{ij}(\boldsymbol{x}(t)) \approx \frac{g_{ij}(\boldsymbol{x}(t))}{\overline{N}_{\rm cell}}. \tag{S57}$$

Eqs. (S54)-(S55), together with the noise covariance result of Eq. (S57), constitute a system of Langevin equations for the genotype frequencies $x_i(t)$. Using the standard relation between Langevin and Fokker-Planck equations[6],

we know that the corresponding Fokker-Planck operator has the form of Eq. (S1), with diffusivity matrix $D_{ij}(\boldsymbol{x})$ defined through Eq. (S57), and the velocity function $v_i(\boldsymbol{x};\lambda)$ defined through Eq. (S55).

Comparing Eqs. (S55) and (S57) to the velocity and diffusivity definitions in our Wright-Fisher Fokker-Planck formalism, main text Eq. (6), we see that they match under the mapping:

$$m_{i\nu} = \hat{m}_{i\nu}(1+s_\nu(\lambda)) \quad \text{for}\, \nu \neq i, \qquad N = \frac{1}{2}\overline{N}_{\text{cell}} = \frac{K}{2}\left(1-\frac{d}{b_0(1-d)}\right). \tag{S58}$$

To derive this mapping, we have used the fact that the diagonal entries of m are defined as $m_{\alpha\alpha} \equiv -\sum_{\beta\neq\alpha} m_{\beta\alpha}$.

8 Potential future applications of CD driving in evolutionary systems

The CD driving examples in the main text and SI were based on anti-malarial drug resistance fitness data from Refs. 7,8. This experimental system is particularly attractive for our purposes, since it is the only one we are aware of to date for which a full "seascape" has been measured: the set of fitnesses for all possible combinations of a set of alleles of interest over a range of environmental conditions, i.e. drug dosages. However there are no reasons why similar experimental techniques cannot be applied to construct seascapes for other systems. For example, single-environment (single dosage) drug-resistance landscapes have been measured in *E. coli* for 16 antibiotics[9], and these could be expanded to include fitnesses across a range of drug doses along the same lines as in Ref. 7.

One incentive for measuring such seascapes would be applications of CD in the treatment of drug resistance in infectious disease and cancer. Let us consider, for instance, a pathogen and two different drugs that can potentially kill or hinder its growth. Prior to drug treatment the pathogen population might consist of a heterogeneous mixture of genetic variants. With exposure to one of the drugs, that mixture changes as the pathogen mutates, with resistant variants coming to dominate due to natural selection. However suppose that these variants resistant to the first drug are in turn extremely susceptible to the second drug. Thus by steering the evolution of the population by use of the first drug, we have created a situation where the second drug can be particularly effective, improving our chances of eradicating the disease using a two-stage sequence of drugs. Clearly for this to work the initial steering stage has to be accomplished quickly[10,11], before the overall population of resistant variants under the first drug rebounds sufficiently to endanger the patient. The aforementioned scenario is not merely hypothetical, but describes a phenomenon known as "collateral sensitivity" between the pathogen and the two drugs in question. It has been observed *in vitro* and in clinical isolates for a variety of infectious diseases, as well as cancer[12,13]. Theoretical and experimental work has demonstrated the prevalence of this phenomenon and the potential for carefully crafted drug sequences to exploit this relationship[14–19]. However a key question in applying such a strategy clinically is how long it would take for the "steering" stage under the first drug to be completed, achieving the desired target distribution of genetic variants ready for the application of the second drug[10,11,20,21]. Given the slow nature of evolution, having the first stage drag out into weeks or months would severely limit the clinical usefulness of this approach. CD speed-up could be useful in overcoming this challenge, by providing a dosage protocol to hasten arrival at the target distribution. This could be one of the key steps in translating this theoretical treatment idea into a clinically actionable therapeutic strategy.

There are also systems outside the context of evolutionary medicine where we expect this protocol to be useful despite the required seascape details. In the context of agriculture, CD driving could potentially be used to speed of up the process of crop breeding. As in *E. coli*, fitness landscapes have been successfully measured in maize[22,23], although not full seascapes. Note that these are technically trait-yield landscapes rather than true fitness landscapes, but that by performing artificial selection based on yield we can effectively use them as fitness landscapes. Of course, if we had sufficient information to use CD driving on a crop, we would also have sufficient information to use genetic engineering instead. It is unclear whether it would ever be advantageous to use CD driving over genetic engineering, but we suspect instances exist where it would be (perhaps due to CD driving's ability to act on an entire population at once, or due to anti-GMO sentiment). A similar argument holds for the potential applicability of CD driving to synthetic biology.

Measuring seascapes in additional systems would be labor-intensive, but tractable if it were the only obstacle standing between us and substantial advances in applications like those mentioned above. If there do exist cases where collecting such full seascapes may be prohibitively difficult, we can try the alternative strategy described in the Discussion section of the main text: quasi-adiabatic *in vitro* experimental trials where we vary external parameters (i.e. drug dosages) extremely gradually, and use sequencing at regular intervals to determine the quasi-equilibrium mean genotype fractions as a function of dosage. This would be generally faster than trying to determine the fitnesses of each genotype at each dosage level, and would provide enough information to construct our CD protocol through Eq. (4) of the main text.

9 Generality of the CD approach: deriving approximate protocols in additional landscape examples

For the empirical seascape used in the main text, describing the response of 16 genotypes to the anti-malarial drug pyrimethamine[7,8], we noted that two of those genotypes dominate at different times during the driving: initially at small concentrations the genotype 1110 ($i=15$) has the largest population, but is eventually overtaken by the wild-type genotype 1111 ($i=16$) at large concentrations as the drug is ramped up. This allowed a particularly simple approach to approximate the CD protocol of main text Eq. (4) by focusing on the the selection coefficient s_{15} that describes the relative fitness of 1110 versus 1111: at every time step we numerically find the drug concentration $\tilde{\lambda}(t)$ where the corresponding selection coefficient $s_{15}(\tilde{\lambda}(t))$ is closest to the value of the CD solution $\tilde{s}_{15}(\lambda(t),\dot{\lambda}(t))$ from Eq. (4), with the constraint that $\tilde{\lambda}(t)$ must be lower than a certain maximum cutoff $\tilde{\lambda}_{\max}$. This approximate CD dosage protocol $\tilde{\lambda}(t)$ is compared to the original dosage protocol $\lambda(t)$ in main text Fig. 3E for different cutoffs $\tilde{\lambda}_{\max}$. While we saw the resulting CD driving worked quite well for this example, how would we generalize this approach to more complex scenarios?

In this section we describe such a general approach for finding the set of experimental control parameters $\tilde{\lambda}(t)$ that best approximates main text Eq. (4) given an arbitrary seascape. We show how the two-genotype approximation described above is a special case of this general approach, and illustrate the latter through two additional driving examples: one using a modified version of the pyrimethamine seascape, and the other using the empirical seascape for cycloguanil, another anti-malarial drug[7]. In both of these cases driving takes us through three different regimes, with different genotypes dominating in each regime.

Let us start with Eq. (18) of the Methods,

$$\begin{aligned}\partial_t\overline{\boldsymbol{x}}(\lambda(t)) &= \boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\delta\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)\\ &= \boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\Big(\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)-\boldsymbol{s}(\lambda(t))\Big).\end{aligned} \tag{S59}$$

The CD solution $\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)$ to this equation is shown in main text Eq. (4). Now let us imagine we want to use the control parameters accessible to us and implement a protocol $\tilde{\lambda}(t)$ that approximates this CD solution. To do this, let us replace $\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)$ with $\boldsymbol{s}(\tilde{\lambda}(t))$ on the right-hand side of Eq. (S59). Subtracting the left and right-hand sides of the resulting equation from each other, and squaring the difference, gives us the following loss function:

$$L(\tilde{\lambda}(t)) = \big[\partial_t\overline{\boldsymbol{x}}(\lambda(t)) - \boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\big(\boldsymbol{s}\big(\tilde{\lambda}(t)\big)-\boldsymbol{s}(\lambda(t))\big)\big]^2. \tag{S60}$$

If we could find a $\tilde{\lambda}(t)$ such that $L(\tilde{\lambda}(t))=0$ at all t then we would have exactly implemented the CD solution of main text Eq. (4). However in practice the control parameters $\tilde{\lambda}(t)$ that we can externally manipulate may not allow this to be perfectly satisfied. Hence we do the next best thing and find $\tilde{\lambda}(t)$ that minimizes the loss function. Note that the loss function can also be rewritten in the following form, where we substitute in the right-hand side of Eq. (S59) for $\partial_t\overline{\boldsymbol{x}}(\lambda(t))$:

$$\begin{aligned}L(\tilde{\lambda}(t)) &= \Big[\boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\Big(\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)-\boldsymbol{s}(\lambda(t))\Big) - \boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\big(\boldsymbol{s}\big(\tilde{\lambda}(t)\big)-\boldsymbol{s}(\lambda(t))\big)\Big]^2\\ &= \Big[\boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)\Big(\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)-\boldsymbol{s}\big(\tilde{\lambda}(t)\big)\Big)\Big]^2.\end{aligned} \tag{S61}$$

In this form we can explicitly see that minimizing the loss function is the same as minimizing the difference between $\tilde{\boldsymbol{s}}\big(\lambda(t),\dot{\lambda}(t)\big)$ from main text Eq. (4) and the experimentally accessible selection coefficient trajectory $\boldsymbol{s}\big(\tilde{\lambda}(t)\big)$. This difference is weighted by the response matrix $\boldsymbol{g}\big(\overline{\boldsymbol{x}}(\lambda(t))\big)$, which governs how much each component of the selection coefficient contributes to the velocity term in the Fokker-Planck equation through Methods Eq. (6). This is not the only possible form of the loss function (one can imagine alternative ways of weighting the selection coefficient differences) but in the numerical examples we investigated this form gave us the best driving results.

Minimizing Eq. (S61) (or equivalently Eq. (S60)) to find $\tilde{\lambda}(t)$ is a general prescription for getting an approximate CD protocol, regardless of the system. How well that protocol will work (how close we can get to true CD driving) is of course a system-specific question, which will depend on the nature and number of control parameters that we can vary and the ranges they can explore (the dimensionality and extent of the space of possible $\tilde{\lambda}(t)$ curves). In general we expect that the more degrees of freedom we have for $\tilde{\lambda}(t)$, the closer we can approach the CD solution. Because simultaneously controlling a large number of parameters is experimentally more challenging, in practice one will seek out the smallest number that can still give a reasonable approximation to CD driving.

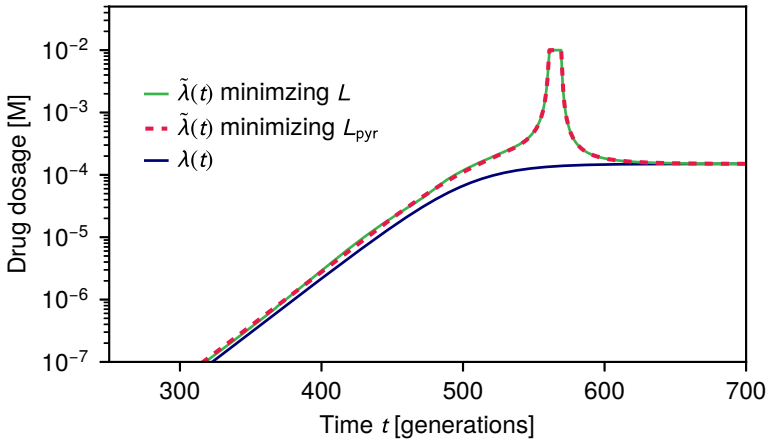

Figure S3. Comparison of the original drug dosage $\lambda(t)$ (blue curve) for the pyrimethamine seascape example of the main text versus the approximate CD dosage protocol $\tilde{\lambda}(t)$ calculated using two methods: minimizing the full loss function $L(\tilde{\lambda}(t))$ of Eq. (S61) (green curve) versus minimzing the simplified loss function $L_{\text{pyr}}(\tilde{\lambda}(t))$ of Eq. (S62) (dashed red curve).

To get some intuition into this approximation scheme, let us consider several examples. For the pyrimethamine seascape used in the main text, minimizing Eq. (S61) turns out to be nearly the same as minimizing the simpler alternative loss function that involves only the $i = 15$ selection coefficient component:

$$L_{\text{pyr}}(\tilde{\lambda}(t)) = \left(\tilde{s}_{15}\big(\lambda(t), \dot{\lambda}(t)\big) - s_{15}\big(\tilde{\lambda}(t)\big)\right)^2. \tag{S62}$$

This is because $i = 15$ difference makes the major contribution to the right-hand side of Eq. (S61), given that the 1110 and 1111 genotypes dominate the population at different stages of the driving. Fig S3 shows the drug protocol $\tilde{\lambda}(t)$ calculated by minimizing Eq. (S61) versus that found from minimizing the simpler loss function of Eq. (S62). The two are nearly identical throughout the entire time range. As can be seen in main text Fig. 3, for this system a single control parameter (pyrimethamine drug dosage) is sufficient to get an excellent approximation to the CD protocol over the entire time range. The genotype distributions under the protocol, $p(\boldsymbol{x}, t)$, are kept close to the IE target distributions $\rho(\boldsymbol{x}; \lambda(t))$, measured via their KL divergence (main text Fig. 3F). Since the distributions are narrow for the large N case, an individual simulation trajectory (solid curves in main text Fig. 3C) will closely follow the mean IE genotype frequencies $\overline{\boldsymbol{x}}\big(\lambda(t)\big)$ (dashed curves in main text Fig. 3C). The excellent performance of the approximate CD protocol in this case likely stems from the fact that the loss function is dominated by a single component ($i = 15$), and varying a one-dimensional control parameter allows us to effectively minimize this component.

Let us now make the situation more complicated in the following way. Take the same pyrimethamine seascape, based on the empirical data of Ref. 7, and make a single parameter alteration: increase the base growth rate (at zero drug concentration) of genotype 0110 by 5%. Everything else in the system remains the same. Under the original drug protocol (blue curve in Fig. S4C) the mean IE genotype fractions (dashed curves in Fig. S4A,B) show a more complex behavior over time: 0110 dominates at small times, followed by a period of 1110 domination, until at large times 1111 takes the lead. As time is varied there are now multiple components on the right-hand side of Eq. (S61) that shift in their relative significance. The result of minimizing $L(\tilde{\lambda}(t))$ is the approximate CD protocol $\tilde{\lambda}(t)$ shown in green in Fig. S4C. This now has two dosage peaks, one around the 0110-1110 transition and the other around the 1110-1111 transition. Under the original protocol the simulation trajectory lags behind the IE expectation (Fig. S4A), but with the approximate CD driving this lag is largely eliminated (Fig. S4B). The overall reduction in lag time (quantified as described in Methods Sec. 4.5.4) is $\Delta t = 1128$ generations, comparable to the reduction achieved in the main text example. Note however that the agreement between simulation and IE curves during the 1110-dominated regime is only partial. Clearly, manipulating a single control parameter (pyrimethamine dosage) is not sufficient to achieve a close enough approximation to the CD protocol to get fine-grained control in this regime. The drug dosage peak near generation 400 is sufficient to drive the interchange between 0110 and 1110, but as a side-effect causes the overshoot of 1111 with respect to its IE expectation. However this overshoot is

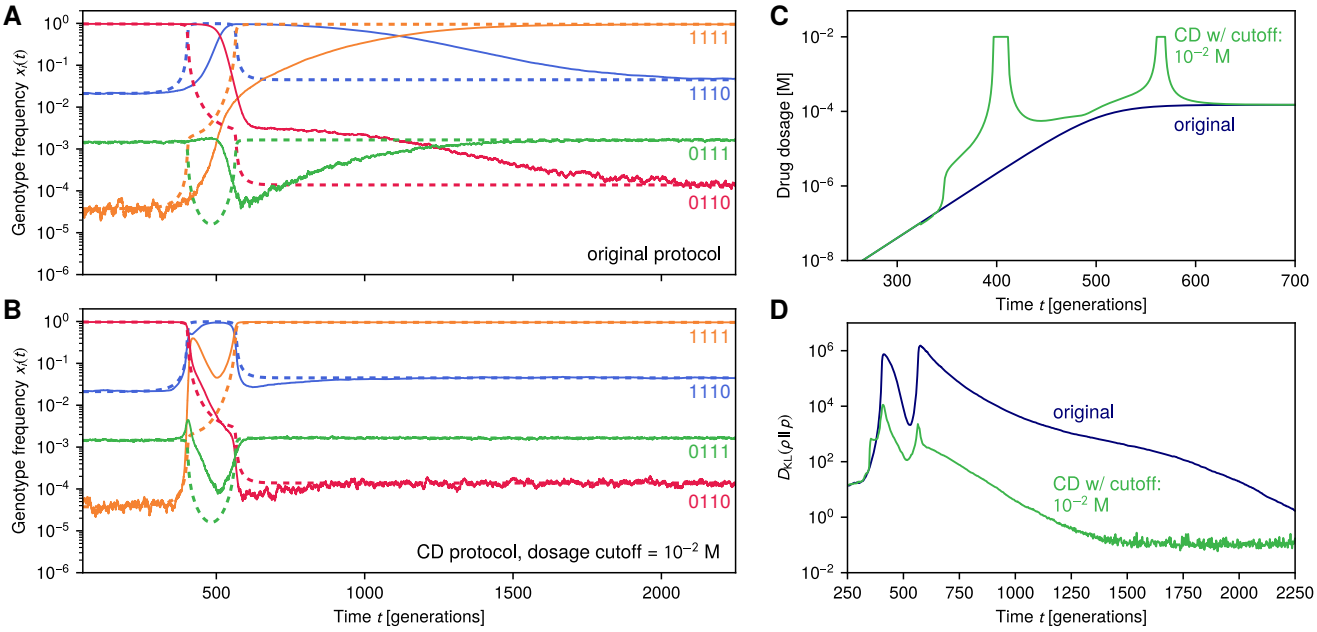

Figure S4. CD driving for an altered 16-genotype pyrimethamine seascape. This is the same seascape as in main text Fig. 3, using the experimental data of Ref. 7, except that genotype 0110 has been modified to have a 5% larger base growth rate under no drug conditions. **(A,B)** Sample simulation trajectories (solid lines) versus IE expectation (dashed lines) for the fraction of 4 representative genotypes without **(A)** and with **(B)** CD driving. The CD driving is implemented approximately through the drug dosage protocol (green curve) shown in panel (**C**) with cutoff 10^{-2} M. The original protocol (blue curve) is shown for comparison. **(D)** Kullback-Leibler divergence between actual and IE distributions versus time, with and without CD driving.

resolved quickly as 1111 approaches its IE curve, and the subsequent driving pushes the system toward the final IE values with little delay. Thus if one were interested only in getting to the final target distribution of genotypes quickly (likely the most common real-world application), this approximate CD protocol would suffice. On the other hand there could be situations where following the exact IE path through genotype space at all times was important, for example if we wanted to steer the system away from problematic intermediate genotypes in some evolutionary medicine scenario. In this case we could benefit from having additional control knobs. For example the possibility of using a second drug alongside pyrimethamine might allow for a better CD approximation.

Fig. S5 tells a qualitatively similar story to Fig. S4, but using the seascape from cycloguanil, a different anti-malarial drug. The cycloguanil data was also experimentally measured in Ref. 7, and we do not make any modifications to the empirical values. We choose an original dosage protocol $\lambda(t)$ with the same form as before (Methods Eq. (22)), but with saturating dosage parameter $a = 1.1 \times 10^{-5}$ M. Under this protocol, we see three regimes during the driving, but in a different combination from the previous example: the dominance shifts from 1110 to 0110 to 0111 as the drug is increased. Again the approximate CD protocol involves two dosage peaks, and gets us to the final target distribution with less lag than the original protocol ($\Delta t = 373$ generations). Though we currently only have data in this system for one drug applied at any given time, we see that despite this limitation we can still do a fairly decent job of eliminating lag, whether that drug be pyrimethamine or cycloguanil. The imperfect control during the intermediate regime could potentially be resolved by a protocol that involves both drugs applied simultaneously (or another external intervention). Thus, while our theory provides the basic framework to understand CD driving in such evolutionary systems, there are a variety of pragmatic questions about balancing driving accuracy versus the complexity of the control protocol that will provide interesting subjects for future work.

References

1. Born, M. & Fock, V. Beweis des adiabatensatzes. *Z. Phys.* **51**, 165–180 (1928).
2. Griffiths, D. J. & Schroeter, D. F. *Introduction to quantum mechanics* (Cambridge University Press, 2018).
3. Risken, H. *The Fokker-Planck Equation* (Springer, 1996).

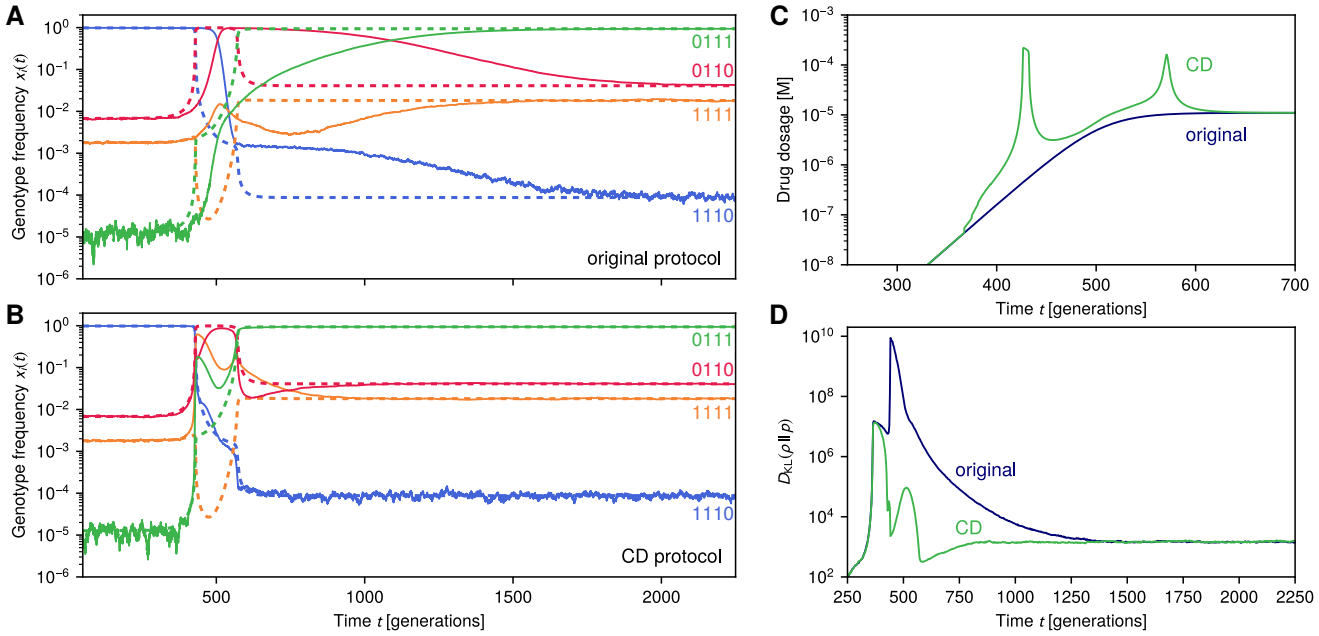

Figure S5. CD driving for a 16-genotype cycloguanil seascape. This is the same 16-genotype system as in the examples of main text Fig. 3 and Fig. S4, except using the anti-malarial drug cycloguanil instead of pyrimethamine. The seascape is based on the experimental data of Ref. 7, without any modifications. **(A,B)** Sample simulation trajectories (solid lines) versus IE expectation (dashed lines) for the fraction of 4 representative genotypes without **(A)** and with **(B)** CD driving. The CD driving is implemented approximately through the drug dosage protocol (green curve) shown in panel (**C**). The original protocol (blue curve) is shown for comparison. **(D)** Kullback-Leibler divergence between actual and IE distributions versus time, with and without CD driving.

4. Ryter, D. On the eigenfunctions of the Fokker-Planck operator and of its adjoint. *Phys. A* **142**, 103–121 (1987).
5. Gillespie, D. T. The chemical Langevin equation. *J. Chem. Phys.* **113**, 297–306 (2000).
6. Gillespie, D. T. The multivariate Langevin and Fokker-Planck equations. *Am. J. Phys.* **64**, 1246–1257 (1996).
7. Ogbunugafor, C. B., Wylie, C. S., Diakite, I., Weinreich, D. M. & Hartl, D. L. Adaptive landscape by environment interactions dictate evolutionary dynamics in models of drug resistance. *PLoS Comp. Biol.* **12**, e1004710 (2016).
8. Brown, K. M. *et al.* Compensatory mutations restore fitness during the evolution of dihydrofolate reductase. *Mol. Biol. Evol.* **27**, 2682–2690.
9. Mira, P. M. *et al.* Rational design of antibiotic treatment plans: a treatment strategy for managing evolution and reversing resistance. *PloS One* **10**, e0122283 (2015).
10. Basanta, D., Gatenby, R. A. & Anderson, A. R. Exploiting evolution to treat drug resistance: combination therapy and the double bind. *Mol. Pharm.* **9**, 914–921 (2012).
11. Gerlee, P. & Altrock, P. M. Extinction rates in tumour public goods games. *J. R. Soc. Interface* **14**, 20170342 (2017).
12. Imamovic, L. & Sommer, M. O. Use of collateral sensitivity networks to design drug cycling protocols that avoid resistance development. *Sci. Transl. Med.* **5**, 204ra132–204ra132 (2013).
13. Dhawan, A. *et al.* Collateral sensitivity networks reveal evolutionary instability and novel treatment strategies in alk mutated non-small cell lung cancer. *Sci. Rep.* **7**, 1–9 (2017).
14. Nichol, D. *et al.* Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS Comp. Biol.* **11**, e1004493 (2015).
15. Zhao, B. *et al.* Exploiting temporal collateral sensitivity in tumor clonal evolution. *Cell* **165**, 234–246 (2016).
16. Barbosa, C., Beardmore, R., Schulenburg, H. & Jansen, G. Antibiotic combination efficacy (ace) networks for a pseudomonas aeruginosa model. *PLoS Biol.* **16**, e2004356 (2018).

17. Nichol, D. *et al.* Antibiotic collateral sensitivity is contingent on the repeatability of evolution. *Nat. communications* **10**, 334 (2019).

18. Maltas, J. & Wood, K. B. Pervasive and diverse collateral sensitivity profiles inform optimal strategies to limit antibiotic resistance. *PLoS Biol.* **17** (2019).

19. Acar, A. *et al.* Exploiting evolutionary steering to induce collateral drug sensitivity in cancer. *Nat. Commun.* **11**, 1–14 (2020).

20. Kaznatcheev, A. Computational complexity as an ultimate constraint on evolution. *Genetics* **212**, 245–265 (2019).

21. Barbosa, C., Roemhild, R., Rosenstiel, P. & Schulenburg, H. Evolutionary stability of collateral sensitivity to antibiotics in the model pathogen pseudomonas aeruginosa. *eLife* **8**, e51481 (2019).

22. Chenu, K. *et al.* Simulating the Yield Impacts of Organ-Level Quantitative Trait Loci Associated With Drought Response in Maize: A "Gene-to-Phenotype" Modeling Approach. *Genetics* **183**, 1507–1523 (2009).

23. Messina, C. D., Podlich, D., Dong, Z., Samples, M. & Cooper, M. Yield–trait performance landscapes: from theory to application in breeding maize for drought tolerance. *J. Exp. Bot.* **62**, 855–868 (2011).