Check for updates

Controlling the speed and trajectory of evolution with counterdiabatic driving

Shamreen Iram [1,7], Emily Dolson [2,7], Joshua Chiel [1,7], Julia Pelesko[1,2], Nikhil Krishnan[2,3],
Özenç Güngör[1], Benjamin Kuznets-Speck[1,4], Sebastian Deffner[5], Efe Ilker [6], Jacob G. Scott [1,2,3] ✉ and
Michael Hinczewski [1] ✉

The pace and unpredictability of evolution are critically relevant in a variety of modern challenges, such as 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 modelling of evolution using fitness landscapes, allowing a degree of prediction for future evolutionary histories. Yet fine-grained control of the speed and 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 behaviour of quantum states for applications like quantum computing and manipulating ultracold atoms. Implementing these ideas for the first time in a biological context, we show how a set of external control parameters (that is, varying drug concentrations and types, temperature and 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 pre-dates our understanding of evolution itself. Recent years have seen growing research efforts towards 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[4,5]. New approaches that have arisen in response include prospective therapies that steer evolution of pathogens towards maximized drug sensitivity[6,7], typically requiring multiple rounds of selective pressures and subsequent evolution under them. Because 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[8,9], in many cases requiring that a system remains in an instantaneous ground state of a time-varying Hamiltonian in applications like cold atom transport[10] and quantum adiabatic computation[11]. 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'[12–17]. 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[18–20]. 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 that a classical stochastic system remains in an instantaneous equilibrium state as external control parameters are changed[21,22]. Extending CD driving ideas to the classical realm has already led to proof-of-concept demonstrations of accelerated equilibration in optical tweezer[23] and atomic force microscope[24] experimental frameworks, and is closely related to optimal, finite-time control problems in stochastic thermodynamics[25,26].

Here we demonstrate the first biological application of CD driving, by using it to control the distribution of genotypes in a Wright–Fisher (WF) model[27] 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 (ABM) of evolving unicellular populations, focusing on a system where 16 possible genotypes compete via a drug dose-dependent fitness landscape derived from experimental measurements.

[1]Department of Physics, Case Western Reserve University, Cleveland, OH, USA. [2]Translational Hematology Oncology Research, Cleveland Clinic, Cleveland, OH, USA. [3]Case Western Reserve University School of Medicine, Cleveland, OH, USA. [4]Biophysics Graduate Group, University of California, Berkeley, CA, USA. [5]Department of Physics, University of Maryland, Baltimore County, Baltimore, MD, USA. [6]Physico-Chimie Curie UMR 168, Institut Curie, PSL Research University, Paris, France. [7]These authors contributed equally: Shamreen Iram, Emily Dolson, Joshua Chiel. ✉e-mail: scottj10@ccf.org; michael.hinczewski@case.edu

Fig. 1 | Schematic overview of counterdiabatic driving for evolutionary systems. a, The trajectory of the selection coefficient vector $\mathbf{s}(\lambda(t))$ as a function of a time-varying control parameter $\lambda(t)$, depicted in the bottom panel, for a hypothetical system of three genotypes ($M = 3$). 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 colour gradient. **b**, The instantaneous equilibrium (IE) distribution of genotypes $\rho(\mathbf{x}; \lambda(t))$ for the three highlighted values of the control parameter from **a**. These distributions are probability densities on the two-dimensional (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 $\bar{\mathbf{x}}(\lambda(t))$. **c**, For driving over finite times, the actual distribution of genotypes $p(\mathbf{x},t)$ will generally lag behind the IE distribution while the control parameter is changing. Thus, at t_f, the distribution $p(\mathbf{x},t_f)$ is still far from $\rho(\mathbf{x};\lambda_f)$, and will only catch up with it at times $t \gg t_f$ as the system re-equilibrates.

Theory

Evolutionary model. We develop our CD driving theory in the framework of a WF 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. Because $\sum_{i=1}^{M} x_i = 1$, we can describe the state of the system through $M - 1$ independent values of x_i, or equivalently through a frequency vector $\mathbf{x} = (x_1, \ldots, x_{M-1})$. 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 $\mathbf{s}$. 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)[6] or other perturbations in fitness secondary to microenvironmental change (for example, 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, $\mathbf{s}(\lambda(t))$, shown schematically in Fig. 1a. Our population thus evolves under a time-dependent fitness landscape, or a so-called 'seascape'[28]. Note that all time variables, unless otherwise noted, are taken to be in units of WF 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.) 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, although this assumption can be relaxed.

Driving the genotype frequency distribution. Given the system described above, we focus on $p(\mathbf{x},t)$, the probability to find genotype frequencies $\mathbf{x}$ 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:

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

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

Table 1 | Analogies between quantum physics and evolutionary dynamics

Quantum physics	Evolutionary dynamics
1. Wavefunction: describes the state of a quantum system. For a simple quantum particle in a spatial region described by coordinates $\mathbf{x}$, this is a function $\psi(\mathbf{x},t)$ whose squared amplitude $\|\psi(\mathbf{x},t)\|^2$ is the probability density of finding the particle at x at time t.	1. Genotype probability distribution: the distribution $p(\mathbf{x},t)$ of genetic variants (genotypes) in a population of organisms at time t, where $\mathbf{x}$ is a vector of genotype fractions.
2. Hamiltonian operator: a differential operator $H(\lambda(t))$ depending on control parameters $\lambda(t)$, defined below. It describes how the wavefunction changes in time through the time-dependent Schrödinger equation, $i\hbar\partial_t H\psi(\mathbf{x},t) = H(\lambda(t))\psi(\mathbf{x},t)$, where $\hbar$ is the reduced Planck constant. $H(\lambda(t))$ involves terms that correspond to the kinetic and potential energies of the quantum particle.	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 $p(\mathbf{x},t)$ changes in time through the Fokker–Planck equation, $\partial_t p(\mathbf{x},t) = \mathcal{L}(\lambda(t))\,p(\mathbf{x},t)$. The full form of $\mathcal{L}(\lambda(t))$ (equations (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 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.	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. Ground state: the lowest energy state of a quantum system. In general, for a Hamiltonian $H(\lambda)$ and given parameter values λ, the energy states (labeled by $n = 0, 1, 2, ...$) correspond to solutions of the time-independent Schrödinger equation: $H(\lambda)\psi_n(\mathbf{x};\lambda) = E_n(\lambda)\psi_n(\mathbf{x};\lambda)$. Here, $E_n(\lambda)$ and $\psi_n(\mathbf{x};\lambda)$ are the energy and wavefunction, respectively, of the nth state. The energies $E_0 < E_1 < ...$ and the ground state corresponds to $n = 0$. If λ is fixed, a system whose wavefunction is $\psi_n(\mathbf{x};\lambda)$ will be stationary, with its wavefunction not changing in time.	4. Equilibrium state: for a given set of parameter values λ, this is the genotype probability $\rho(\mathbf{x};\lambda)$, 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 $\psi_n(\mathbf{x};\lambda)$ and eigenvalues $-\kappa_n(\lambda) \le 0$ for $n = 0, 1, 2, ...$ defined through the following equation: $\mathcal{L}(\lambda)\,\psi_n(\mathbf{x};\lambda) = -\kappa_n(\lambda)\psi_n(\mathbf{x};\lambda)$. The equilibrium state corresponds to $n = 0$, with eigenvalue $-\kappa_0(\lambda) = 0$ and $\rho(\mathbf{x};\lambda) \equiv \psi_0(\mathbf{x};\lambda)$.
5. Adiabatic theorem: if we start in the nth energy state, $\psi(\mathbf{x},t=0) = \psi_n(\mathbf{x};\lambda(0))$ 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 the nth energy state corresponding to the instantaneous value of the parameters, $\psi(\mathbf{x},t) = \psi_n(\mathbf{x};\lambda(t))$. This is true as long as there is always a non-zero difference between $E_n(\lambda(t))$ and any $E_m(\lambda(t))$ for $m \neq n$ at all t.	5. Adiabatic theorem: if we start at equilibrium $p(\mathbf{x},t=0) = \rho(\mathbf{x};\lambda(0))$ for some initial control parameters $\lambda(0)$, and then vary $\lambda(t)$ infinitesimally slowly (adiabatically), the theorem (derived in the Supplementary Information) states that at later times we will always remain in the equilibrium state corresponding to the instantaneous value of the parameters, $p(\mathbf{x},t) = \rho(\mathbf{x};\lambda(t))$.

A summary of the connections between quantum and evolutionary concepts used in our theory is shown, where each numbered item in the quantum column on the left has its analogue in the evolutionary column on the right.

to the Schrödinger equation (though not vice versa)[29–31]. For a particular value of the control parameter λ, the analogue of the quantum ground-state wavefunction is the eigenfunction $\rho(\mathbf{x};\lambda)$ with eigenvalue zero, the solution of the equation

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

In the evolutionary context, $\rho(\mathbf{x};\lambda)$ 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 $p(\mathbf{x},t)$ obeying equation (1) will approach $\rho(\mathbf{x};\lambda)$ in the limit $t \to \infty$.

Consider the following control protocol, where we start at one control parameter value, $\lambda(t) = \lambda_0$ for $t \le t_0$, and finish at another value, $\lambda(t) = \lambda_f$ for $t \ge t_f$, with some arbitrary driving function $\lambda(t)$ in the interval $t_0 < t < t_f$. We assume the system starts in one equilibrium distribution, $p(\mathbf{x},t_0) = \rho(\mathbf{x};\lambda_0)$, and we know that it will eventually end at a new equilibrium, $p(\mathbf{x},t) \to \rho(\mathbf{x};\lambda_f)$ 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, $p(\mathbf{x},t) = \rho(\mathbf{x};\lambda(t))$ for all t. This result, derived in the Supplementary Information, is the analogue of the quantum adiabatic theorem[32] 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). Figure 1b presents schematic snapshots of $\rho(\mathbf{x};\lambda(t))$ 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: $p(\mathbf{x},t) \neq \rho(\mathbf{x};\lambda(t))$ 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 $p(\mathbf{x},t)$ moving towards but not able to catch up with $\rho(\mathbf{x};\lambda(t))$, as illustrated in Fig. 1c. For $t > t_f$, once $\lambda(t)$ stops changing, the system will eventually settle into equilibrium at $\rho(\mathbf{x};\lambda_f)$ in the long-time limit.

Control and counterdiabatic driving. This lag can be an obstacle if one wants to control the evolution of the system over finite time intervals. Because 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, because in practice we could imagine

the interval $t_0 < t < t_f$ as just one step of a multi-stage protocol involving distinct interventions (that is, 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 Supplementary Section 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, $p(\mathbf{x}, t_0) = \rho(\mathbf{x}; \lambda_0)$, to another, $p(\mathbf{x}, t_f) = \rho(\mathbf{x}; \lambda_f)$, over a finite time $t_f - t_0$.

In the context of quantum adiabatic computing[11], the typical focus is on the initial ground state (which has to be easy to realize experimentally) and the final ground state (because 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, $\mathbf{s}(\lambda_0)$, to a final fitness landscape, $\mathbf{s}(\lambda_f)$, 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. However, among the remaining set of realizable trajectories, some may be more desirable than others (that is, have different evolutionary consequences[33] or trade-offs[34]). Each trajectory defines a continuous sequence of IE distributions $\rho(\mathbf{x}; \lambda(t))$, and for each distribution there is a mean genotype frequency $\overline{\mathbf{x}}(\lambda(t))$, 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 $\overline{\mathbf{x}}(\lambda(t))$ for $t_0 < t < t_f$ stays far away from certain regions of the genotype frequency space. This, in turn, restricts the $\mathbf{s}(\lambda(t))$ 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 $p(\mathbf{x}, t) = \rho(\mathbf{x}; \lambda(t))$ for some chosen control protocol $\lambda(t)$ in $t_0 < t < t_f$. 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 guide our system. Clearly, we will not achieve success by just directly implementing $\lambda(t)$, because $p(\mathbf{x}, t)$, obeying equation (1), will generally lag behind $\rho(\mathbf{x}; \lambda(t))$[35]. 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[18–20]. 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 $\mathcal{L}(\lambda(t))$ in equation (1) with a different operator $\widetilde{\mathcal{L}}(\lambda(t), \dot{\lambda}(t))$, which depends on both $\lambda(t)$ and its time derivative $\dot{\lambda}(t) \equiv \mathrm{d}\lambda(t)/\mathrm{d}t$. This CD operator satisfies

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

Thus, by construction, $p(\mathbf{x}, t) = \rho(\mathbf{x}; \lambda(t))$ is a solution to the Fokker–Planck equation with the new operator. Additionally, to be consistent with the slow adiabatic driving limit discussed above, $\widetilde{\mathcal{L}}(\lambda(t), 0) = \mathcal{L}(\lambda(t))$, 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 $\widetilde{\mathcal{L}}(\lambda(t), \dot{\lambda}(t))$ in this way is the easy part: figuring out how to implement a new control protocol to realize $\widetilde{\mathcal{L}}(\lambda(t), \dot{\lambda}(t))$ 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 $\mathbf{s}(\lambda(t))$ with a frequency-dependent version, $\tilde{\mathbf{s}}(\mathbf{x}; \lambda(t), \dot{\lambda}(t))$. 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, $\tilde{\mathbf{s}}(\mathbf{x}; \lambda(t), \dot{\lambda}(t)) \approx \tilde{\mathbf{s}}(\lambda(t), \dot{\lambda}(t))$. 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 non-zero mutation rates where $\alpha \neq \beta$, then this corresponds to $\mu N \gg 1$, $N \gg 1$ (refs. [36–39]). In this regime multiple genotypes can generally coexist in the population at equilibrium (although one may be quite dominant), which is particularly relevant for pathogenic populations, especially ones spreading through space[40–42]. Remarkably, there is a simple analytical expression that provides an excellent approximation to $\tilde{\mathbf{s}}(\lambda(t), \dot{\lambda}(t))$ in this case:

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

where $\overline{x}_M(\lambda(t)) \equiv 1 - \sum_{i=1}^{M-1} \overline{x}_i(\lambda(t))$. We see that the new selection coefficient protocol is defined through the target mean genotype frequency trajectory $\overline{\mathbf{x}}(\lambda(t))$, and reduces to the original protocol when $\dot{\lambda}(t) \to 0$. Moreover, as we show in the examples below for specific systems, equation (4) can at least in certain cases be implemented through physically realistic manipulations of the environment, like time-varying drug dosages. Although 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 modelled as a sequence of mutant fixations, one can also formulate a CD theory based on a discrete Markov state description[43].

Results

Two genotypes. The simplest example of our CD theory is for a two-genotype ($M = 2$) system, where the dynamics are one-dimensional (1D) and described by a single frequency x_1 and selection coefficient $s_1(\lambda(t))$. As shown in the Methods, equation (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 favours genotype 1. Figure 2a shows $p(x_1, t)$ from a numerical solution of equation (1) using this protocol, compared against the IE distribution $\rho(x_1; \lambda(t))$, solved using equation (2), at three time snapshots. To validate the Fokker–Planck approach, we also designed an ABM, as described in the Methods, 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 (see Methods), one can directly compare the $p(x_1, t)$ results from the ABM simulations (circles) to the Fokker–Planck numerical solution of equation (1) (curves), which show excellent agreement. In the absence of CD driving, as expected, $p(x_1, t)$ lags behind $\rho(x_1; \lambda(t))$, with the latter shifting rapidly to larger x_1 frequencies as the fitness of genotype 1 increases.

To eliminate this lag, we implemented the alternative selection coefficient trajectory of equation (20). Figure 2b shows a comparison between $\tilde{s}_1(\lambda(t), \dot{\lambda}(t))$ and the original $s_1(\lambda(t))$. 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 $\rho(x_1; \lambda(t))$. Figure 2c shows the same snapshots as in Fig. 2a, but now with CD driving: we see that the actual and IE distributions nearly perfectly overlap at all times. To quantify the effectiveness of the CD protocol, we measured the degree of overlap through the Kullback–Leibler (KL) divergence[35,44], defined for any two probability distributions $p(\mathbf{x})$ and $q(\mathbf{x})$ as $D_{\mathrm{KL}}(p||q) = \int \mathrm{d}\mathbf{x} \; p(\mathbf{x}) \log_2(p(\mathbf{x})/q(\mathbf{x}))$. Expressed in bits, the KL divergence is always ≥ 0, and equals 0 for identical distributions. Figure 2d shows $D_{\mathrm{KL}}(\rho||p)$ for both the original and

Fig. 2 | Counterdiabatic driving results for two genotypes. a,c, We plot three time snapshots of the actual probability distribution $p(x_1,t)$ versus the IE distribution $\rho(x_1;\lambda(t))$ 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 equation (1) for $p(x_1,t)$ and 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 equation (20). The three snapshot times (50, 100 and 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.

CD protocols, with the latter dramatically reducing the divergence across the time interval of driving.

Multiple genotypes via agent-based modelling. 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 (four alleles), with selection coefficients based on a well-characterized experimental system: the fitness effects of the antimalarial 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 labelled 'original' in Fig. 3e). With our choice of simulation parameters (see Methods), 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. Figure 3a,b shows the results in the absence of CD driving, with each genotype labelled by a four-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 1,500 generations to resolve. This is quantified in the KL divergence $D_{KL}(\rho\|p)$ in Fig. 3f, which rapidly increases by five 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 2,000. 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 15D space, as well as the fact that distributions estimated from an ensemble of simulations (1,000 independent 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 equation (4). However this involves guiding the system along a fitness trajectory in a 15D 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 that 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} that 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 equation (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

Fig. 3 | CD driving eliminates evolutionary lag in a 16-genotype simulation. a,d, Three snapshots of the evolving agent-based population model without (**a**) and with (**d**) CD driving. Each of the three 16-genotype (four 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 (**e**) with a cutoff of 10^{-2} M. **b,c,** Corresponding sample simulation trajectories (solid lines) versus IE expectation (dashed lines) for the fraction of four representative genotypes without (**b**) and with (**c**) CD driving. The latter significantly reduces the nearly 1,500-generation lag. **e,** The original drug protocol versus the CD protocol with different possible dosage cutoffs. **f,** KL divergence between the 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.

seek to achieve the closest possible protocol to the one described by equation (4), given the experimental constraints. In Supplementary Section 9 and Extended Data Figs. 1 and 2 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.

Figure 3e shows CD drug protocols with three different cutoffs: 1×10^{-2}, 1×10^{-3} and 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 the Methods. For a cutoff of 10^{-2} M the lag is reduced

by $\Delta t = 1{,}210$ generations. Notably, although 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 equation (4) exactly, we can still arrive at the target distribution faster through an approximate CD protocol.

Discussion

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 is discussed in more detail in Supplementary Section 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 1D 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 $\mathbf{s}(\boldsymbol{\lambda}(t))$, where $\boldsymbol{\lambda}(t) = (\lambda_1(t), \lambda_2(t), \ldots)$ is a multidimensional vector of control parameters at time t: $\lambda_1(t)$ is the concentration of one drug, $\lambda_2(t)$ is the concentration of another drug (or nutrient), and so on. More generally, one could explore how fundamental differences among fitness landscapes (that is, the difficulties in reaching local optima in so-called 'hard' landscapes[45]) influence the types of intervention needed to achieve driving and their effectiveness.

Even given accurately measured fitnesses, one might be hampered by imperfect estimation of other system parameters, such as 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 Supplementary Information for details). Although 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 equation (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. Because the prescription is dominated by the most frequently observed genotypes, this approach is well suited for cases where we do not have complete fitness information about all possible mutations, opening up CD driving to an even wider range of systems.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41567-020-0989-3.

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 Comput. 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. *Antimicrobial Resistance: Global Report on Surveillance* (WHO, 2014).
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. Nichol, D. et al. Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS Comput. Biol.* **11**, e1004493 (2015).
7. Maltas, J. & Wood, K. B. Pervasive and diverse collateral sensitivity profiles inform optimal strategies to limit antibiotic resistance. *PLoS Biol.* **17**, e3000515 (2019).
8. Bason, M. G. et al. High-fidelity quantum driving. *Nat. Phys.* **8**, 147–152 (2012).
9. Zhou, B. B. et al. Accelerated quantum control using superadiabatic dynamics in a solid-state lambda system. *Nat. Phys.* **13**, 330–334 (2017).
10. Walther, A. et al. Controlling fast transport of cold trapped ions. *Phys. Rev. Lett.* **109**, 080501 (2012).
11. Farhi, E. et al. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **292**, 472–475 (2001).
12. Torrontegui, E. et al. Shortcuts to adiabaticity. *Adv. At. Mol. Opt. Phys* **62**, 117–169 (2013).
13. Deffner, S., Jarzynski, C. & del Campo, A. Classical and quantum shortcuts to adiabaticity for scale-invariant driving. *Phys. Rev. X* **4**, 021013 (2014).
14. Deffner, S. Shortcuts to adiabaticity: suppression of pair production in driven Dirac dynamics. *New J. Phys.* **18**, 012001 (2015).
15. Acconcia, T. V., Bonança, M. V. S. & Deffner, S. Shortcuts to adiabaticity from linear response theory. *Phys. Rev. E* **92**, 042148 (2015).
16. Campbell, S. & Deffner, S. Trade-off between speed and cost in shortcuts to adiabaticity. *Phys. Rev. Lett.* **118**, 100601 (2017).
17. Guéry-Odelin, D. et al. Shortcuts to adiabaticity: concepts, methods and applications. *Rev. Mod. Phys.* **91**, 045001 (2019).
18. Demirplak, M. & Rice, S. A. Adiabatic population transfer with control fields. *J. Phys. Chem. A* **107**, 9937–9945 (2003).
19. Demirplak, M. & Rice, S. A. Assisted adiabatic passage revisited. *J. Phys. Chem. B* **109**, 6838–6844 (2005).
20. Berry, M. V. Transitionless quantum driving. *J. Phys. A Math. Theory* **42**, 365303 (2009).
21. Patra, A. & Jarzynski, C. Shortcuts to adiabaticity using flow fields. *New J. Phys.* **19**, 125009 (2017).
22. Li, G., Quan, H. & Tu, Z. Shortcuts to isothermality and nonequilibrium work relations. *Phys. Rev. E* **96**, 012144 (2017).
23. Martínez, I. A., Petrosyan, A., Guéry-Odelin, D., Trizac, E. & Ciliberto, S. Engineered swift equilibration of a Brownian particle. *Nat. Phys.* **12**, 843–846 (2016).
24. Le Cunuder, A. et al. Fast equilibrium switch of a micro mechanical oscillator. *Appl. Phys. Lett.* **109**, 113502 (2016).

25. Schmiedl, T. & Seifert, U. Optimal finite-time processes in stochastic thermodynamics. *Phys. Rev. Lett.* **98**, 108301 (2007).

26. 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).

27. Wright, S. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proc. Sixth Int. Congress on Genetics* Vol. 1, 356–366 (Univ. Chicago Press, 1932).

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

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

30. Van Kampen, N. G. *Stochastic Processes in Physics and Chemistry* (Elsevier, 1992).

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

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

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

34. Li, Y., Petrov, D. A. & Sherlock, G. Single nucleotide mapping of trait space reveals Pareto fronts that constrain adaptation. *Nat. Ecol. Evol* **3**, 1539–1551 (2019).

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

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

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

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

39. Sniegowski, P. D. & Gerrish, P. J. Beneficial mutations and the dynamics of adaptation in asexual populations. *Philos. Trans. R. Soc. B* **365**, 1255–1263 (2010).

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

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

42. Krishnan, N. & Scott, J. G. Range expansion shifts clonal interference patterns in evolving populations. Preprint at https://www.biorxiv.org/content/10.1101/794867v2 (2019).

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

44. Kullback, S. & Leibler, R. A. On information and sufficiency. *Ann. Math. Stat.* **22**, 79–86 (1951).

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

Methods

Fokker–Planck description of the WF evolutionary model. The underlying evolutionary dynamics of our model are based on the canonical haploid WF model with mutation and selection, and we adopt the formalism of recent approaches[28,46] that generalize Kimura's original two-allele diffusion theory[47] 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, ABMs 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[28,46,47] for evolutionary population dynamics is the assumption that genotype frequencies change only by small amounts in each generation. We can thus take the genotype frequency vector $\mathbf{x}$ to be a continuous variable that follows a stochastic trajectory. The key quantities describing these stochastic dynamics are the lowest-order moments of $\delta\mathbf{x}$, 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 $v_i(\mathbf{x};\lambda(t)) \equiv \langle \delta x_i \rangle$. Note that in general $v_i(\mathbf{x};\lambda(t))$ will be a function of the genotype frequencies $\mathbf{x}$ at the current time step, and also have a dependence on the control parameter $\lambda(t)$ through the selection coefficient vector $\mathbf{s}(\lambda(t))$ (which influences $\langle \delta x_i \rangle$).

In non-evolutionary contexts $v_i(\mathbf{x};\lambda(t))$ 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 $D_{ij}(\mathbf{x})$ to describe the covariance of the genotype changes, defined through $2D_{ij}(\mathbf{x}) \equiv \langle \delta x_i \delta x_j \rangle - \langle \delta x_i \rangle \langle \delta x_j \rangle$. As shown in the Supplementary Information, to the lowest-order approximation, $D_{ij}(\mathbf{x})$ is independent of $\mathbf{s}(\lambda(t))$, and hence is not an explicit function of $\lambda(t)$. If we are interested in the dynamics on timescales much larger than a single generation, the probability $p(\mathbf{x},t)$ to observe a genotype state $\mathbf{x}$ at time t obeys a multivariate Fokker–Planck equation[48]:

$$\partial_t p(\mathbf{x},t) = -\partial_i(v_i(\mathbf{x};\lambda(t))p(\mathbf{x},t)) + \partial_i\partial_j(D_{ij}(\mathbf{x})p(\mathbf{x},t))$$
$$\equiv \mathcal{L}(\lambda(t))p(\mathbf{x},t) \tag{5}$$

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 $M-1$ (for example, the term $\partial_i\partial_j(D_{ij}(\mathbf{x})p(\mathbf{x},t)) \equiv \sum_{i=1}^{M-1}\sum_{j=1}^{M-1}\partial_i\partial_j(D_{ij}(\mathbf{x})p(\mathbf{x},t))$). The right-hand side of equation (5) defines the Fokker–Planck differential operator $\mathcal{L}(\lambda(t))$ in equation (1). To correspond to genotype fractions, vectors $\mathbf{x}$ 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 $\mathbf{x} \in \Delta$, then the wild-type fraction $x_M = 1 - \sum_{j=1}^{M-1} x_j$ lies between 0 and 1. Normalization of $p(\mathbf{x},t)$ takes the form $\int_\Delta d\mathbf{x}\, p(\mathbf{x},t) = 1$, where the integral is over the volume of the simplex Δ.

To complete the description of the model, we need expressions for the functions $v_i(\mathbf{x};\lambda(t))$ and $D_{ij}(\mathbf{x})$. Given a WF evolutionary model, these take the following form (see Supplementary Information for a detailed derivation):

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

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

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

Instantaneous equilibrium distributions. The IE distribution $\rho(\mathbf{x};\lambda(t))$ is defined through equation (2), $\mathcal{L}(\lambda(t))\rho(\mathbf{x};\lambda(t)) = 0$. Because we evaluate the effectiveness of our driving by comparing the actual distribution $p(\mathbf{x},t)$ to the IE distribution, it is useful to know the form of $\rho(\mathbf{x};\lambda(t))$. Unfortunately, it is generally not possible to find an IE analytical expression, except in some specific cases[28,46]. 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[28,46]:

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

where $\Phi(\mathbf{x};\lambda(t))$ is an effective 'potential' given by

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

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, $m_{\beta\alpha} \approx \mathcal{O}(\mu)$, for all non-zero matrix entries where $\alpha \neq \beta$, with $\mu N \gg 1$, $N \gg 1$. In this case we know that $\rho(\mathbf{x};\lambda(t))$ is approximately a multivariate normal distribution of the form

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

Here, $\overline{x}_i(\lambda) = \int_\Delta d\mathbf{x}\, x_i\rho(\mathbf{x};\lambda)$ 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 $\Sigma_{ij} \equiv \overline{x_ix_j} - \overline{x}_i\overline{x}_j$. To make practical use of equation (10), we need a method to estimate $\overline{x}_i(\lambda)$ and $\Sigma(\lambda)$. As shown in the Supplementary Information, this can be done through an approximate numerical solution to a set of exact equations involving the moments of $\rho(\mathbf{x};\lambda)$.

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

$$\partial_t\rho(\mathbf{x};\lambda(t)) = \widetilde{\mathcal{L}}(\lambda(t),\dot\lambda(t))\rho(\mathbf{x};\lambda(t))$$
$$= -\partial_i(\tilde{v}_i(\mathbf{x};\lambda(t),\dot\lambda(t))\rho(\mathbf{x};\lambda(t))) + \partial_i\partial_j(D_{ij}(\mathbf{x})\rho(\mathbf{x};\lambda(t))) \tag{11}$$

with a modified velocity function

$$\tilde{v}_i(\mathbf{x};\lambda(t),\dot\lambda(t)) = m_{i\mu}x_\mu + g_{ij}(\mathbf{x})\tilde{s}_j(\mathbf{x};\lambda(t),\dot\lambda(t))$$
$$= v_i(\mathbf{x};\lambda(t)) + g_{ij}(\mathbf{x})\delta\tilde{s}_j(\mathbf{x};\lambda(t),\dot\lambda(t)) \tag{12}$$

Using the fact that $\mathcal{L}(\lambda(t))\rho(\mathbf{x};\lambda(t)) = 0$, because $\rho(\mathbf{x};\lambda(t))$ is the IE distribution of the original operator $\mathcal{L}$, we can rewrite equation (11) as

$$\partial_t\rho(\mathbf{x};\lambda(t)) = -\partial_i\left(\rho(\mathbf{x};\lambda(t))g_{ij}(\mathbf{x})\delta\tilde{s}_j(\mathbf{x};\lambda(t),\dot\lambda(t))\right)$$
$$\equiv -\partial_i\mathcal{J}_i \tag{13}$$

where $\mathcal{J}_i \equiv \rho(\mathbf{x};\lambda(t))g_{ij}(\mathbf{x})\delta\tilde{s}_j(\mathbf{x};\lambda(t),\dot\lambda(t))$ is a probability current. In this form, equation (13) looks like a continuity equation, describing the local transport of probability density due to the current field $\mathcal{J}$. 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 $\mathbf{n}$ is normal to the boundary at the point. The perturbation $\delta\tilde{\mathbf{s}}(\mathbf{x};\lambda(t),\dot\lambda(t))$ that satisfies equation (13) and the boundary condition defines an exact CD protocol for the evolutionary system.

Given an arbitrary continuous time sequence of IE distributions $\rho(\mathbf{x};\lambda(t))$, such a perturbation always exists. In fact, from a formal mathematical standpoint[49], any perturbation of the following form is a solution (note that for clarity we do not use Einstein summation in this case):

$$\delta\tilde{\mathbf{s}}(\mathbf{x};\lambda(t),\dot\lambda(t))$$
$$= \mathbf{g}^{-1}(\mathbf{x})\left[\frac{1}{\rho(\mathbf{x};\lambda(t))}\left(-\sum_{i=1}^{M-1}\hat{\mathbf{x}}_i w_i \int_0^{x_i} dx_i'\, \partial_t\rho(x_1,\ldots,x_i',\ldots,x_{M-1};\lambda(t))\right. \right.$$
$$\left.\left.+\mathbf{f}(\mathbf{x};\lambda(t),\dot\lambda(t))\right)\right] \tag{14}$$

Here $\mathbf{g}^{-1}(\mathbf{x})$ is the inverse of the matrix $\mathbf{g}(\mathbf{x})$, $\hat{\mathbf{x}}_i$ 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 equation (14) that make the solution potentially non-unique: (1) the weights w_i can be any real numbers, so long as $\sum_{i=1}^{M-1} w_i = 1$; (2) $\mathbf{f}(\mathbf{x};\lambda(t),\dot\lambda(t))$ is an $(M-1)$-dimensional vector function that has zero divergence, $\partial_i f_i = 0$. However, we have additional constraints on this function $\mathbf{f}$: it has to be compatible with the vanishing of the current orthogonal to the boundary, $\mathcal{J}_i n_i = 0$. For $M=2$, where necessarily $w_1 = 1$, these constraints mean that only $\mathbf{f} = 0$ 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 $\mathbf{f}$ uniquely, and hence we get many possible allowable CD solutions, all of which satisfy equation (13). This in turn means that we can always find CD Fokker–Planck operators $\widetilde{\mathcal{L}}(\lambda(t),\dot\lambda(t))$ that satisfy equation (3).

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

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

By Gauss's theorem, $\int_\Delta d\mathbf{x}\, \partial_i(x_k\mathcal{J}_i) = \int_{\partial\Delta} d\sigma\, x_k\mathcal{J}_i n_i$, 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, that is, $\mathcal{J}_i n_i = 0$, so the first term in equation (15) is zero. Plugging the definition of $\mathcal{J}_k$ into the second term, we get

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

or equivalently

$$\partial_t \bar{\mathbf{x}}(\lambda(t)) = \langle \mathbf{g}(\mathbf{x}) \delta\tilde{\mathbf{s}}(\mathbf{x}; \lambda(t), \dot{\lambda}(t)) \rangle \qquad (17)$$

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

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

$$\partial_t \bar{\mathbf{x}}(\lambda(t)) = \mathbf{g}(\bar{\mathbf{x}}(\lambda(t))) \delta\tilde{\mathbf{s}}(\lambda(t), \dot{\lambda}(t)) \qquad (18)$$

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

Counterdiabatic driving for the two-genotype example. For the $M = 2$ system, the exact IE distribution is given by equations (8) and (9). In the large-population, frequent-mutation limit we can estimate the mean frequency $\bar{x}_1(\lambda(t))$ corresponding to this distribution as

$$\bar{x}_1(\lambda(t)) \approx \frac{-m_{12} - m_{21} + s_1(\lambda(t)) + \sqrt{(m_{12} + m_{21} - s_1(\lambda(t)))^2 + 4m_{12}s_1(\lambda(t))}}{2s_1(\lambda(t))} \qquad (19)$$

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

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

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

Agent-based model. *Model description.* For the 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 wild type) 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 WF generations t will be discussed below.) At each simulation time step, every cell in the population undergoes the following process: (1) with probability d it dies; (2) if it survives, the cell divides with a genotype-dependent probability

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

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 $\hat{m}_{ji}$, $j \neq i$.

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 flow chart of Supplementary Fig. 4. The $M = 2$ selection coefficient $s_1(\lambda(\tau))$ and other model parameters are as described above.

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 of $K = 5 \times 10^6$. 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 non-zero, the probability is $\hat{m}_{ji} = 2.5 \times 10^{-4}$, 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 $\tau \approx 10^4$, and then plateaus at later times. The dosage follows the equation

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

with parameters $a = 1.5 \times 10^{-4}$ M, $b = 2 \times 10^{-3}$ and $c = 10,110$. 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 antimalarial drug pyrimethamine[2,3]. To calculate the distributions of genotype frequencies, every simulation was repeated 1,000 times.

Mapping the agent-based model simulations to a Fokker–Planck equation. To implement the CD driving protocol, derived for WF 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 Supplementary Information, 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[48,50]. The resulting approximate correspondence is summarized as follows: (1) a duration of τ ABM simulation time steps maps to $t \approx \tau d$ WF generations, where d is the ABM death rate. (2) 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. (3) The effective 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).

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, $p(\mathbf{x}, t)$, and the IE one, $\rho(\mathbf{x}; \lambda(t))$, defined as $D_{\text{KL}}(\rho||p) = \int d\mathbf{x}\, \rho(\mathbf{x}; \lambda(t))\log_2(\rho(\mathbf{x}; \lambda(t))/p(\mathbf{x}, t))$. 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 (equations (8) and (9)). Hence the 1D integral for $D_{\text{KL}}(\rho||p)$ can be numerically evaluated. For $M = 16$ the situation becomes more complicated. There is no analytical solution for $\rho(\mathbf{x}; \lambda(t))$, but we do have a good approximation in terms of the multivariate normal distribution of equation (10), expressed in terms of the mean vector $\bar{\mathbf{x}}(\lambda(t))$ and covariance matrix $\Sigma(\lambda(t))$, which are calculated using the moment approach described in the Supplementary Information. The ABM simulation results are also normally distributed in this parameter regime, and hence there is a corresponding simulation mean $\bar{\mathbf{x}}_{\text{sim}}(t)$ and covariance $\Sigma_{\text{sim}}(t)$ that can be calculated at each time t. These are calculated from the ensemble of 1,000 simulations that are run for each parameter set. The integral for the KL divergence $D_{\text{KL}}(\rho||p)$ between the simulation and IE multivariate normal distributions can then be evaluated directly, yielding

$$\begin{aligned} D_{\text{KL}}(\rho||p) = \\ \frac{1}{2\ln 2}\Big[\ln\frac{\det\Sigma_{\text{sim}}(t)}{\det\Sigma(\lambda(t))} - M + 1 + \text{tr}\ \left(\Sigma_{\text{sim}}^{-1}(t)\Sigma(\lambda(t))\right) \\ + (\bar{\mathbf{x}}_{\text{sim}}(t) - \bar{\mathbf{x}}(\lambda(t)))^T \Sigma_{\text{sim}}^{-1}(t)(\bar{\mathbf{x}}_{\text{sim}}(t) - \bar{\mathbf{x}}(\lambda(t))) \Big] \end{aligned} \qquad (23)$$

Because $\Sigma_{\text{sim}}(t)$ 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 $\Sigma_{\text{sim}}^{-1}(t)$.

We can use the curves of $D_{\text{KL}}(\rho||p)$ 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. The lag time savings are $\Delta t = t_{\text{eq}}^{\text{orig}} - t_{\text{eq}}^{\text{CD}}$, where $t_{\text{eq}}^{\text{orig}}$ and $t_{\text{eq}}^{\text{CD}}$ are respectively the times at which probability distributions in the original and CD protocols reach their final IE target values. In terms of $D_{\text{KL}}(\rho||p)$, there is minimum value $D_{\text{KL}}^{\text{eq}}$ attained at long times when $p(\mathbf{x}, t)$ has converged with $\rho(\mathbf{x}; \lambda(t))$. Note that this value is not precisely zero because of numerical noise associated with estimation of the distribution $p(\mathbf{x}, t)$ from a finite number of simulations. At long times when $D_{\text{KL}}(\rho||p)$ approaches $D_{\text{KL}}^{\text{eq}}$, the final approach can be fit well by the following exponential decay function:

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

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

Data availability
The raw numerical data for the figures in the main text and Supplementary Information, as well as the code to generate the figures, are available via GitHub at

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.

References

46. Baxter, G. J., Blythe, R. A. & McKane, A. J. Exact solution of the multi-allelic diffusion model. *Math. Biosci.* **209**, 124–170 (2007).
47. Kimura, M. Stochastic processes and distribution of gene frequencies under natural selection. *Cold Spring Harb. Symp. Quant. Biol.* **20**, 33–53 (1955).
48. Gillespie, D. T. The multivariate Langevin and Fokker–Planck equations. *Am. J. Phys.* **64**, 1246–1257 (1996).
49. Sahoo, S. Inverse vector operators. Preprint at https://arxiv.org/pdf/0804.2239.pdf (2008).
50. Gillespie, D. T. The chemical Langevin equation. *J. Chem. Phys.* **113**, 297–306 (2000).

Acknowledgements

Author contributions

S.I. and J.P. performed mathematical analysis, wrote the two-allele code, peformed simulations, analysed the data and wrote the manuscript. J.C., E.I., O.G., B.K.-S. performed mathematical analysis, analysed the data and wrote the manuscript. E.D. and N.K. wrote the multidimensional ABM code, performed the simulations, analysed data and wrote the manuscript. J.G.S. analysed the data and wrote the manuscript. M.H. performed the mathematical analysis and simulations, wrote code, analysed the data and wrote the manuscript. S.D. wrote the manuscript, and S.D., E.I., J.G.S. and M.H. contributed to developing the overall theoretical framework.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41567-020-0989-3.

Supplementary information is available for this paper at https://doi.org/10.1038/s41567-020-0989-3.

Correspondence and requests for materials should be addressed to J.G.S. or M.H.

Peer review information *Nature Physics* thanks Ken Funo and Daniel Weinreich for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Extended Data Fig. 1 | 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. [2], 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.

Extended Data Fig. 2 | 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 Extended Data Fig. 1, except using the antimalarial drug cycloguanil instead of pyrimethamine. The seascape is based on the experimental data of Ref. [2], 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.