Understanding cellular growth strategies via optimal control

Mononen, Tommi

2023-01-04

Authors for correspondence:
Tommi Mononen
e-mail: tommi.mononen@helsinki.fi
Ville Mustonen
e-mail: v.mustonen@helsinki.fi

Understanding cellular growth strategies via optimal control

Tommi Mononen[1], Teemu Kuosmanen[1], Johannes Cairns[1] and Ville Mustonen[1,2]

[1]Department of Computer Science, Organismal and Evolutionary Biology Research Programme and [2]Institute of Biotechnology, University of Helsinki, Helsinki 00014, Finland

TM, 0000-0002-3603-0813; TK, 0000-0002-8724-5048; JC, 0000-0003-1329-2025; VM, 0000-0002-7270-1792

Evolutionary prediction and control are increasingly interesting research topics that are expanding to new areas of application. Unravelling and anticipating successful adaptations to different selection pressures becomes crucial when steering rapidly evolving cancer or microbial populations towards a chosen target. Here we introduce and apply a rich theoretical framework of optimal control to understand adaptive use of traits, which in turn allows eco-evolutionarily informed population control. Using adaptive metabolism and microbial experimental evolution as a case study, we show how demographic stochasticity alone can lead to lag time evolution, which appears as an emergent property in our model. We further show that the cycle length used in serial transfer experiments has practical importance as it may cause unintentional selection for specific growth strategies and lag times. Finally, we show how frequency-dependent selection can be incorporated to the state-dependent optimal control framework allowing the modelling of complex eco-evolutionary dynamics. Our study demonstrates the utility of optimal control theory in elucidating organismal adaptations and the intrinsic decision making of cellular communities with high adaptive potential.

1. Introduction

Microbial experimental evolution is used for detailed investigation of evolution in laboratory conditions [1]. Following the lead of the long-term evolution experiment of *Escherichia coli* project, spanning more than 70 000 generations, the field has blossomed during the last decade [2]. A typical experimental design assigns selection over a microbial population (or a community), and the population will then evolve under the chosen selection pressure. Clonal microbial populations are inoculated into medium and by a serial transfer protocol at the end of each culture cycle, a fixed proportion is transferred into fresh medium. This procedure continues over a selected number of cycles and the results are contrasted with the wild-type (ancestral) population or control populations propagated under identical laboratory conditions except for lacking the selection pressure of interest.

A trait that increases relative fitness will increase in frequency. However, in general, this can happen together with other traits that might be just incidental or compensatory (or only reveal themselves in different conditions). Traits are often connected with each other through genetic mechanisms, and adaptations cannot, therefore, necessarily occur independently, which is typically optimal. Such evolutionary constraints can slow down or prevent the occurrence of further adaptations [3]. Accepting that the traits themselves evolve under (possibly strong) constraints, we can still ask how these traits could be best used.

Optimal control theory [4] gives us the possibility to test how to optimally use traits, that evolution has already found, in a given environmental context. We can investigate, for example, the benefits of an adaptive use of lower and

higher metabolic rates. Microbes can sense population density via specific signalling molecules (quorum sensing) [5,6] or more general environmental cues (e.g. oxidative stress, single-stranded DNA from damaged cells, pH changes and accumulation of metabolic by-products at toxic levels), and react to nutrient deprivation by lowering their metabolic level [7] and entering to stasis [8]. As soon as a target is selected, for example, maximizing population size and environmental conditions are set (e.g. carrying capacity), we can derive the optimal strategy to use a set of given traits. By comparing such optimal solutions to wet laboratory evolution experiments and biological literature, we can assess how close populations are to using the optimal solution, and thus learn about growth strategies and their constraints.

Here, we will use optimal control theory to better understand growth strategies and their evolutionary consequences. We choose to use the traditional Bellman equations and discrete approach instead of commonly used ODE approaches, because we have to capture also small population size dynamics correctly.

Applying optimal control in the context of evolution is a re-emerging field [9], with current focus of optimal control on optimal therapies, directed trait evolution and host–pathogen systems [10–14]. By contrast, our setting here is from the perspective of the intrinsic optimal decision making within a population (see also [15]) in a context typical for laboratory evolution studies. Earlier applications can be found from the field of behavioural ecology, including studies to solve optimal age- and size-dependent life-history strategies [16–18]. Control theory has also been used to study how evolving populations use information and possible modes of transmitting genetic material [19,20], and there is a rising interest on the systems biology side as well [21,22].

We use the already mentioned metabolism example as a case study. We will compute and analyse an optimal solution for a microbe that is able to use two metabolic modes: lower and higher. The aim is to learn how to use these two modes to maximize the population size (yield) under a stochastic logistic growth model. The optimal solution can be seen as the expected outcome of an evolved microbial phenotype that has gone through a serial transfer experiment. Therefore, the phenotype will be adapted to the cycle length used [23] and maximizes population size at the end of a cycle. This type of selection can be introduced in laboratory conditions by establishing k parallel single cell populations each having their own compartment. At the end of the cycle, all populations are mixed, and the populations with higher growth become enriched when transferred into fresh media of the next k parallel populations [24]. Hence, we hypothesize that this kind of enrichment process can lead to cycle length adaptation. The length of lag phase and change in a population growth rate can make this adaptation process possible [25–27]. Additionally we assume that, for example, limited space, restricted nutrition or metabolic by-products can create a carrying capacity limitation in bacterial populations or communities.

We then change the context from monoculture to head-to-head competitions between different growth strategies. Adaptive phenotypes compete against other non-adaptive phenotypes having only a single metabolic mode. We hypothesize that adaptive phenotypes are the most competitive in these experiments.

2. Results

2.1. Optimal growth strategy in a monoculture

2.1.1. Monoculture growth model

Life-history strategies in cell populations can be characterized by each strategy's birth and death rates. Assume that a cell has two alternative metabolic modes, the *defensive* and the *aggressive* mode, with intrinsic birth and death rates (β_{def}, δ_{def}) and (β_{aggr}, δ_{aggr}), respectively. When using the defensive mode, the cell increases its survival at the expense of proliferation whereas in the aggressive mode the cell invests more of its resources to rapid proliferation, constituting a fundamental life-history trade-off between survival and proliferation. Hence we assume that the intrinsic birth and death rates satisfy the following basic properties:

$$
\left.
\begin{array}{ll}
\text{(i)} & 0 < \delta_{\text{def}} < \delta_{\text{aggr}} \\
\text{(ii)} & 0 < \beta_{\text{def}} - \delta_{\text{def}} < \beta_{\text{aggr}} - \delta_{\text{aggr}}.
\end{array}
\right\} \quad (2.1)
$$

The first assumption guarantees that the defensive mode has longer expected lifetime, whereas the second assumption guarantees that the aggressive strategy has a higher intrinsic *per capita* growth rate.

Finite resources necessarily imply the existence of some environmental feedback mechanism which regulates population density. Indeed, let us denote the total population size by N and let K be the population size at which the *per capita* growth rate becomes zero (the carrying capacity). We model the population growth by a discrete stochastic logistic model, where we assume that the carrying capacity realizes solely by increasing the death rate while allowing the birth rate to remain constant. The effective birth and death rates for different growth strategies are then

$$
\left.
\begin{array}{l}
\widetilde{\beta}(N_t, t) = \beta_{m(N_t, t)} \\
\text{and} \quad \widetilde{\delta}(N_t, t) = \delta_{m(N_t, t)} + \left(\beta_{m(N_t, t)} - \delta_{m(N_t, t)} \right) \dfrac{N_t}{K},
\end{array}
\right\} \quad (2.2)
$$

where $m : [0, Kl] \times [0, T] \to \{\text{def, aggr}\}$ is the mode which the cell uses at time $t \in [0, T]$ if the population is at size $N_t \in [0, Kl]$. The parameter $l > 1$ gives the range in which the strategy is determined above the carrying capacity, which is important to handle the stochastic fluctuations. Kl has to be larger than any typically reached population size (K + fluctuations), and allows us to make the system computable. Otherwise the state space would extend to arbitrarily large populations with probability of finding a population from the large sizes extremely small.

We discretize the continuous-time stochastic system by drawing the number of births and deaths from the corresponding Poisson distributions at each time step. Each simulation starts from initial population size N_{initial} and is updated at the next time step according to the rule

$$
\left.
\begin{array}{l}
N_{t+1} = N_t + \text{births}(t) - \text{deaths}(t), \\
\text{births}(t) \sim \text{Poisson}(N_t \widetilde{\beta}(N_t, t)) \\
\text{and} \quad \text{deaths}(t) \sim \text{Poisson}(N_t \widetilde{\delta}(N_t, t))
\end{array}
\right\} \quad (2.3)
$$

and the simulation proceeds until the chosen end-time T is reached. A population goes extinct, if $N_{t+1} \le 0$ (an absorbing boundary condition). The problem discussed here is to determine the optimal growth strategy $m_{\text{opt}}(N_t, t)$ for each t and N

which maximizes the expected final population size N_T (yield).

A strategy is called *pure*, if the same mode is used at all conceivable times and population sizes. If different modes are used at different times, we say that the strategy is *adaptive*. If the adaptive strategy is such that at each time point it chooses the mode which maximizes its current *per capita* growth rate, we say that the adaptive strategy is *greedy*. If instead the adaptive strategy restrains to maximize the instantaneous growth at some earlier time points in order to maximize the end population size (the cumulative *per capita* growth rate), we say that the adaptive strategy is *anticipative*. Next we present an optimal control framework which can be used to determine the optimal growth strategies and elucidate the situations where adaptive and anticipative strategies may evolve and replace pure strategies.

2.1.2. Forward propagation of probabilities

In order to compare different growth strategies, we need summary statistics from the process defined in equations (2.2) and (2.3). One approach is to estimate them from averages over individual simulation runs. Alternatively, we can propagate directly the probability distributions of population sizes over time. The change in population size is distributed as

$$\Delta N(t) \sim \text{Skellam}(\text{births}(t) - \text{deaths}(t); \mu_1, \mu_2) \quad (2.4)$$

where $\mu_1 = N_t \widetilde{\beta}(N_t, t)$ and $\mu_2 = N_t \widetilde{\delta}(N_t, t)$. Skellam distributions can be easily computed using convolution of two Poisson distributions. The transition probability from state N_t to state N_{t+1} using strategy $m(N_t, t)$ can then be written as

$$W(N_{t+1}|N_t; m(N_t, t)) = P(N_t + \Delta N(t)|N_t; m(N_t, t)), \quad (2.5)$$

and therefore forward propagation of probabilities can be computed as

$$\left. \begin{aligned} P(N_{\text{initial}}) &= 1 \\ \text{and} \qquad P(N_{t+1}) &= \sum_{N_t} P(N_t) W(N_{t+1}|N_t; m(N_t, t)). \end{aligned} \right\} \quad (2.6)$$

The Skellam distribution is defined also on the negative side (more deaths than births). Since negative population sizes are biologically unfeasible, we use an absorbing boundary condition at $N_t = 0$.

2.1.3. Stochastic optimal control solution

We assume that a cell can switch between aggressive and defensive modes without costs, choosing whichever mode is deemed to be more beneficial at that time. Further, only the final population size at the end of the finite time horizon ($t = T$) is considered to be important. Such requirement arises naturally during the serial transfer experiments, where the populations are diluted based on their final frequency, but also more generally in biotechnology, where the yield of beneficial bacteria or their metabolites is maximized within a given time-frame (see e.g. [28]). The optimal growth strategy leading to the largest population sizes, can be solved efficiently using the Bellman optimality equation [4].

Formally, we define the expected gain $J(N_t, t)$ describing how large population sizes we can reach in the future if we are in the population size N_t at time t and follow the optimal strategy until the end-time is reached. We compare the gains of aggressive and defensive modes at each time point and always select the one that leads to larger population sizes at the end. We can solve the optimal strategy $m_{\text{opt}}(N_t, t)$ iteratively starting from the end by using a backward-recurrence over all reachable population sizes

$$\left. \begin{aligned} J(N_t, t) &= \max_{m(N_t, t)} \sum_{N'} J(N', t+1) W(N'|N_t; m(N_t, t)) \\ \text{and} \quad m_{\text{opt}}(N_t, t) &= \underset{m(N_t, t)}{\text{argmax}} \ J(N_t, t), \end{aligned} \right\} \quad (2.7)$$

where $T \geq t \geq 1$ and with a boundary condition $J(N, T) = N$ defining our optimization target of maximizing the population size at the end T. Hence, this dynamic programming problem can be solved using equations (2.2)–(2.5) and (2.7), together with the boundary condition. Equation (2.7) can also be augmented with other terms that keep track of path rewards or costs at intermediate times and a cost term for use of control (here mode switching). Although including these could be interesting, e.g. to account for resource consumption of the dividing cells, we here focus on the minimal version where only the end population size is important and no cost is assigned to switch of growth mode.

The optimal strategy m_{opt} assigns the optimal growth mode for all conceivable states, hence achieving a complete control rule for every possible trajectory. We call this strategy a *control map*.

2.1.4. Control maps for monoculture

First, we investigated the basic behaviour of different growth strategies in a non-competitive setting to form a better understanding of their dynamics. We computed the control maps using the following intrinsic birth and death rates for the defensive and aggressive modes:

	birth rate (β)	death rate (δ)	growth rate
aggressive mode	0.125	0.075	0.05
defensive mode	0.025	0.005	0.02

The pure aggressive strategy leads to much faster growth than the pure defensive strategy, but has also much shorter life expectancy ($1/0.075 \approx 13.33$ time units compared with $1/0.005 = 200$ time units at initial growth phase). Combined, the average lifetime reproductive success is much higher in the defensive mode than in the aggressive mode ($0.025/0.005 = 5$ compared with $0.125/0.075 \approx 1.667$). Clearly, the chosen parameter values are to a degree arbitrary, but nevertheless the control maps remain qualitatively similar for all parameters satisfying the basic assumptions, which yield the growth-survival trade-off. Interestingly, we find that it is the cycle length T (the growth time before dilution), which changes the control maps qualitatively. For very short cycle lengths, the optimal growth strategy is to minimize the lag phase and grow as aggressively as possible when below the carrying capacity (figure 1*a*). For longer cycle lengths, when the population has enough time to reach the carrying capacity, the optimal strategy starts to use the defensive strategy both at low densities and above carrying capacity (figure 1*b*,*c*). Applying the correct mode becomes increasingly important closer to the time horizon, especially when the system is far from the target. The relative importance of applying the correct mode can be quantified by comparing the expected gains at each point.
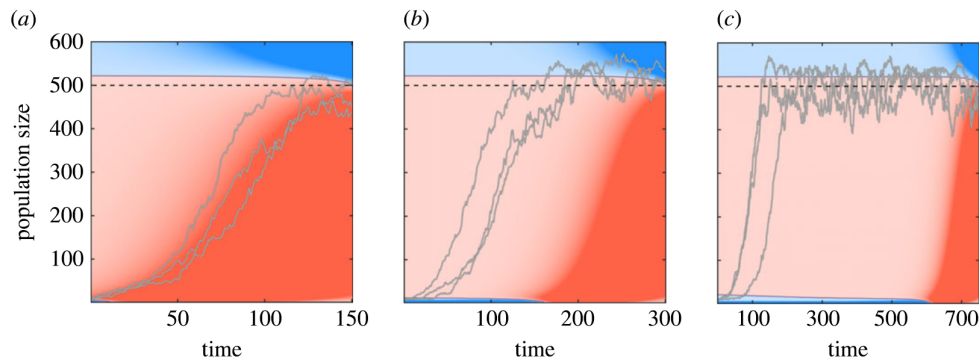
Figure 1. Control maps change qualitatively with respect to the end-time T: (a) For very short cycle lengths, the optimal growth strategy minimizes the lag and uses the aggressive mode (red shading) even when facing stochastic extinction risk ($T = 150$). (b) For longer cycle lengths another defensive mode region appears, where the optimal strategy uses the defensive mode (blue shading at a tiny population region). This creates a distinct lag phase which decreases the stochastic extinction risk ($T = 300$). (c) For very long cycle lengths the lag phase becomes even longer up to a point ($T = 750$). Carrying capacity ($K = 500$) is shown with a dashed black line. The intensity of the shading reflects the expected gain of using the optimal mode compared with the alternative.

2.1.5. Boundaries between two control modes

A boundary (contours separating blue and red regions in figure 1) between the two modes creates a directional flow from the aggressive side (red region) to the defensive side (blue region). The reason for the directional flow is that population size fluctuations are smaller in the defensive mode than in the aggressive mode. Therefore, a big jump from the aggressive side to the defensive side is more likely to occur compared with a reverse jump. Hence we call a boundary *minimizing* if larger population sizes above the threshold boundary use the aggressive mode, as such boundaries create a force that opposes population growth. By contrast, a boundary is *maximizing* if larger population sizes use the defensive mode, in which case the boundary creates a positive force towards larger population sizes.

Let us look at the two boundaries more carefully (figure 2). The minimizing boundary prolongs the apparent lag phase, slowing down the growth before initiating the exponential growth phase (figure 2*b*). The maximizing boundary on the other hand benefits from population size fluctuations that exceed carrying capacity (figure 2*a*). This boundary essentially tries to trap the population size above the carrying capacity by using the defensive mode whenever the boundary is exceeded (in effect saving resources by using the defensive mode).

The boundaries between the control modes change in time and the anticipative growth strategy (as given by the control map) leverages this time-dependence to maximize the expected yield. The final time slice of the control map ($t = T - 1$) corresponds to the greedy strategy, where the maximizing boundary always lies precisely at the carrying capacity. Thus, the length of the time interval becomes irrelevant and the greedy strategy stays constant over time. This corresponds to a simple adaptive strategy predicted by the classical r/K selection theory, where the aggressive mode is favoured at low densities and the defensive mode at high densities [29].

2.1.6. Cycle length induces selective pressure for lag time evolution

Next we compare the pure aggressive and defensive strategies with the greedy and the anticipatory growth strategy by propagating the corresponding probability distributions of the population sizes forward in time (figure 3). When

the cycle length is sufficiently long, both the aggressive and defensive strategies reach the carrying capacity and produce Gaussian-shaped densities around the carrying capacity. However, since the aggressive strategy has higher growth rate, there is more variance in the final population sizes. Adaptive phenotypes on the other hand can slightly exceed the carrying capacity on average via the efficient switching of the two modes. This is evident by observing the median population sizes (figure 3) and the probability distributions over the final population sizes. The optimal anticipative strategy gains higher yields than the greedy strategy in the monoculture due to its time-dependent control boundaries, which reduce stochastic extinctions and spur aggressive growth also slightly above the carrying capacity.

By defining the apparent lag time as the median time when the optimal strategy switches from the defensive mode to the aggressive mode, we can determine the optimal lag time from the propagated probability distributions as a function of the cycle length. Although the lag time is conventionally defined as the time until the first cell division, in practice the lag time is often determined by fitting slopes to a measured growth curve. Thus, we argue that our definition of lag cannot be distinguished from the observed lag times and that the lag time is an evolvable trait rather than a static or stochastic property [30].

The embedded panel in figure 3 shows the optimal lag time as a function of the cycle length. We note that for very short cycle lengths there is selection to minimize the lag. Then as the cycle length increases, we see that the optimal lag also increases linearly until finally plateauing at very long cycle lengths. This result shows how the used cycle length can generate selective pressure for specific growth strategies with fine-tuned lag times and hence give rise to the optimal, anticipative growth strategy.

2.2. Competition under a serial transfer protocol

2.2.1. Biculture competition model

We next investigated the different growth strategies under pairwise (head-to-head) competitions between the anticipative phenotype (monoculture optimal) and the other strategies. We assume that both of the adaptive phenotypes, greedy and anticipative, are able to sense their own, N_t^a, and the competitor's population sizes, N_t^c, and therefore the
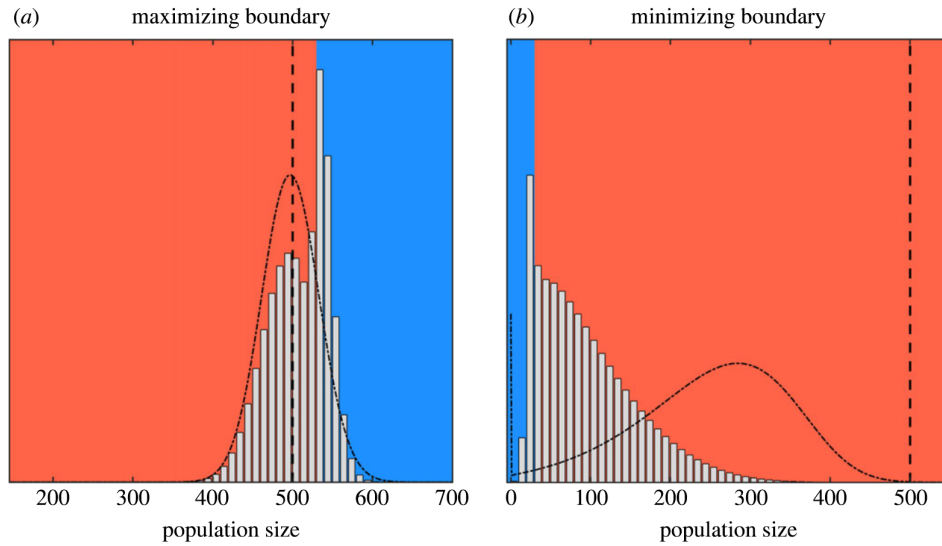
Figure 2. The propagated probability distributions of the optimal strategy (grey bins) compared with the aggressive strategy (dash-dotted line). The maximizing boundary traps the population above the carrying capacity, whereas the minimizing boundary traps the population below the boundary while reducing extinctions. The boundary itself is the reason for the emergence of the lag phase. (a) Distribution of the population sizes after $t = 200$ time steps starting from initial population size $N_{\text{initial}} = 500$. (b) Distribution of the population sizes after $t = 85$ time steps starting from initial population size $N_{\text{initial}} = 10$.
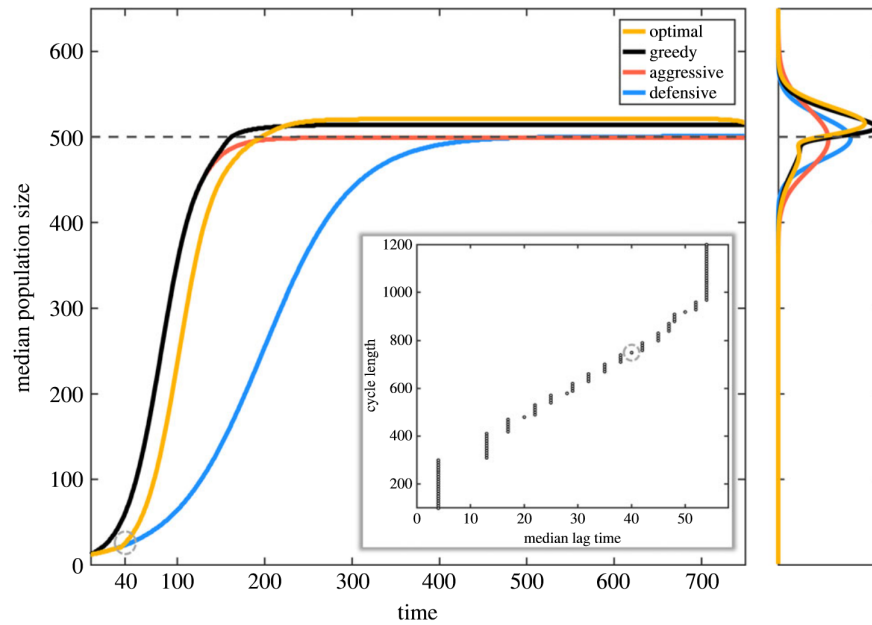
Figure 3. Median population sizes as a function of time for the aggressive, defensive, greedy and optimal (in monoculture) strategies. We note that the adaptive strategies reach higher final population sizes, the optimal marginally beating the greedy strategy. The pure aggressive and defensive strategies have the same median final population size, but the aggressive strategy displays higher variance. The lag time can be defined as the median time when the optimal strategy switches from the defensive to the aggressive strategy (see the dashed circle in the main panel). The embedded panel shows lag times as the function of cycle lengths (the dashed circle matches the cycle length of the main panel). For very short cycles, there is selection to eliminate the lag. For cycles longer than $T > 300$, there appears a distinct lag phase which increases linearly until around $T = 1000$ after which the lag time plateaus.

community density $N_t = N_t^a + N_t^c$. However, they either do not have access to any other information regarding the competitor's behaviour or lack the computational capabilities to use such information systematically to their advantage.

The competitor thus has the following death and birth rates:

$$\left. \begin{array}{l} \widetilde{\beta^c}(N_t^c, t) = \beta_{m(N_t, t)} \\[2mm] \text{and} \quad \widetilde{\delta^c}(N_t^c, t) = \delta_{m(N_t, t)} + \left(\beta_{m(N_t, t)} - \delta_{m(N_t, t)}\right)\dfrac{N_t}{K}, \end{array} \right\} \quad (2.8)$$

with the $m(N_t, t)$ strategy being either pure (defensive or aggressive) or under greedy control, the mode is defensive if $N_t > K$ or aggressive if $N_t \leq K$. Similarly the anticipative strategy has the following death and birth rates:

$$\left. \begin{array}{l} \widetilde{\beta^a}(N_t^a, N_t^c, t) = \beta_{m(N_t^a, N_t^c, t)} \\[2mm] \text{and} \quad \widetilde{\delta^a}(N_t^a, N_t^c, t) = \delta_{m(N_t^a, N_t^c, t)} + \left(\beta_{m(N_t^a, N_t^c, t)} - \delta_{m(N_t^a, N_t^c, t)}\right)\dfrac{N_t}{K}, \end{array} \right\} \quad (2.9)$$
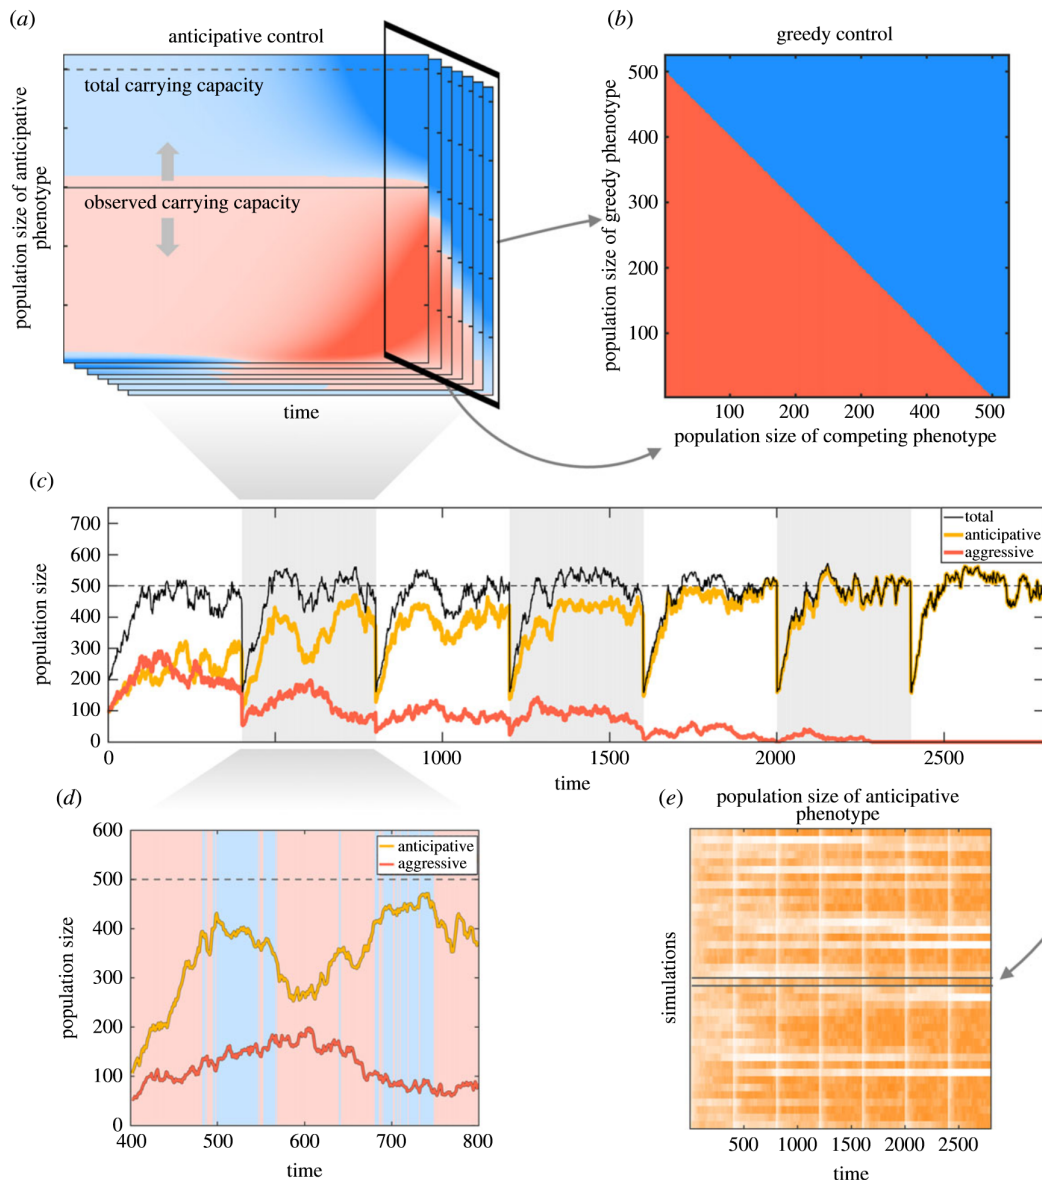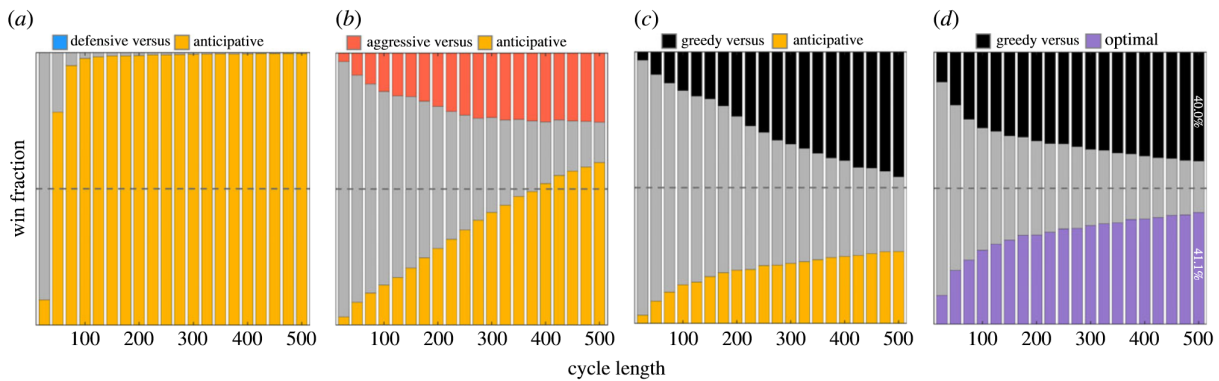
Figure 4. Biculture competition between phenotypes over multiple serial transfers. (a) Two phenotypes interact with each other via population densities. The observed carrying capacity by the optimal phenotype changes according to the population size of the competing phenotype. Therefore, the control map is a stack of slices (a control tensor), where the population size of the competing phenotype is fixed for each slice (see text). (b) The greedy strategy is an intersecting slice over the last control time point. (c) Competition between the anticipative and aggressive phenotypes takes place via mutual carrying capacity. Visible cycles emerge, because at the end of each cycle only a fraction of population (here 20%) is transferred into a fresh media. (d) The second cycle is enlarged to show the mode switches of the anticipative phenotype. The aggressive and defensive modes are visualized with red and blue backgrounds. (e) A simulation matrix shows 40 runs over time (the example run in (c) is emphasized). In most cases only one of the phenotypes survives through all seven cycles (a saturated orange colour denotes a large population size of the anticipative phenotype and the pure white colour its extinction).

with the adaptive strategy $m(N_t^a, N_t^c, t)$ calculated using equation (2.7) as before but now for each N_t^c value separately. Thus, in our competitive setting, the control map becomes a three-dimensional control tensor (figure 4a). At time t, the optimal monoculture phenotype uses the control map with a current population size of N_t^c. Therefore, the adaptive strategy can accurately sense the instantaneous carrying capacity $N_t/K = N_t^a/K + N_t^c/K$ and anticipate its own dynamics in approaching it. But it is not able to factor in the dynamics of the competitor. When competitor's population size (N_t^c) goes to zero, we recover the monoculture results. Similarly, if the competitor dynamics are slow compared with the time horizon T, population size N_t^c is approximately a

constant, and the anticipative strategy is optimal. However, for intermediate cases with fast competitor dynamics the inability for the anticipative strategy to take into account the dynamics of N_t^c leads to interesting results.

The growth strategy for the greedy phenotype is achieved via a straightforward expansion of the monoculture case. We take a slice over fixed time $t = T - 1$ (the last controlled time point) from the control tensor (figure 4a,b). Instead of a vector (with respect to the monoculture case), we have a control map that does not depend on time but only the population sizes of two phenotypes. The control map for the greedy phenotype is in our model symmetrical, changing its strategy from aggressive to defensive whenever the community density exceeds the carrying capacity.

Figure 5. Influence of cycle lengths on pairwise competitions: (*a–c*) the results of pairwise competitions (with six cycles) between the anticipative phenotype and the other phenotypes; (*d*) a pairwise competition between the greedy and the biculture optimal phenotype (10 cycles). A grey colour indicates mutual existence of both phenotypes. Dashed horizontal lines denote the 50% win fraction. (*a*) The anticipative phenotype overruns the competition between the defensive and the anticipative phenotype. However, with very short cycles, the defensive populations can mostly persist over a whole run. (*b*) The aggressive phenotype wins a small proportion of competitions against the anticipative, even with long cycles, as it grows faster than the anticipative phenotype in the beginning. Overall, the anticipative phenotype is stronger as it can adapt to changing observed carrying capacity. (*c*) The greedy phenotype combines faster growth and adaptivity. Therefore, it is a stronger competitor than the anticipative phenotype, which is optimized without the knowledge of competitor dynamics. (*d*) The optimal biculture phenotype with the knowledge of competitor dynamics performs just slightly better than the greedy phenotype. The length of competitions (10 cycles) is longer to show the difference between almost equally strong competitors. The panels were computed using $K = 500$ and 20 000 simulations per cycle length. Each new cycle was initiated using a fixed 0.1 K transfer proportion (50 individuals).

To demonstrate the setting, we first competed the anticipative (monoculture optimal) and aggressive phenotypes in an *in silico* experimental evolution simulation over several cycles, as shown in figure 4*c*. The anticipative phenotype takes advantage of the two different growth modes to outgrow its competitor (figure 4*d*). In the beginning of each cycle, the community size is diluted to a fixed fraction from where the community starts its new growth (figure 4*c*). In a wet laboratory, this corresponds to taking a sample from a microbial community and transferring it to fresh medium. After several cycles, the stronger competitor is eventually likely to conquer the whole community. Figure 4*e* shows 40 simulation runs with a fixed cycle length, where the aggressive phenotype competes against the anticipative phenotype. Initial population sizes for both phenotypes were set equal. After seven consecutive cycles the anticipative phenotype wins over 26 out of 40 competitions, loses four and in the rest of the simulations both phenotypes coexist. We note that also an inferior phenotype can sometimes win competitions due to random fluctuations.

2.2.2. Cycle length modifies the competitive fitness between the growth strategies

We next demonstrate that the anticipative (monoculture optimal) phenotype makes suboptimal decisions by assuming that the effective, at that time perceived, carrying capacity stays the same.

The cycle length has a large impact on the relative fitnesses of different phenotypes (figure 5). If the cycle length is too short, populations do not have time to grow and the whole community can end up extinct due to dilutions at the end of cycles. Even if a community survives, all phenotypes using the aggressive growth at low densities—aggressive, greedy and anticipative—are equally fit on average.

Adaptive phenotypes benefit from their two modes only when the cycle length is long enough for communities to

reach carrying capacity. Within these intermediate cycle lengths, the adaptive phenotypes start to gain the competitive advantage over non-adaptive phenotypes. The longer the cycle length, the greater the advantage.

The maximizing boundary of the greedy strategy (located at carrying capacity) seems to be more optimal in competition as the anticipative strategy slightly loses to it. The reason is that the community size, which matches carrying capacity is more often reached than community sizes even greater. Therefore, switching later (at larger population sizes) becomes a weaker strategy.

Smaller transfer ratios increase randomness and thereby weaken the relevance of mode switching for survival. In the beginning of each cycle, random fluctuations in population sizes become more important than success in the previous round.

2.2.3. Optimal biculture competitor

Finally, we computed the optimal strategy for the biculture scenario by straightforward extension of equation (2.7) to factor in the full dynamics of the competitor (see electronic supplementary material, S1–S3). The optimal strategy computed this way outperforms the greedy strategy playing competitor as shown in figure 5*d*. Interestingly, when approaching the joint carrying capacity the optimal anticipative strategy pre-emptively switches to using the defensive mode slightly before the carrying capacity is reached (see electronic supplementary material, figure S1). Therefore, the actual change between monoculture and biculture optimal strategies could be approximated by changing the location of the maximizing boundary to be somewhat below the carrying capacity in biculture. This simplicity leaves room for a genetically hard-coded rule to implement a winning control map with only sensing of the carrying capacity necessary and not requiring sophisticated computation as given by electronic supplementary material, equations S1–S3.

3. Discussion

In this study, we showed how optimal control theory provides a powerful methodology to assess optimal growth (life history) strategies in different settings and how they could be realized via evolution. We demonstrated in our adaptive metabolism case study how the ability to switch between aggressive and defensive growth modes can help monocultures reach higher yields and drive non-adaptive strains extinct. The adaptive phenotypes won competitions simply by exploiting the interplay of demographic stochasticity and the carrying capacity. Furthermore, we noted how the used cycle length of the serial transfer experiments has practical importance, as it may cause unintentional selection and lag time evolution. The experiments are likely to enrich strains fine-tuned for the specific time horizon.

These examples show how the study of rapid short-term evolution, often driven by human intervention, requires handling of stochastic population dynamics and an explicit notion of time (e.g. time-dependent interventions in drug therapies) to understand how populations can and will respond. However, it has also been argued that evolutionary processes are not always—and in every turn—optimizing the fitness [31]. Despite the fact that responses to selection can be seen as one type of optimization, this must be kept in mind.

In laboratory evolution experiments, microbial populations have been exposed, for example, to starvation and other harsh environments, predation, competition, environmental and temporal fluctuations and time pressure [32–36]. After an experiment, an evolved population will be investigated to reveal the genetic and phenotypic changes that have accumulated under the chosen selection regime. We can test in a shared environment how well the evolved microbes fare in a community consisting of wild-type microbes or other evolved phenotypes. These other phenotypes can introduce additional selective factors, which make the monoculture-evolved phenotype maladapted to the community [37–39]. This may cause the evolved microbes to be lost from the community [40], which is why we included the competition aspect into this paper.

The calculated control maps illuminate the optimal use of traits, giving rise to adaptive life-history strategies. They show how the optimal anticipative strategy with time-dependent control boundaries converges to the greedy strategy at the end of the time horizon. The greedy strategy itself can be (in the simple model considered) viewed as the optimal adaptive strategy at the deterministic limit. Hence, the control maps elegantly show the value of anticipation in any given model. Here we showed how anticipating demographic stochasticity and the limits of growth (carrying capacity) can prove to be useful, but certainly the rewards of anticipation can be much higher. For example, the adaptive immune system can be seen to effectively anticipate reinfection by the already seen pathogen [12].

Adaptive strategies require constant monitoring of environmental cues (such as quorum sensing in bacteria) and hence at least some sort of sensory system, which consumes energy (incurring a fitness cost). Anticipative strategies require even more sophisticated information processing. Hence, adaptive and anticipative strategies must have a substantial advantage over pure strategies to compensate the necessary investments in maintaining such systems [41,42]. Alternatively, the corresponding control maps must be simple enough to be approximately executed by some rule which is hard-coded genetically [12]. Or accessible by a heterogeneous population executing a mixed strategy based on the coexistence of pure strategies in a population as a form of bet-hedging. Interestingly, a recent study focusing on cellular decision making under volatile stochastic environments used methods from mathematical finance, as well as optimal control, to analyse how volatility affects the success of persister phenotypes [15].

We hypothesize that the intrinsic lag phase observed in bacterial growth may be an example of such a hard-coded rule, which can efficiently reduce the risk of stochastic extinctions. While the lag phase may serve also many other important functions which enable future proliferation, it is nevertheless worthwhile to note how the tendency for considerable lag phase arises in our model purely from demographic stochasticity alone. Given the observed variation in microbial lag times and the widespread potential for lag time evolution [30,43,44], our minimal model provides a simple explanation for why we observe distinct lag phases during microbial growth in the first place. Indeed, playing the defensive mode first rather than jumping directly to the aggressive mode may be highly advantageous for the establishment of new mutants and in the context of dispersal to new environments. Thus, the lag phase could be optimized in conjunction with selection of other beneficial traits, which may, therefore, either promote or constrain its evolvability [45].

One particularly interesting direction for future research would be to further investigate the extent of anticipative strategies in biology: are biological algorithms predominantly greedy and when can we expect anticipative behaviour to evolve? As our case study suggests, seemingly anticipative behaviour may emerge simply by the process of natural selection and adaptation to the prevailing environment. However, as our biculture competition model suggests, such optimal monocultures may be easily invaded by a greedy strategy (which may decrease the resource efficiency of the community).

Incorporating more complex and general environmental feedback and resource dynamics to our modelling framework will certainly allow for important insights. In these cases, it may be necessary to adopt other mathematical tools than the Bellman equations [21], if computational burden becomes otherwise too large. Therapies against cancer cell, microbial, pest and parasite populations can all be seen via the lens of systematic growth manipulation. Our setting here could be useful in that context in better understanding the optimal response of the target population to these interventions. Going beyond our case study, we also envision how the control theoretic approach could be highly useful also more generally to test various evolutionary hypotheses, elucidate organismal adaptations, predict likely evolutionary trajectories and steer evolution to a desired direction by manipulating the evolutionary drivers of the traits in question.

Authors' contributions. T.M.: conceptualization, formal analysis, investigation, methodology, software, visualization, writing—original draft; T.K.: conceptualization, formal analysis, investigation, methodology, software, writing—original draft; J.C.: investigation, writing—review and editing; V.M.: conceptualization, formal analysis, funding acquisition, investigation, methodology, supervision, visualization, writing—original draft.

9

References

1. McDonald MJ. 2019 Microbial experimental evolution—a proving ground for evolutionary theory and a tool for discovery. *EMBO Rep.* **20**, e46992. (doi:10.15252/embr.201846992)

2. Barrick JE, Yu DS, Yoon SH, Jeong H, Oh TK, Schneider D, Lenski RE, Kim JF. 2009 Genome evolution and adaptation in a long-term experiment with *Escherichia coli*. *Nature* **461**, 1243–1247. (doi:10.1038/nature08480)

3. Futuyma DJ. 2010 Evolutionary constraint and ecological consequences. *Evolution* **64**, 1865–1884. (doi:10.1111/j.1558-5646.2010.00960.x)

4. Todorov E. 2007 Optimal control theory (chapter 12). In *Bayesian brain: probabilistic approaches to neural coding* (eds K Doya, S Ishii, A Pouget, RPN Rao). Cambridge, MA: The MIT Press.

5. Darch SE, West SA, Winzer K, Diggle SP. 2012 Density-dependent fitness benefits in quorum-sensing bacterial populations. *Proc. Natl Acad. Sci. USA* **109**, 8259–8263. (doi:10.1073/pnas.1118131109)

6. An JH, Goo E, Kim H, Seo Y-S, Hwang I. 2014 Bacterial quorum sensing and metabolic slowing in a cooperative population. *Proc. Natl Acad. Sci. USA* **111**, 14 912–14 917. (doi:10.1073/pnas.1412431111)

7. Boutte CC, Crosson S. 2013 Bacterial lifestyle shapes stringent response activation. *Trends Microbiol.* **21**, 174–180. (doi:10.1016/j.tim.2013.01.002)

8. Nyström T. 2003 Conditional senescence in bacteria: death of the immortals. *Mol. Biol.* **48**, 17–23.

9. Schaffer WM. 1983 The application of optimal control theory to the general life history problem. *Am. Nat.* **121**, 418–431. (doi:10.1086/284070)

10. Fischer A, Vázquez-García I, Mustonen V. 2015 The value of monitoring to control evolving populations. *Proc. Natl Acad. Sci. USA* **112**, 1007–1012. (doi:10.1073/pnas.1409403112)

11. Lässig M, Mustonen V, Walczak AM. 2017 Predicting evolution. *Nat. Ecol. Evol.* **1**, 1–9. (doi:10.1038/s41559-017-0077)

12. Lässig M, Mustonen V. 2020 Eco-evolutionary control of pathogens. *Proc. Natl Acad. Sci. USA* **117**, 19 694–19 704. (doi:10.1073/pnas.1920263117)

13. Kuosmanen T, Cairns J, Noble R, Beerenwinkel N, Mononen T, Mustonen V. 2021 Drug-induced resistance evolution necessitates less aggressive treatment. *PLoS Comput. Biol.* **17**, e1009418. (doi:10.1371/journal.pcbi.1009418)

14. Nourmohammad A, Eksin C. 2021 Optimal evolutionary control for artificial selection on molecular phenotypes. *Phys. Rev. X* **11**, 011044.

15. Browning AP, Sharp JA, Mapder T, Baker CM, Burrage K, Simpson MJ. 2021 Persistence as an optimal hedging strategy. *Biophys. J.* **120**, 133–142. (doi:10.1016/j.bpj.2020.11.2260)

16. Sibly R, Calow P, Nichols N. 1985 Are patterns of growth adaptive? *J. Theor. Biol.* **112**, 553–574. (doi:10.1016/S0022-5193(85)80022-9)

17. Houston AI, McNamara JM. 1999 *Models of adaptive behaviour: an approach based on state*. Cambridge, UK: Cambridge University Press.

18. Mangel M, Clark CW. 2019 *Dynamic modeling in behavioral ecology*. Princeton, NJ: Princeton University Press.

19. Rivoire O, Leibler S. 2011 The value of information for populations in varying environments. *J. Stat. Phys.* **142**, 1124–1166. (doi:10.1007/s10955-011-0166-2)

20. Rivoire O, Leibler S. 2014 A model for the generation and transmission of variations in evolution. *Proc. Natl Acad. Sci. USA* **111**, E1940–E1949. (doi:10.1073/pnas.1323901111)

21. Sharp JA, Burrage K, Simpson MJ. 2021 Implementation and acceleration of optimal control for systems biology. *J. R. Soc. Interface* **18**, 20210241. (doi:10.1098/rsif.2021.0241)

22. Tsiantis N, Banga JR. 2020 Using optimal control to understand complex metabolic pathways. *BMC Bioinf.* **21**, 472. (doi:10.1186/s12859-020-03808-8)

23. Li Y, Petrov DA, Sherlock G. 2019 Single nucleotide mapping of trait space reveals pareto fronts that constrain adaptation. *Nat. Ecol. Evol.* **3**, 1539–1551. (doi:10.1038/s41559-019-0993-0)

24. Bachmann H, Fischlechner M, Rabbers I, Barfa N, Branco dos Santos F, Molenaar D, Teusink B. 2013 Availability of public goods shapes the evolution of competing metabolic strategies. *Proc. Natl Acad. Sci. USA* **110**, 14 302–14 307. (doi:10.1073/pnas.1308523110)

25. Oxman E, Alon U, Dekel E. 2008 Defined order of evolutionary adaptations: experimental evidence. *Evol. Int. J. Org. Evol.* **62**, 1547–1554. (doi:10.1111/evo.2008.62.issue-7)

26. Augustin J-C, Brouillaud-Delattre A, Rosso L, Carlier V. 2000 Significance of inoculum size in the lag time of listeria monocytogenes. *Appl. Environ. Microbiol.* **66**, 1706–1710. (doi:10.1128/AEM.66.4.1706-1710.2000)

27. Levy S, Kafri M, Carmi M, Barkai N. 2011 The competitive advantage of a dual-transporter system. *Science* **334**, 1408–1412. (doi:10.1126/science.1207154)

28. St. John PC, Crowley MF, Bomble YJ. 2017 Efficient estimation of the maximum metabolic productivity of batch systems. *Biotechnol. Biofuels* **10**, 28. (doi:10.1186/s13068-017-0709-0)

29. Reznick D, Bryant MJ, Bashey F. 2002 *r*- and *k*-selection revisited: the role of population regulation in life-history evolution. *Ecology* **83**, 1509–1520. (doi:10.1890/0012-9658(2002)083[1509:RAKSRT]2.0.CO;2)

30. Bertrand RL. 2019 Lag phase is a dynamic, organized, adaptive, and evolvable period that prepares bacteria for cell division. *J. Bacteriol.* **201**, e00697-18. (doi:10.1128/JB.00697-18)

31. Gould SJ, Lewontin RC. 1979 The spandrels of san marco and the panglossian paradigm: a critique of the adaptationist programme. *Proc. R. Soc. B* **205**, 581–598.

32. Dragosits M, Mattanovich D. 2013 Adaptive laboratory evolution—principles and applications for biotechnology. *Microb. Cell Fact.* **12**, 1–17. (doi:10.1186/1475-2859-12-64)

33. Nair RR, Vasse M, Wielgoss S, Sun L, Yu YTN, Velicer GJ, 2019 Bacterial predator-prey coevolution accelerates genome evolution and selects on virulence-associated prey defences. *Nat. Commun.* **10**, 4301. (doi:10.1038/s41467-019-12140-6)

34. Kim W, Racimo F, Schluter J, Levy SB, Foster KR. 2014 Importance of positioning for microbial evolution. *Proc. Natl Acad. Sci. USA* **111**, E1639–E1647. (doi:10.1073/pnas.1323632111)

35. Karve SM, Daniel S, Chavhan YD, Anand A, Kharola SS, Dey S. 2015 *Escherichia coli* populations in unpredictably fluctuating environments evolve to face novel stresses through enhanced efflux activity. *J. Evol. Biol.* **28**, 1131–1143. (doi:10.1111/jeb.12640)

36. Travisano M, Maeda M, Fuji F, Kudo T. 2018 Rapid adaptation to near extinction in microbial experimental evolution. *J. Bioecon.* **20**, 141–152. (doi:10.1007/s10818-017-9257-8)

37. Goodnight CJ. 1990 Experimental studies of community evolution II: the ecological basis of the response to community selection. *Evolution* **44**, 1625–1636.

38. Xie L, Yuan AE, Shou W. 2019 Simulations reveal challenges to artificial community selection and possible strategies for success. *PLoS Biol.* **17**, e3000295. (doi:10.1371/journal.pbio.3000295)

39. De Meester L et al. 2018 Analysing eco-evolutionary dynamics—the challenging complexity of the real world. *Funct. Ecol.* **33**, 43–59. (doi:10.1111/fec.2019.33.issue-1)

40. Manhart M, Shakhnovich EI. 2018 Growth tradeoffs produce complex microbial communities on a single limiting resource. *Nat. Commun.* **9**, 3214. (doi:10.1038/s41467-018-05703-6)

41. Acerenza L. 2016 Constraints, trade-offs and the currency of fitness. *J. Mol. Evol.* **82**, 117–127. (doi:10.1007/s00239-016-9730-3)

42. Ruparell A. et al. 2016 The fitness burden imposed by synthesising quorum sensing signals. *Sci. Rep.* **6**, 33101. (doi:10.1038/srep33101)

43. Moreno-Gámez S, Kiviet DJ, Vulin C, Schlegel S, Schlegel K, van Doorn GS, Ackermann M. 2020 Wide lag time distributions break a trade-off between

reproduction and survival in bacteria. *Proc. Natl Acad. Sci. USA* **117**, 18 729–18 736. (doi:10.1073/pnas.2003331117)

44. Blount ZD, Borland CZ, Lenski RE. 2008 Historical contingency and the evolution of a key innovation in an experimental

population of *Escherichia coli*. *Proc. Natl Acad. Sci. USA* **105**, 7899–7906. (doi:10.1073/pnas.0803151105)

45. Adkar BV, Manhart M, Bhattacharyya S, Tian J, Musharbash M, Shakhnovich EI. 2017 Optimization of lag phase shapes the evolution of a bacterial

enzyme. *Nat. Ecol. Evol.* **1**, 1–6. (doi:10.1038/s41559-017-0149)

46. Mononen T, Kuosmanen T, Cairns J, Mustonen V. 2023 Understanding cellular growth strategies via optimal control. Figshare. (doi:10.6084/m9.figshare.c.6350197)