PNAS PLUS

The Red Queen and King in finite populations

Carl Veller[a,b,1], Laura K. Hayward[c], Christian Hilbe[d], and Martin A. Nowak[a,b,e]

[a]Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138; [b]Program for Evolutionary Dynamics, Harvard University, Cambridge, MA 02138; [c]Department of Mathematics, Columbia University, New York, NY 10027; [d]Institute of Science and Technology Austria, 3400 Klosterneuburg, Austria; and [e]Department of Mathematics, Harvard University, Cambridge, MA 02138

Edited by Steven A. Frank, University of California, Irvine, CA, and accepted by Editorial Board Member Joan E. Strassmann May 23, 2017 (received for review February 6, 2017)

In antagonistic symbioses, such as host–parasite interactions, one population's success is the other's loss. In mutualistic symbioses, such as division of labor, both parties can gain, but they might have different preferences over the possible mutualistic arrangements. The rates of evolution of the two populations in a symbiosis are important determinants of which population will be more successful: Faster evolution is thought to be favored in antagonistic symbioses (the "Red Queen effect"), but disfavored in certain mutualistic symbioses (the "Red King effect"). However, it remains unclear which biological parameters drive these effects. Here, we analyze the effects of the various determinants of evolutionary rate: generation time, mutation rate, population size, and the intensity of natural selection. Our main results hold for the case where mutation is infrequent. Slower evolution causes a long-term advantage in an important class of mutualistic interactions. Surprisingly, less intense selection is the strongest driver of this Red King effect, whereas relative mutation rates and generation times have little effect. In antagonistic interactions, faster evolution by any means is beneficial. Our results provide insight into the demographic evolution of symbionts.

rate of evolution | symbiosis | mutualism | antagonism | Müllerian mimicry

Antagonistic symbioses can be conceptualized well in the following simple, constant-sum game (1, 2):

		Player 2	
		C	D
Player 1	A	$1,0$	$0,1$
	B	$0,1$	$1,0$

[1]

Here, player 1's available strategies are A and B, player 2's are C and D, the first payoff in each cell is player 1's, and the second one is player 2's. Interacting populations of "player 1s" and "player 2s" constitute an antagonistic symbiosis, which we expect to evolve according to arms-race dynamics (3). To see why, suppose we start with population 1 all playing strategy A and population 2 all playing strategy C, so that population 1 is doing well at the expense of population 2. Now, a mutant in population 2 who plays strategy D does better (payoff 1) than other members of population 2 (payoff 0) and can take over, at population 1's expense (their payoff decreases from 1 to 0). Population 1 is then expected to switch to strategy B, after which population 2 switches to strategy C, and so on. As the Red Queen said to Alice, "it takes all the running *you* can do, to keep in the same place" (ref. 4, p. 42).

Arms-race dynamics characterize many interactions both in the natural world and in human behavior. In a host–parasite interaction, for example, the host is selected to develop immunity to the parasite, which in turn selects for new "resistance" mutations in the parasite, which selects for the host to develop immunity to the new mutant parasite, and so on. In Batesian mimicry, a palatable species (of butterfly, for example) evolves to mimic the warning display of an unpalatable species, so that it is mistaken by predators for the unpalatable species. This selects for the unpalatable species to evolve a new display, to not be mistaken for the palatable species, which is then under selection to mimic the new display, and so on (5).

Antagonistic interactions play out at an intragenomic level as well. Examples, among many others (6, 7), are the X-linked male-meiotic driving gene *Dox* and its autosomal suppressor *Nmy* in *Drosophila simulans* (8, 9), centromeric repeat sequences seeking to drive in female meiosis and the centromere-histone genes that keep them in check (10, 11), and the mammalian recombination-specifying gene *Prdm9* and the binding sequences of its protein [which are under positive selection to escape PRDM9 binding (12–15)]. Consistent with arms-race dynamics, the conflicting elements in these examples all show genetic evidence of rapid evolution (10, 11, 16–23).

Mutualistic symbioses with a degree of conflict (24, 25) can also be conceptualized in a simple game (26),

		Player 2	
		C	D
Player 1	A	$2,1$	$0,0$
	B	k,k	$1,2$

[2]

with $0 \leq k < 2$. From either of the two mutualistic coordination states (A, C) and (B, D), neither population is under selection to deviate to one of the noncoordination states (A, D) and (B, C). But population 1 prefers coordination state (A, C), whereas population 2 prefers coordination state (B, D). Of most interest is the case $k < 1$, where game **2** is a true mutualism (both populations prefer the two coordination states to the two noncoordination states).

This game also describes many interactions in both nature and humans. Mutualisms often involve a division of labor (27), such as the production of different nutrients in a microalgal–microbial partnership (28), the separate production of a "poison" and its "antidote" by two linked genetic elements in a gamete-killing meiotic drive complex (6, 29), or the economic production of different goods by trading partners under increasing returns to scale (30). In these cases, it is beneficial to both interactants that

Significance

When two populations interact, when does it pay to evolve rapidly, and can it ever be an advantage to evolve slowly? We address these questions using evolutionary game theory. In antagonistic interactions (e.g., host–parasite), we find that faster evolution by any means is beneficial—the "Red Queen" effect. In certain mutualisms, slower evolution is favored in the long run. This "Red King" effect is driven by differences in how efficiently natural selection acts in the two populations, rather than by differences in their generation times or mutation rates. Our results clarify the role of evolutionary rate in symbiont evolution.

Author contributions: C.V., L.K.H., C.H., and M.A.N. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. S.A.F. is a guest editor invited by the Editorial Board.

[1]To whom correspondence should be addressed. Email: carlveller@fas.harvard.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1702020114/-/DCSupplemental.

all tasks be carried out, but there can be conflict over who should do which (31). For example, it is better to be the antidote producer than the poison producer in the meiotic drive complex, because the gains of their partnership are shared equally, but its disruption (by recombination) is costlier to the poison producer.

Another example, Müllerian mimicry, involves two poisonous species (of butterfly, say) evolving to share a common pattern, one pattern being easier for their predators to learn than two (32). But if the species' original patterns evolved to suit differences in their respective habitats, behaviors, or genomic backgrounds, then there might be conflict over which pattern to converge on, with each species enjoying an evolutionary advantage if convergence is on its original pattern.

In both antagonistic and mutualistic symbioses, the relative rates of evolution of the participating populations are thought to be important determinants of their relative evolutionary success. A faster rate of evolution could be achieved through various means (3, 24–26). First, it could derive from a shorter generation time, allowing more generations over which to adapt (2). Second, it could derive from a higher rate of mutation, more rapidly generating new variants that allow a population to escape an unfavorable state of interaction (24, 33, 34). Finally, it could derive from more effective action of natural selection. This could be because the stakes are higher for one population than the other—"the rabbit is running for his life, while the fox is only running for his dinner" (ref. 3, p. 493)—or because selection acts more effectively in large populations, owing (roughly speaking) to a reduced effect of random drift (35) (a precise formulation of this statement is given in *SI Appendix*, section S5.2).

Common wisdom holds that a faster rate of evolution is advantageous in antagonistic interactions (36), the better to "keep ahead in the race" (ref. 3, p. 492). But it has not, to our knowledge, been clearly demonstrated which individual biological parameters drive this Red Queen effect in theory.

In an important paper, Bergstrom and Lachmann (26) (hereafter, B&L) demonstrated that a slower rate of evolution might be advantageous in some mutualistic interactions, an effect they called the Red King effect. They studied two populations interacting according to the mutualism game **2**, with deterministic evolutionary dynamics operating in each of the two populations, both of infinite size and without the possibility of mutations.

In their model, x is the proportion of population 1 playing strategy A (so that the proportion $1-x$ play B), and y is the proportion of population 2 playing D ($1-y$ play C). Then the expected payoff to a member of population 1 who plays A when interacting with a random member of population 2 is, from the payoff matrix **2**, $\pi_1^A = 2(1-y)$. On the other hand, the payoff to a member playing B is $\pi_1^B = k(1-y)+y$, and the average payoff in population 1 is $\bar{\pi}_1 = x\pi_1^A + (1-x)\pi_1^B$. The analogous quantities in population 2, π_2^C, π_2^D, and $\bar{\pi}_2$, are calculated similarly. The strategy frequencies in the two populations evolve according to replicator dynamics:

$$\dot{x} = mx\left[\pi_1^A(y) - \bar{\pi}_1(x,y)\right],\ \ \dot{y} = ny\left[\pi_2^D(x) - \bar{\pi}_2(x,y)\right]. \quad \textbf{[3]}$$

Here $m, n > 0$ are parameters that calibrate the relative rates of evolution of the two populations by determining how responsive their respective evolutionary dynamics are to fitness differences among strategies.

These dynamics are deterministic—evolutionary trajectories are fully determined once their starting points are known—and evolution necessarily leads to one of the two coordination equilibria, (A, C) ($x=1, y=0$; population 1's preference) or (B, D) ($x=0, y=1$; population 2's preference), where evolution then halts.

B&L (26) showed that, when $k>1$ in game **2** and population 1 evolves slower than population 2 ($m<n$), the set of starting points from which evolution proceeds to population 1's favored equilibrium (A, C) (the equilibrium's "basin of attraction") is larger than the basin of attraction of the equilibrium (B, D). In this sense, slower evolution is beneficial. When $k<1$, the opposite result holds: If population 1 evolves slower, the basin of attraction of (A, C) is smaller than that of (B, D).

Here, we construct a finite-population model of symbiosis evolution, incorporating all of the biological determinants of evolutionary rate. This model allows us to extend B&L's (26) results for mutualisms in several ways and to apply a similar analysis to antagonistic symbioses, leading to a richer picture of the evolutionary dynamics of symbioses.

First, our model allows for an explicit characterization of which evolutionary rate parameters influence the relative success of the interacting populations. This is unknown for mutualisms [B&L's (26) general rate parameters m and n have no clear biological interpretation], has not been fully disentangled for antagonistic symbioses, and in both cases is crucial for standard empirical measures of evolutionary rate [such as the substitution rate at neutral genetic loci, which depends on mutation rate and generation time, but is insensitive to changes in population size (37)].

Second, our model allows us to uncover a key influence of evolutionary timescale on symbiosis evolution. In the mutualism game **2**, we show that the short-run behavior of our stochastic evolutionary dynamics is very similar to that of the replicator dynamics studied by B&L (26): From a given starting state, the dynamics rapidly converge to, or near to, one of the two equilibria of the game. Which equilibrium is most likely to be approached depends critically on the starting point, just as in the dynamics of B&L (26). On a longer timescale, however, we show the evolutionary dynamics of this game to be of a very different nature. The long-run dynamics involve transitions between equilibria, driven by sporadic mutation, and eventually become independent of where the dynamics first started. As we show, in mutualisms, many of the conclusions concerning the short-run dynamics are either annulled or reversed in the long run.

Our main results are summarized in Table 1.

A Finite-Population Model of Symbiosis Evolution

Populations 1 and 2 are of sizes N_1 and N_2 and interact according to a two-player, two-strategy game such as games **1** and **2**. The "populations" here can be broadly construed: They could be all of the individuals of two species in a symbiosis, for example, or all of the alleles at two distinct loci among the genomes of a single species.

The evolutionary process occurs in discrete time steps. Each time step, individuals in each population receive their average payoffs from interacting with a random member of the opposite population. An individual's payoff π in population l is translated to a nonnegative fitness value $1 + w_l\pi$, so that the "selection strengths" $w_1, w_2 > 0$ calibrate the effectiveness of natural selection in the two populations.

In each time step, a "birth–death event" occurs in one of the populations. Individuals in the two populations have relative generation times g_1 and g_2: In a given time step, the birth–death event occurs in population l with probability proportional to N_l/g_l, independently across time steps.

A birth–death event in a population involves choosing an individual to reproduce, with probability proportional to fitness, and an individual to die, with each one equally likely. These can be the same individual. A single offspring of the reproducing individual replaces the individual that was chosen to die. This within-population process, the Moran process (38, 39), has been used as a model both of biological evolution (40, 41) and of imitation learning (42–44). In the Moran process, one "generation" of a population typically corresponds to the number of birth–death events that is about the same as the population's size (41). In our two-population framework, we label the number of time steps

Table 1. The biological parameters that drive the Red Queen and King effects in antagonistic and mutualistic symbioses

	Antagonistic symbiosis		Mutualism, *k* small		Mutualism, *k* large	
	Short run	Long run	Short run	Long run	Short run	Long run
1/generation time	♕Q	♕Q	♕Q	No effect	♔K	No effect
Mutation rate	No effect	♕Q	No effect	No effect	No effect	No effect
Selection strength	♕Q	♕Q	♕Q	♔K	♔K	♕Q
Population size	♕Q	♕Q	♕Q	♔K	♔K	♕Q

For each of the parameters that determine evolutionary rate, we ask whether the population with the larger parameter value, which therefore evolves faster, is more successful (♕Q, a Red Queen effect) or less successful (♔K, a Red King effect) in the interaction, holding the other rate parameters constant and equal between the two populations. The short-run results are numerically computed for particular parameter values (Figs. 1–3). The long-run results are for the weak-mutation limit and are exact (see text for details). For selection strength and population size in the mutualisms, (*i*) we set one population's selection strength to w (or both populations', when studying the effect of population size) and assume that the larger parameter value is larger by a small amount, (*ii*) we define "*k* small" as $k < 1/(1+w)$ and "*k* large" as $k > 1/(1+w)$, and (*iii*) we assume the populations to be sufficiently large.

equal to the sum of the populations' sizes as a common generation: Each individual experiences on average one birth–death event per generation.

Finally, to account for the possibility of mutation, we assume that, in a birth–death event in population l, the offspring inherits its parent's strategy with probability $1-\varepsilon\mu_l$ or mutates to the alternative strategy with probability $\varepsilon\mu_l$. Thus, μ_1 and μ_2 represent the relative mutation rates of the two populations (they are unitless), whereas the parameter ε calibrates the overall frequency of mutations in the two populations (in units "per replication").

In sum, given a payoff matrix such as **[1]** and **[2]** and rate parameters N_1, N_2, g_1, g_2, w_1, w_2, μ_1, μ_2, and ε and provided $\mu_1, \mu_2, \varepsilon > 0$, the above defines an ergodic Markov chain over a state space comprising all possible strategy compositions of the two populations. If $0 \leq i \leq N_1$ is the number of A strategists in population 1 at some point in time (so that $N_1 - i$ individuals play B), and $0 \leq j \leq N_2$ is the number of C strategists in population 2 ($N_2 - j$ play D), then the population state is (i, j). The behavior of the evolutionary process is characterized by the probability of moving from state (i, j) to (i', j') in one time step, for all such pairs of states—these probabilities are provided in *SI Appendix*, section S1.

Two regimes are of interest in these dynamics, corresponding roughly to their short-run and long-run behavior.

Given some initial population state (i^0, j^0), the short-run dynamics (*i*) are not affected much by mutations, instead being governed mostly by selection strengths, population sizes, and generation times; (*ii*) depend critically on the starting point (i^0, j^0); (*iii*) have trajectories that are similar to those of the replicator dynamics, especially when the populations are large; and (*iv*) in coordination games, like game **2**, converge rapidly to or near one of the pure population states (in which each population is monomorphic) associated with the equilibria of the game. For illustration, Figs. 2 *A–D*, *Upper* and 3 *A–D*, *Upper* and especially *SI Appendix*, Fig. S1 display the similarity of the short-run behavior of our dynamics to that of the replicator dynamics, in the context of the mutualism game **2** [compare figure 2 of B&L (26)].

The long-run dynamics are of a very different nature. Because the evolutionary process, as we have defined it, is ergodic, the probability that the system is in some state in the future eventually becomes independent of where the system started (45). This effect is illustrated for the antagonistic symbiosis game in Fig. 1 *A–D*, *Lower* and for the mutualism game in Figs. 2 *A–D*, *Lower* and 3 *A–D*, *Lower*. The state of the system comes to depend not on the early dynamics that emanate from the starting point, but instead on infrequent transitions between equilibria, driven by mutations.

The object of interest in these long-run dynamics is their stationary distribution, the proportion of time spent in each population state in the long run. Equivalently, the stationary distribution tells us, were we to observe many independent instances of equivalent symbioses evolving, what proportion of these symbioses we should expect to find in each possible state at some fixed point in time in the long run (45).

Whereas we numerically study the stationary distribution of our evolutionary process for large mutation rates and selection strengths, we also invoke recent methodological advances in evolutionary game theory (46, 47) to study it analytically in certain limits. Chief among these is the "weak-mutation limit," $\varepsilon \to 0$, which approximates the case where mutations are very infrequent: $N_1\varepsilon\mu_1, N_2\varepsilon\mu_2 \ll 1$ (44, 46). This is a common assumption in the population genetics and evolutionary game theory literatures (44, 48, 49) and is realistic when populations are small or when the individual mutation rate is very small. For example, in genetical evolution, the mutations we are considering might be single-nucleotide substitutions, which occur at rates of order 10^{-8} or smaller per generation for most organisms (50); if multiple nucleotide substitutions are required to change strategies, then the relevant mutation rates are even lower. Therefore, whereas the case of infrequent mutations is certainly not general, it is a relevant and interesting case to consider.

In the weak-mutation limit, the dynamics converge to an embedded dynamical process over just the four pure states (42, 46); the stationary distribution collapses to a probability distribution over these pure states, $\boldsymbol{\lambda} = \left[\lambda_{(A,C)}, \lambda_{(A,D)}, \lambda_{(B,C)}, \lambda_{(B,D)}\right]$. In the embedded dynamics, transitions between pure states occur with probabilities determined by the relative frequency of appearance of mutants in the two populations and the probabilities of fixation of these mutants (46) (details in *SI Appendix*, section S4).

In this paper, we examine the influence of individual rate parameters—mutation rate, generation time, selection strength, and population size—on the outcomes of antagonistic and mutualistic symbioses. For the majority of our analysis, we take the following "all-else equal" approach: For each rate parameter, we hold equal and constant for the two populations all but that parameter and then ask whether the population with the larger value of that parameter is, by some relevant criterion, evolutionarily more successful.

This all-else equal approach is motivated by two considerations: First, it allows for simple mathematical characterization of the long-run dynamics in the weak-mutation limit. Second, it allows us to ask in a clear way, "What is the contribution of a given parameter to the success of the populations?" For illustration, suppose we see that a large population of parasites does

well against a small host population. We might wonder whether it is the parasite's large population size or its rapid generation time that is the main reason for its success. To get at an answer, we would ask what the effect of the parasite's large population size would be if we eliminated any benefit from a faster generation time—i.e., by setting the parasite's generation time equal to its host's. We would then do the same for generation time, by setting the parasite and the host population sizes equal.

This approach ignores possible interactions between rate parameters—for example, if the parasite's larger population size is beneficial only if the parasite has a faster generation time than its host's. For antagonistic symbioses, in the weak-mutation limit, we will be able to relax the all-else equal simplification to prove more general results about the long-run effect of the various rate parameters. For mutualisms, our long-run weak-mutation analytical results do require the all-else equal assumption. We give some numerical suggestion that they hold qualitatively when we relax this assumption.

Antagonistic Symbioses

In this section, we study the evolutionary dynamics of populations interacting according to the antagonistic symbiosis game **1**. An appropriate measure of the relative success of population 1 in a given population state is the proportion of (A, C) and (B, D) matchings (favorable to population 1) minus the proportion of (A, D) and (B, C) matchings (unfavorable to population 1). If this quantity is positive, then population 1 has a larger average payoff than population 2 in that state.

Strong Mutation, Strong Selection. We begin by numerically studying the short- and long-run dynamics, when mutations are not very rare and selection is not very weak.

We first set the two populations' sizes, selection strengths, and mutation rates to be equal ($N_1 = N_2$, $w_1 = w_2$, $\mu_1 = \mu_2$) and vary their relative generation times g_1 and g_2. For the parameter values that we consider, we find a weak Red Queen effect in the short run (Fig. 1*A*): When population 1 has a longer generation time, the set of initial states from which, after 50 generations, population 1 is on average more successful is greater than the set of initial states for which the reverse is true. We find a more pronounced Red Queen effect in the long run, when the dynamics have become independent of where they started (Fig. 1*A*).

We then set the populations' generation times equal ($g_1 = g_2$), as well as their sizes ($N_1 = N_2$) and mutation rates ($\mu_1 = \mu_2$), and allow their selection strengths w_1 and w_2 to vary. We now find very strong Red Queen effects in both the short run and the long run (Fig. 1*C*). We find similar Red Queen effects, especially in the long run, when populations differ in their mutation rates (Fig. 1*B*) or in their sizes (Fig. 1*D*).

We have spoken loosely of the "short run" and "long run." Whereas 50 generations are certainly short run, the long run should be more precisely defined as the time by which the distribution over population states is close to the stationary distribution. In the language of Markov chains, this is known as the "mixing time" of the dynamical process (51).

For small population sizes, the time evolution of the probability distribution over states can feasibly be computed given any starting point. *SI Appendix*, Fig. S6 shows that, in the antagonistic dynamics, for the parameter values used in Fig. 1 (with small populations, of size 50), we may speak of the long run as being after 100–1,000 generations. For larger population sizes, it is not computationally feasible to compute the time evolution of the probability distribution over states, and we must resort to approximate analytical arguments.

We give such arguments in *SI Appendix*, section S3. Here, we summarize their conclusions. We assume that mutations are infrequent. When selection acts weakly in at least one of the populations l ($N_l w_l < 1$), the mixing time is approximately proportional to $1/(\varepsilon\mu_l)$. When selection acts strongly in both populations ($N_1 w_1, N_2 w_2 > 1$), then the mixing time is proportional to

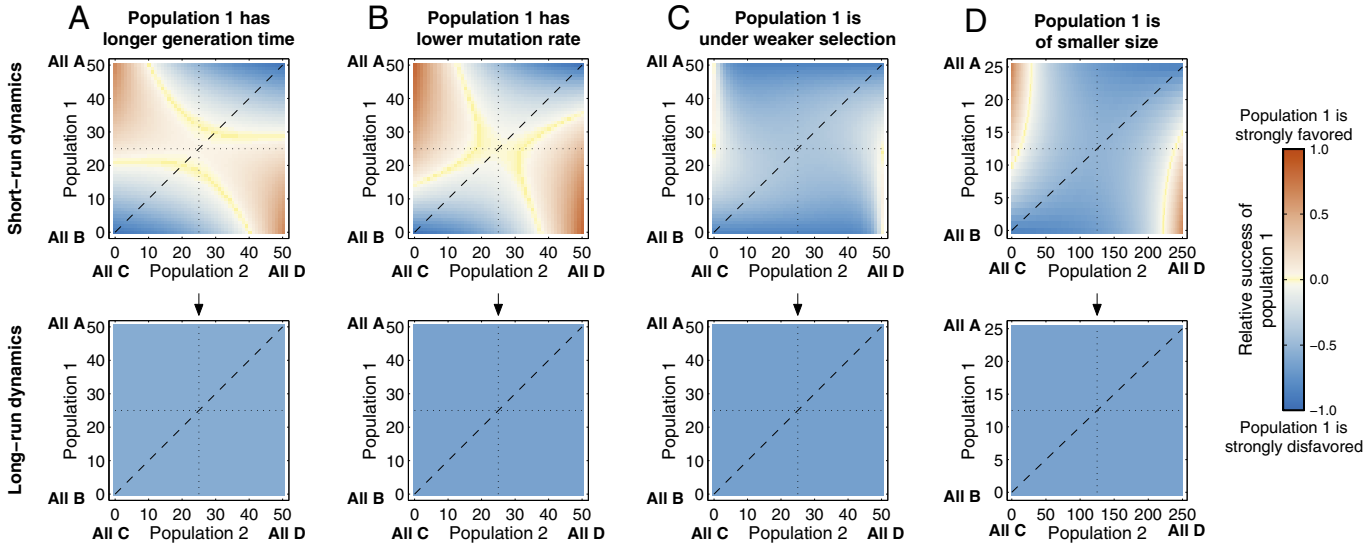

Fig. 1. Evolutionary dynamics of antagonistic symbioses, when population 1 evolves slower than population 2 owing to (*A*) a longer generation time, (*B*) a lower mutation rate, (*C*) weaker selection, and (*D*) a smaller population size. Each panel shows the numerically computed dynamics, assuming that both populations coincide in their evolutionary rate parameters except for the parameter explicitly varied. *A–D*, *Upper* show, for each initial population state, population 1's relative success after 50 generations. Except when the populations differ only in their mutation rates, a short-run Red Queen effect operates in all cases: As the blue area covers more than 50% of the square, population 1's slower evolution leads to its typically being disfavored after 50 generations for a larger fraction of initial population states. The short-run Red Queen effect is strongest for selection strength and population size. *A–D*, *Lower* show population 1's relative success after 50, 000 generations, by which time the starting configuration no longer influences the dynamics: The panels have a uniform color. A strong long-run Red Queen effect is observed in all cases: Population 1's slower evolution causes a larger fraction of long-run time to be spent in states unfavorable to it. Baseline parameters are $N_1 = N_2 = 50$, $w_1 = w_2 = 0.05$, $g_1 = g_2 = 1$, $\mu_1 = \mu_2 = 1$, and $\varepsilon = 0.001$. (*A*) $g_1 = 10$; (*B*) $\mu_1 = 0.1$; (*C*) $w_2 = 0.5$; (*D*) $N_1 = 25$, $N_2 = 250$. A generation is $N_1 + N_2$ elementary time steps of the Moran process.

$1/\min_l(N_l w_l \varepsilon \mu_l)$. In this latter case, the mixing time decreases with increasing population size, owing to a higher substitution rate of beneficial mutations (which drive the arms-race dynamics) in larger populations (*SI Appendix*, section S5.2). That the mixing time does not increase rapidly with increasing population size, and in fact decreases in the case of strong selection, indicates that the stationary distribution will be relevant on realistic timescales.

Weak Mutation. To gain a greater understanding of the above results, we study the analytically tractable case of rare mutations, $\varepsilon \to 0$, using the methodology developed in ref. 46.

In this "weak-mutation limit," the long-run stationary distribution collapses to a distribution over just the four pure states, $\boldsymbol{\lambda} = [\lambda_{(A,C)}, \lambda_{(A,D)}, \lambda_{(B,C)}, \lambda_{(B,D)}]$. The relative success of population 1, as defined above, then simplifies to $(\lambda_{(A,C)} + \lambda_{(B,D)}) - (\lambda_{(A,D)} + \lambda_{(B,C)})$, which, by symmetry of the underlying states $[\lambda_{(A,C)} = \lambda_{(B,D)}, \lambda_{(A,D)} = \lambda_{(B,C)}]$, is proportional to $\lambda_{(A,C)} - \lambda_{(A,D)}$.

First, we examine the influence of generation time and mutation rate on the success of population 1. We fix $N_1 = N_2 = N$ and $w_1 = w_2 = w$ and write $\gamma = (1+w)^{N-1} > 1$. Define $r_1 = (\mu_1/g_1)/(\mu_2/g_2)$, the relative arrival rate of mutations in population 1. The stationary distribution of the evolutionary process, $\boldsymbol{\lambda} = [\lambda_{(A,C)}, \lambda_{(A,D)}, \lambda_{(B,C)}, \lambda_{(B,D)}]$, is then

$$\boldsymbol{\lambda} = \left[\frac{1+\gamma r_1}{r_1+\gamma}, 1, 1, \frac{1+\gamma r_1}{r_1+\gamma}\right] \Big/ \bar{\lambda}, \qquad [4]$$

where $\bar{\lambda}$ ensures that $\boldsymbol{\lambda}$ sums to one (calculations in *SI Appendix*, section S5.4). The relative success of population 1 is

$$\lambda_{(A,C)} - \lambda_{(A,D)} = \frac{(\gamma-1)(r_1-1)}{(r_1+\gamma)\bar{\lambda}}, \qquad [5]$$

which is increasing in r_1. Because both a higher mutation rate and a shorter generation time for population 1 increase r_1, they are both associated with greater evolutionary success in Red Queen interactions. In fact, this result can be shown to hold far more generally: It does not require that the population sizes and selection strengths be set equal for the two populations and holds for all standard evolutionary dynamics, including those exhibiting frequency dependence (39, 40, 47, 52, 53) (proof in *SI Appendix*, section S5.3).

A larger population size and a stronger selection strength of population 1 can also be shown to increase its relative success, $\lambda_{(A,C)} - \lambda_{(A,D)}$, although the results cannot be written as neatly as in Eq. **4** (*SI Appendix*, sections S5.1 and S5.2). Again, these results hold more generally: For many evolutionary processes, including the Moran and Wright–Fisher processes, they hold even if the generation times, mutation rates, and selection strengths are not set equal for the two populations.

Mutualistic Symbioses

In this section, we study the evolutionary dynamics of populations interacting according to the mutualism game **2**. We take the measure of relative success of population 1 in a given population state to be the proportion of (A, C) matchings in that state minus the proportion of (B, D) matchings. If this quantity is positive, then population 1 has a larger average payoff than population 2 in that state.

Strong Mutation, Strong Selection. As we did for the antagonistic game in the previous section, we begin by numerically studying the short- and long-run dynamics, when mutations are not rare and selection is not weak. We consider small and large values of k, specifically $k = 1/2$ and $k = 3/2$.

When $k = 1/2$, and we set the populations' sizes, selection strengths, and mutation rates equal ($N_1 = N_2$, $w_1 = w_2$, $\mu_1 = \mu_2$) and vary their relative generation times g_1 and g_2, we find a weak Red Queen effect in the short run for the parameter values we consider: When population 1 has a longer generation time, the set of initial states from which, after 50 generations, the proportion of (A, C) matchings is expected to be less than that of (B, D) matchings is larger than the set of initial states for which the reverse is true (Fig. 2*A*). This is similar to B&L's (26) general replicator-dynamics result for $k < 1$.

In the long run, however, the average proportion of (A, C) matchings relative to (B, D) matchings becomes independent of where the dynamics started; this effect is clear after 50,000 generations (Fig. 2 *A–D*, *Lower*). In this case, we find that differences

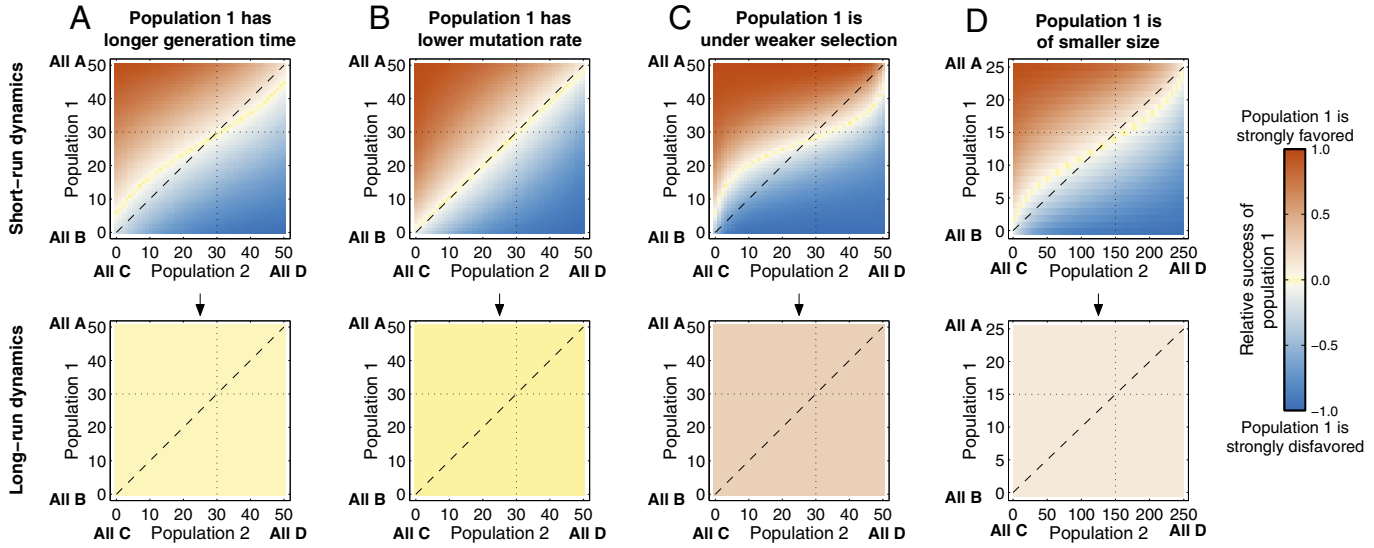

Fig. 2. Evolutionary dynamics of mutualisms when k is small, and population 1 evolves slower than population 2 owing to (*A*) a longer generation time, (*B*) a lower mutation rate, (*C*) weaker selection, and (*D*) a smaller population size. Population parameters are the same as in Fig. 1, and $k = 1/2$. A short-run Red Queen effect is observed for generation time, selection strength (especially), and population size. Relative differences in mutation rate again have no discernible short-run effect. In the long run, differences in mutation rate and generation time again have no effect, whereas a Red King effect is found for selection strength and population size.

in generation time between the two populations have almost no effect on their relative success (Fig. 2*A*).

We also find that differences in the populations' relative mutation rates μ_1 and μ_2, when their other rate parameters are equal, confer almost no advantage to either population, in both the short run and the long run (Fig. 2*B*).

If, instead, we allow the populations' selection strengths w_1 and w_2 to differ, setting their other rate parameters equal, we observe stronger effects. Now, there is a pronounced Red Queen effect in the short run, which reverses to a strong Red King effect in the long run (Fig. 2*C*). A similar result is found when the populations have equal rate parameters except for their population sizes: When population 1 is smaller (so that it evolves slower—see *SI Appendix*, section S5.2 for a precise statement of this), it suffers a small disadvantage in the short run (a Red Queen effect), but an advantage in the long run (a Red King effect) (Fig. 2*D*).

When $k = 3/2$, these results reverse (Fig. 3). We now find a Red King effect in the short run, which is largest when rate differences between the populations derive from differences in selection strength (Fig. 3*C*). This is similar to B&L's (26) general replicator-dynamics result for $k > 1$, and indeed, Fig. 3 *A*, *B*, and *D*, *Upper* and *SI Appendix*, Fig. S1 look very similar to B&L's figure 2 (26). Again, generation time and mutation rate differences have little effect on the long-run dynamics (Fig. 3 *A* and *B*). In contrast, a strong Red Queen effect is found in the long run when the populations differ in their selection strengths (Fig. 3*C*), with population size having a weak Red Queen effect in the long run (Fig. 3*D*).

To summarize, when mutation rates are not very small, we find, for $k = 1/2$, a Red Queen effect in the short run and a Red King effect in the long run; for $k = 3/2$, we find a Red King effect in the short run and a Red Queen effect in the long run. The long-run effects are driven by selection strengths and population sizes; generation times and mutation rates have little effect on the long-run dynamics.

Again, the appropriate definition of the long run is how long it takes for the dynamics to get close to their stationary distribution. This can be computed exactly for small population sizes. In the mutualism game **2**, for the parameter values we have considered in Figs. 2 and 3, *SI Appendix*, Fig. S6 suggests that the long run could be considered any time after about 1,000–10,000 generations.

For larger population sizes, such computations are not feasible, and we must again resort to approximate analytical arguments. These are detailed in *SI Appendix*, section S3. We summarize their conclusions here. Mutations are assumed to be infrequent. When selection acts weakly in both populations ($N_1 w_1, N_2 w_2 < 1$), the mixing time is approximately proportional to $1/\min_l(\varepsilon\mu_l)$. When selection acts strongly in at least one population l, then the mixing time is approximately proportional to $e^{N_l w_l}/(N_l w_l \varepsilon \mu)$.

The exponential term in this last expression is a result of requiring substitutions against selection for the process to mix. It means that the mixing time will be prohibitively long when populations are large and selection acts strongly in them [mixing times increasing exponentially with population size have also been observed in single-population coordination games (54, 55), where transitions between equilibria require evolution against selection too]. In these cases, our stationary distributions will not be empirically relevant; evolution over realistic timescales will involve movement to an equilibrium and then stasis there, as in B&L's (26) analysis. When populations are not large and selection is not very strong, then the stationary distribution will still be reached on a realistic timescale. This will also be true when the effective sizes of the populations are not large (56) or when the populations are subdivided into small subpopulations (57, 58), properties that hold for many mutualistic symbionts (59–61).

Weak Mutation. Again, analytical results can be obtained for the long-run dynamics in the weak-mutation limit ($\varepsilon \to 0$), using the methodology developed in ref. 46. In this limit, the stationary distribution collapses to a distribution over just the pure states, $\boldsymbol{\lambda} = \left[\lambda_{(A,C)}, \lambda_{(A,D)}, \lambda_{(B,C)}, \lambda_{(B,D)}\right]$, and population 1's relative success, as defined above, simplifies to the quantity $\lambda_{(A,C)} - \lambda_{(B,D)}$.

We first examine the influence of generation time and mutation rate. Setting $N_1 = N_2 = N$ and $w_1 = w_2 = w$, but not

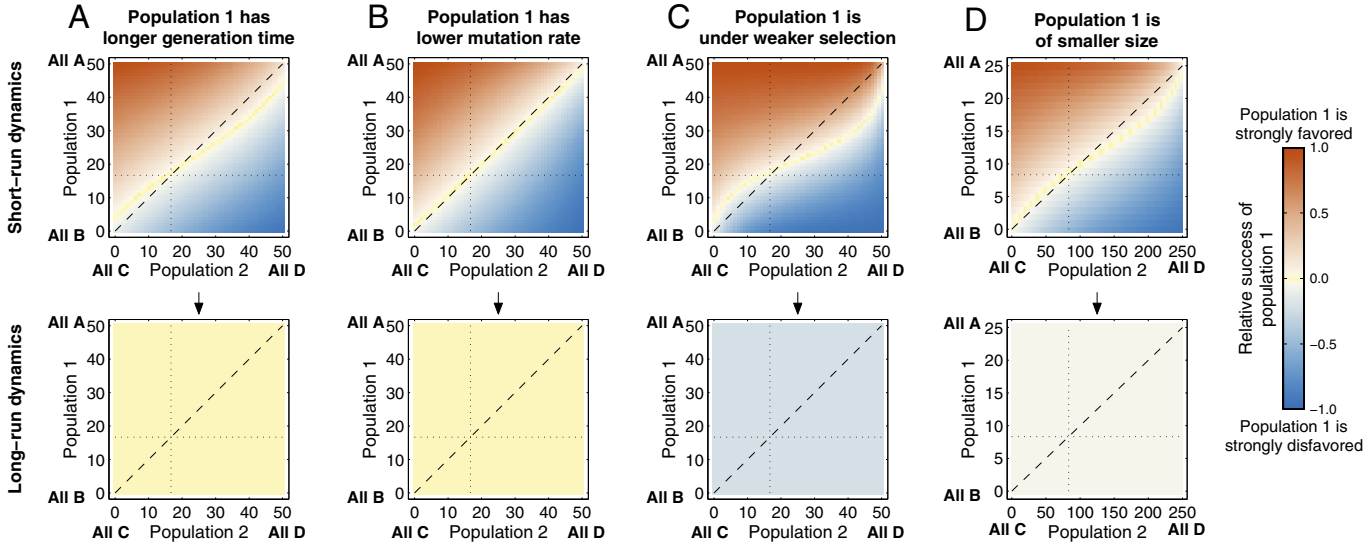

Fig. 3. Evolutionary dynamics of mutualisms when k is large and population 1 evolves slower than population 2 owing to (*A*) a longer generation time, (*B*) a lower mutation rate, (*C*) weaker selection, and (*D*) a smaller population size. Population parameters are the same as in Fig. 2, but now $k = 3/2$. A short-run Red King effect is observed for generation time, selection strength, and population size, with that for selection strength most pronounced. These short-run Red King effects are analogous to those found by B&L (26). Relative differences in mutation rate have no discernible short-run effect. In the long run, differences in mutation rate and generation time have no effect on the relative success of the populations, whereas a Red Queen effect is found for selection strength and population size.

specifying k, g_1, g_2, μ_1, or μ_2, the stationary distribution of the evolutionary process is

$$\boldsymbol{\lambda} = \left[1, \left(\frac{1}{1+w}\right)^{N-1}, \left(\frac{1+kw}{1+2w}\right)^{N-1}, 1\right] \Big/ \bar{\lambda}, \qquad [6]$$

where $\bar{\lambda}$ ensures that $\boldsymbol{\lambda}$ sums to one (calculations in *SI Appendix*, section S6.1). As expected, most time in the long run is spent in the two coordination states (A, C) and (B, D), because $w > 0$ and $0 \leq k < 2$. Less intuitively, for small values of k ($< 1/[1 + w]$), more time is spent in the worse noncoordination state (A, D) than in the better noncoordination state (B, C).

Similar to the numerical results we obtained for the case where mutations are not very rare, neither the generation times nor the relative mutation rates of the two populations have any influence on the stationary distribution $\boldsymbol{\lambda}$ in the weak-mutation limit: If populations 1 and 2 are of the same size, and if natural selection acts equally efficiently in them, then they will be equally successful in the long run, no matter their relative generation times or mutation rates. This weak-mutation result in fact holds for all standard evolutionary processes, including those exhibiting within-population frequency dependence, such as the dynamics that govern Müllerian mimicry (*SI Appendix*, section S6.1).

It turns out (mathematical details in *SI Appendix*, section S6.1) that transitions against selection play an important role in this weak-mutation result. For example, whereas a higher relative mutation rate (or shorter generation time) in population 1 renders the transition from (A, D) to (B, D) more likely than that from (A, D) to (A, C), they also make the reverse transition, from (B, D) to (A, D), more likely than that from (A, C) to (A, D). Because of the symmetry of the fitness changes in these two sets of transitions, these mutation effects (and, similarly, generation effects) cancel out exactly. In words, with a higher mutation rate, population 1 is typically the one that evolves to a coordination equilibrium (in the direction of selection), but it is also typically the one that evolves back out of it (against the direction of selection).

Numerical calculations suggest that this weak-mutation result is broadly robust to inequalities in population sizes and selection strengths (*SI Appendix*, Fig. S10 *C* and *D*).

The directional effects of selection strength and population size limit are more subtle. For selection strength, we equate population 1 and 2's sizes at N and fix population 2's selection strength at $w_2 = w$. Then, if N is sufficiently large, and population 1 has a slightly higher selection strength than population 2, population 1 is less successful when $k < 1/(1 + w)$ and more successful when $k > 1/(1 + w)$ (proofs in *SI Appendix*, section S6.3). Numerical calculations suggest that the "sufficiently large" population size is not very large: Results for $N = 100$ agree with those above (*SI Appendix*, Fig. S10*A*).

The results are similar if we equate the selection strengths in the two populations at w and set population 1's size slightly higher than that of population 2: Population 1 is less successful when $k < 1/(1 + w)$ and more successful when $k > 1/(1 + w)$ (proofs in *SI Appendix*, section S6.4). Again, numerical calculations suggest that these results are robust to inequality in mutation rates and generation times and larger differences in the two populations' sizes when both are sufficiently large (*SI Appendix*, Fig. S10 *B*, *E*, and *F*).

Intuition for these results can be gained by noting that any switch from one coordination state to the other involves two transitions: one transition out of the original equilibrium (against the direction of selection) and a subsequent transition into the new equilibrium state (in the direction of selection). Whether weaker selection (or a smaller population size) favors population 1 depends on which nonequilibrium state is more often passed during these two transitions. This is driven predominantly by the relative probabilities of the transitions against selection, because these probabilities are much more sensitive to population size, selection strength, and the payoffs involved.

Of the transitions against selection, transitions to (B, C) always involve the fixation of mutants of payoff k in populations of payoff 2, whereas transitions to (A, D) always involve the fixation of mutants of payoff 0 in populations of payoff 1.

Therefore, when k is large, the passed nonequilibrium state is usually (B, C) [indeed, note that (B, C) has a higher weight than (A, D) in the stationary distribution **6** precisely when $k > 1/(1 + w)$]. Focusing therefore on transitions through (B, C), if population 1 experiences weaker selection (or is of smaller size), the transition from (A, C) to (B, C) occurs more easily than that from (B, D) to (B, C). So we expect more time to be spent in population 1's disfavored equilibrium state (B, D) in the long run—a Red Queen effect.

Contrariwise, if k is small, the passed nonequilibrium state is usually (A, D). So, if selection is less effective in population 1, more time is spent in its favored equilibrium state (A, C)—a Red King effect—because transitions from this state to (A, D), involving fixation of a selected-against mutant in population 2, are relatively very rare.

In contrast to the minor influence of mutation rates and generation times in the long run, the effects of population size and selection strength are very large. For example, when $N_1 = N_2 = 100$ and $w_2 = 0.15$, a small increase in w_1 from 0.15 to 0.16 increases $\lambda_{(A,C)}/\lambda_{(B,D)}$ (another measure of population 1's relative success, useful for extreme values) by 35% when $k = 1.5$ and decreases it by 30% when $k = 0.8$ (*SI Appendix*, Fig. S10*A*). Changes of similar magnitude are seen when holding $w_1 = w_2 = 0.15$ and increasing N_1 from 100 to 105 (*SI Appendix*, Fig. S10*B*). These effects, caused by small changes in population size and selection strength, are larger than those caused by varying relative mutation rates and generation times across 10 orders of magnitude (*SI Appendix*, Fig. S10 *C* and *D*). Larger changes in population size and selection strength have enormous effects on the long-run relative success of the two populations (*SI Appendix*, Fig. S10 *E* and *F*).

Weak Selection. Returning to the effect of mutation rate and generation time, we have shown analytically that these have no effect on the long-run stationary distribution in the weak-mutation limit. To analytically study the long-run effect of mutation rates when they need not be negligibly small, we make use of another recent advance in evolutionary game theory (47). We impose that the population sizes N, selection strengths w, and generation times g of the two populations be equal. Their mutation rates, $\varepsilon\mu_1$ and $\varepsilon\mu_2$, need not be equal, and we do not impose that $\varepsilon \to 0$. Instead, we assume $w \ll 1$; that is, we assume that selection operates very weakly in the two populations. Long-run Red King and Queen effects consistent with those we have found above for selection strength and population size would have that slower evolution (via a smaller mutation rate) is favored when $k < 1$ and disfavored when $k > 1$ [the threshold for k that was previously relevant, $1/(1 + w)$, tends to 1 as $w \to 0$].

Unlike in the weak-mutation case, the stationary distribution of this process places nonnegligible weight on every possible population state, including those where one or both populations are polymorphic. Again, the relative success of population 1 depends on the long-run proportion of (A, C) and (B, D) matchings, which, following the notation in ref. 47, we denote by $\langle p_A q_C \rangle$ and $\langle p_B q_D \rangle$, respectively. Using equation 29 in ref. 47, we calculate the long-run advantage to population 1 in the weak-selection limit,

$$\langle p_A q_C \rangle - \langle p_B q_D \rangle = A(1 - k)\varepsilon(\mu_2 - \mu_1), \qquad [7]$$

where $A = wN(N - 1)/[8(1 + [N - 1]\varepsilon\mu_1)(1 + [N - 1]\varepsilon\mu_2)] > 0$ (calculations in *SI Appendix*, section S7). So, when $k < 1$, population 1 does better when it evolves slower ($\mu_1 < \mu_2$); when $k > 1$,

population 1 does better when it evolves faster ($\mu_1 > \mu_2$). We have thus recovered long-run Red King and Queen effects of mutation consistent with those for selection strength and population size. When mutation rates are small ($\varepsilon \ll 1$), the two populations do approximately equally well, consistent with our weak-mutation results above.

Discussion

We have placed the evolution of symbioses, both antagonistic and mutualistic, into a simple finite-population model and studied the effect of parameters that influence the rate of evolution of the participating populations on those populations' relative success.

Among these "rate parameters," mutation rate and generation time are perhaps the most responsive (34, 62–68). Selection strengths, although also a clear determinant of evolutionary rate, seem more of a fixed property of an interaction, but could change over time if one population reduces its dependence on the interaction (3) or actively increases the dependence of its interactant (69). Population size is the most complicated of the rate parameters, because changes in the relative success of interacting populations (possibly driven by differences in their population sizes) are expected to manifest themselves in subsequent changes in the populations' sizes (70).

It is important to realize that, although we have studied differences in the relative success of interacting populations caused by different evolutionary rate parameters, our results do not say anything directly about how the rate parameters themselves should evolve. For example, it is conceivable that a slower generation time leads to greater success for a population (i.e., that a Red King effect holds), but that a faster generation time is selected for within that population. In this example, slow evolution could be interpreted as a public good.

This has important implications for how results such as ours and those of B&L (26) should be interpreted in studies of molecular evolution (e.g., ref. 71). Although generation times and mutation rates determine the rate of nucleotide substitution at neutral sites (37), were we to find that a slower generation time or a lower mutation rate leads to greater success for a population in a mutualism, this would not necessarily imply that species involved in mutualisms should exhibit lower neutral substitution rates.

Examining the evolution of the rate parameters themselves would require a more complicated model than the one we have studied. In a population genetic setting, one could posit in each population two distinct genetic loci: a "strategy locus," the alleles at which determine which game strategies their bearers play, and a "rate locus," the alleles at which determine the value of a rate parameter such as mutation rate. The evolutionary dynamics would then inform how both the success of the populations and their evolutionary rates evolve. One setting where we might expect the evolutionary rate within a population to evolve to increase that population's relative success is if the interacting populations are subdivided into many isolated symbioses (on "islands"), between which migration occurs [this model is similar to the "haystack model" considered by B&L (26)]. A rate parameter value that improves a subpopulation's success on an island causes that subpopulation to send out more migrants—who bear the rate parameter value—to colonize other islands. Thus, in symbioses with a high degree of population structure, we might expect the results we have found in this paper to be informative of how evolutionary rate will evolve within populations.

In antagonistic symbioses, we have shown that faster evolution through any of the rate parameters leads to greater evolutionary success in our model, in both the short run and the long run—a clear Red Queen effect. Mutation rate and generation time play similar roles in determining the long-run outcome of antagonistic interactions when mutations are infrequent. As far as our results have implications for evolution within populations, this suggests that one population might compensate for a longer generation time by evolving an elevated mutation rate. Consistent with this, bacterial strains subjected to antagonistic interactions with bacteriophages often exhibit much higher mutation rates than control strains (72). A similar explanation has been suggested for the finding that, in mammals, generation time is positively correlated with the rate of intrachromosomal recombination (73), which, like mutation, is a generator of variation (1). With particular respect to selection strength, the "life–dinner principle" (3) holds in our model: In an antagonistic interaction, the population for whom the interaction matters more (the rabbit, not the fox) is evolutionarily more successful.

In mutualistic symbioses, we have uncovered an important influence of evolutionary timescale on the relative success of interacting populations, with results in the short run and long run often being in opposition. In the short run, the stochastic evolutionary dynamics that we have studied are similar to the deterministic dynamics studied by B&L (26), and our short-run results replicate theirs: We find a short-run Red Queen effect when k is small in game **2** and a short-run Red King effect when k is large. Our analysis extends B&L's (26) by allowing us to determine which of the biological rate parameters drive these effects: to wit, generation times, population sizes, and (especially) selection strengths.

In the long run, we find a Red King effect when k is small and a Red Queen effect when k is large, contrary to the short-run results. Among the rate parameters that could drive these long-run results, mutation rate and generation time in fact have little to no effect, unless mutation is frequent. This is a surprising result, given that mutation rate and generation time are perhaps the most prominent determinants of evolutionary rate. This result depends on the possibility of evolution against selection, always present in stochastic evolutionary dynamics. It could not be discovered using deterministic evolutionary dynamics, because in such dynamics, evolution always proceeds in the direction of selection. Deterministic dynamics admit little notion of the strength of selection: Selection is either "positive" or "negative." In reality, selection is a continuum—as the Red Queen said to Alice, "*I* could show you hills, in comparison with which you'd call that a valley" (ref. 4, p. 37)—and deleterious mutations do sometimes fix because of random drift (37, 74–77).

The long-run Red King and Queen effects in mutualisms, for small k and large k, respectively, instead operate predominantly in our model through the efficiency of selection (which increases with selection strength and population size), not the generation of variation (mutation and generation rates).

Which results, short-run or long-run, are relevant to a specific case depends on which evolutionary timescale is appropriate. Again, this question is treated quantitatively in *SI Appendix*, section S3. In mutualisms, the long-run dynamics are characterized by transitions between the equilibria—involving evolution first against, and then with, selection—so the applicability of our long-run results depends on the timescale of these transitions. When both populations are very large and under strong selection, the time it takes for a transition between equilibria to occur can be so long as to be empirically irrelevant; here, the short-run dynamics are of more interest. When the populations have small effective population sizes (owing to small census size or population structure, for example) or selection acts weakly in them, equilibrium transitions can be frequent, and the long-run dynamics apply.

In our analysis, we have made several modeling decisions and simplifying assumptions, described below.

We have assumed that interactions are pairwise, for simplicity. The consideration of multiplayer games (2, 78–83) is a desirable extension and has been shown to influence the outcome of mutualism games in infinite populations (80) and antagonism and Snowdrift games in finite populations (2).

In deriving our main analytical results, we assumed that mutations are infrequent. This is a realistic assumption in many cases, but not always. *SI Appendix*, Fig. S5 shows that mutation rates in the two populations do not have to be exceedingly small for our weak-mutation analytical results to be an excellent quantitative match (46, 84, 85), and in *SI Appendix*, Figs. S2–S5 suggest that, even for very large mutation rates, the long-run relative success of the populations qualitatively matches our weak-mutation results in most cases. The only notable exceptions, all for the mutualism game **2**, are (*i*) for large k, when the populations have different selection strengths, a strong long-run Red Queen effect for small mutation rates reverses to a weak Red King effect when mutation rates are very large (*SI Appendix*, Figs. S2*C* and S5*C*); (*ii*) when the populations have different mutation rates and both their mutation rates are very large, then, contrary to our weak-mutation result where differences in mutation rates have no effect, we find a long-run Red King effect for small k (*SI Appendix*, Figs. S3*B* and S5*B*) and a Red Queen effect for large k (*SI Appendix*, Figs. S2*B* and S5*C*) (exception *ii* is consistent with our weak-selection result, Eq. **7**); and (*iii*) for small k, generation time has a Red Queen effect for intermediate and large mutation rates (*SI Appendix*, Figs. S3*A* and S5*B*). A very recent methodological advance in stochastic evolutionary dynamics (86) suggests that tractable analysis of the stationary distribution outside the weak-mutation regime may soon be possible.

In finite-population evolutionary game theory, game payoffs π must be translated to positive fitnesses f (40), with this translation calibrated by a strength of selection w. We have used the commonly used linear translation $f = 1 + w\pi$, but others are also possible. The exponential translation $f = \exp(w\pi)$ has the advantage that it guarantees fitnesses always to be positive, no matter the range of game payoffs, and also sometimes allows for neater characterization of long-run dynamics (87–90). Results are expected to be similar for the two translations, and they are identical in the weak-selection limit $w \to 0$. *SI Appendix*, Fig. S5 suggests that our results are essentially unchanged when we use either translation.

The dynamics in our model, based on the Moran process, are highly stylized, especially in their assumption of constant population sizes. It would be interesting to study evolution in games **1** and **2**, using more complicated intergenerational population dynamics (91–93).

Finally, the games themselves are particular simplifications of more complex interactions. One key simplification in games **1** and **2** is that strategies are discrete: A, B, C, D. This will be relevant for many antagonistic symbioses; for example, it is probably a reasonable summary that a pathogen is either resistant or not to a host's defenses. For mutualisms, we have motivated the discrete game **2** with examples that exhibit such discreteness, such as division of labor and Müllerian mimicry. However, many mutualisms are likely to involve continuous strategies: for example, how much energy to expend on a cooperative task, as in the degree to which an ant colony protects from herbivores the plant that houses and feeds it (94) and in turn the energy the plant dedicates to housing and feeding the ants. For such cases, an alternative, continuous-strategy model specification is appropriate.

An attractive option is a simple modification of a Nash bargaining game (95). Players 1 and 2 choose activity levels x and y. If $x + y > 1$, then both receive zero payoff. If $x + y \leq 1$, then player 1 receives payoff $\alpha x + (1 - \alpha)y$, and player 2 receives $\alpha y + (1 - \alpha)x$. $\alpha \in (0, 1)$ calibrates the degree to which players' payoffs depend on their own actions. As long as $x + y < 1$, both players' payoffs increase if either player increases its activity, so that the game is mutualistic. But when α is small, each player would prefer to have a lower activity than the other, whereas when α is large, each would prefer to have the larger activity. In contrast to the discrete-strategy mutualism game, where there were two discrete equilibria, here there is a continuous path of equilibria (the line $x + y = 1$). In *SI Appendix*, section S8, we study the coevolutionary dynamics of populations of player 1s and 2s. The short-run dynamics in this game involve evolution to the equilibrium line, whereas the long-run dynamics involve drift-like movement along and around it.

In both the short and long runs, the faster-evolving population is at an advantage when α is large (a Red Queen effect; *SI Appendix*, Fig. S11), but at a disadvantage when α is small (a Red King effect; *SI Appendix*, Fig. S12). The short-run effects are driven by all rate parameters. But similar to what we found for the discrete mutualism game, the long-run effects are driven predominantly by selection strength and population size—differences in mutation rate and generation time have little effect. This is because selection strength and population size contribute differentially here to a case of drift-induced selection: Stochastic jumps off the equilibrium path return with an average directional bias. Drift-induced selection is a newly recognized phenomenon that has recently gained attention in the stochastic dynamics literature (96–98). It can be studied analytically using diffusion approximations and separations of timescales (99). Such a study in the context of continuous-strategy mutualisms would be an important extension of our preliminary analysis.

In this vein, studying the finite-population dynamics of other games in which slow evolution has been hypothesized to be beneficial (100) would also be desirable.

ACKNOWLEDGMENTS. C.V. is grateful to Philipp Altrock for introducing him to the Red King literature. We thank Carl Bergstrom, Kirill Borusyak, Kooskia Burns-Edelman, Richard Childers, Michael Doebeli, David Haig, Jonathan Libgober, Alex McAvoy, Sam Sinai, Luis Zaman, and participants at the Harvard Economics Department's Games and Markets lunch for helpful comments.

1. Nee S (1989) Antagonistic co-evolution and the evolution of genotypic randomization. *J Theor Biol* 140:499–518.
2. Damore J, Gore J (2011) A slowly evolving host moves first in symbiotic interactions. *Evolution* 65:2391–2398.
3. Dawkins R, Krebs J (1979) Arms races between and within species. *Proc Biol Sci* 205:489–511.
4. Carroll L (1871) *Through the Looking-Glass* (Macmillan, London).
5. Turner J (1987) The evolutionary dynamics of Batesian and Muellerian mimicry: Similarities and differences. *Ecol Entomol* 12:81–95.
6. Burt A, Trivers R (2006) *Genes in Conflict* (Harvard Univ Press, Cambridge, MA).
7. McLaughlin RN, Malik HS (2017) Genetic conflicts: The usual suspects and beyond. *J Exp Biol* 220:6–17.
8. Tao Y, Masly J, Araripe L, Ke Y, Hartl D (2007) A sex-ratio meiotic drive system in *Drosophila simulans*. I: An autosomal suppressor. *PLoS Biol* 5:e292.
9. Tao Y, et al. (2007) A sex-ratio meiotic drive system in *Drosophila simulans*. II: An X-linked distorter. *PLoS Biol* 5:e293.
10. Henikoff S, Ahmad K, Malik H (2001) The centromere paradox: Stable inheritance with rapidly evolving DNA. *Science* 293:1098–1102.
11. Malik H, Henikoff S (2002) Conflict begets complexity: The evolution of centromeres. *Curr Opin Genet Dev* 12:711–718.
12. Boulton A, Myers R, Redfield R (1997) The hotspot conversion paradox and the evolution of meiotic recombination. *Proc Natl Acad Sci USA* 94:8058–8063.
13. Baudat F, et al. (2010) PRDM9 is a major determinant of meiotic recombination hotspots in humans and mice. *Science* 327:836–840.
14. Parvanov E, Petkov P, Paigen K (2010) *Prdm9* controls activation of mammalian recombination hotspots. *Science* 327:835–835.
15. Úbeda F, Wilkins J (2011) The Red Queen theory of recombination hotspots. *J Evol Biol* 24:541–553.
16. Kingan SB, Garrigan D, Hartl DL (2010) Recurrent selection on the Winters sex-ratio genes in *Drosophila simulans*. *Genetics* 184:253–265.
17. Myers S, et al. (2010) Drive against hotspot motifs in primates implicates the *PRDM9* gene in meiotic recombination. *Science* 327:876–879.
18. Ponting C (2011) What are the genomic drivers of the rapid evolution of PRDM9?. *Trends Genet* 27:165–171.
19. Jeffreys AJ, Cotton VE, Neumann R, Lam KW (2013) Recombination regulator PRDM9 influences the instability of its own coding sequence in humans. *Proc Natl Acad Sci USA* 110:600–605.
20. Schwartz J, Roach D, Thomas J, Shendure J (2014) Primate evolution of the recombination regulator PRDM9. *Nat Commun* 5:4370.

21. Lesecque Y, Glémin S, Lartillot N, Mouchiroud D, Duret L (2014) The Red Queen model of recombination hotspots evolution in the light of archaic and modern human genomes. *PLoS Genet* 10:e1004790.
22. Davies B, et al. (2016) Re-engineering the zinc fingers of PRDM9 reverses hybrid sterility in mice. *Nature* 530:171–176.
23. Baker Z, et al. (2017) Repeated losses of PRDM9-directed recombination despite the conservation of PRDM9 across vertebrates. *eLife* 6:e24133.
24. Doebeli M, Knowlton N (1998) The evolution of interspecific mutualisms. *Proc Natl Acad Sci USA* 95:8676–8680.
25. Herre E, Knowlton N, Mueller U, Rehner S (1999) The evolution of mutualisms: Exploring the paths between conflict and cooperation. *Trends Ecol Evol* 14:49–53.
26. Bergstrom C, Lachmann M (2003) The Red King effect: When the slowest runner wins the coevolutionary race. *Proc Natl Acad Sci USA* 100:593–598.
27. Leigh E (2010) The evolution of mutualism. *J Evol Biol* 23:2507–2528.
28. Hom E, Aiyar P, Schaeme D, Mittag M, Sasso S (2015) A chemical perspective on microalgal–microbial interactions. *Trends Plant Sci* 20:689–693.
29. Haig D, Grafen A (1991) Genetic scrambling as a defence against meiotic drive. *J Theor Biol* 153:531–558.
30. Krugman P (1979) Increasing returns, monopolistic competition, and international trade. *J Int Econ* 9:469–479.
31. Wahl L (2002) Evolving the division of labour: Generalists, specialists and task allocation. *J Theor Biol* 219:371–388.
32. Sherratt T (2008) The evolution of Müllerian mimicry. *Naturwissenschaften* 95: 681–695.
33. Taddei F, et al. (1997) Role of mutator alleles in adaptive evolution. *Nature* 387: 700–702.
34. Sniegowski P, Gerrish P, Johnson T, Shaver A (2000) The evolution of mutation rates: Separating causes from consequences. *BioEssays* 22:1057–1066.
35. Lanfear R, Kokko H, Eyre-Walker A (2014) Population size and the rate of evolution. *Trends Ecol Evol* 29:33–41.
36. Van Valen L (1973) A new evolutionary law. *Evol Theor* 1:1–30.
37. Kimura M (1983) *The Neutral Theory of Molecular Evolution* (Cambridge Univ Press, Cambridge, UK).
38. Moran P (1958) Random processes in genetics. *Math Proc Camb Philos Soc* 54:60–71.
39. Taylor C, Fudenberg D, Sasaki A, Nowak M (2004) Evolutionary game dynamics in finite populations. *Bull Math Biol* 66:1621–1644.
40. Nowak M (2006) *Evolutionary Dynamics* (Harvard Univ Press, Cambridge, MA).
41. Hartl DL, Clark AG (2007) *Principles of Population Genetics* (Sinauer, Sunderland, MA), 4th Ed.
42. Fudenberg D, Imhof LA (2006) Imitation processes with small mutations. *J Econ Theor* 131:251–262.
43. Hauert C, Traulsen A, De Silva H, Nowak M, Sigmund K (2008) Public goods with punishment and abstaining in finite and infinite populations. *Biol Theor* 3:114–122.
44. Fudenberg D, Nowak MA, Taylor C, Imhof LA (2006) Evolutionary game dynamics in finite populations with strong selection and weak mutation. *Theor Popul Biol* 70: 352–363.
45. Karlin S, Taylor H (1975) *A First Course in Stochastic Processes* (Academic, New York).
46. Veller C, Hayward L (2016) Finite-population evolution with rare mutations in asymmetric games. *J Econ Theor* 162:93–113.
47. Ohtsuki H (2010) Stochastic evolutionary dynamics of bimatrix games. *J Theor Biol* 264:136–142.
48. McCandlish D, Stoltzfus A (2014) Modeling evolution using the probability of fixation: History and implications. *Q Rev Biol* 89:225–252.
49. Waxman D, Gavrilets S (2005) 20 questions on adaptive dynamics. *J Evol Biol* 18: 1139–1154.
50. Lynch M (2010) Evolution of the mutation rate. *Trends Genet* 26:345–352.
51. Levin D, Peres Y, Wilmer E (2009) *Markov Chains and Mixing Times* (Am Math Soc, Providence, RI).
52. Nowak M, Sasaki A, Taylor C, Fudenberg D (2004) Emergence of cooperation and evolutionary stability in finite populations. *Nature* 428:646–650.
53. Sekiguchi T, Ohtsuki H (2015) Fixation probabilities of strategies for bimatrix games in finite populations. *Dyn Games Appl* 7:93–111.
54. Black A, Traulsen A, Galla T (2012) Mixing times in evolutionary game dynamics. *Phys Rev Lett* 109:028101.
55. Kandori M, Mailath G, Rob R (1993) Learning, mutation, and long run equilibria in games. *Econometrica* 61:29–56.
56. Charlesworth B (2009) Fundamental concepts in genetics: Effective population size and patterns of molecular evolution and variation. *Nat Rev Genet* 10:195–205.
57. Whitlock MC (2003) Fixation probability and time in subdivided populations. *Genetics* 164:767–779.
58. Ellison G (1993) Learning, local interaction, and coordination. *Econometrica* 61: 1047–1071.
59. Moran NA, Wernegreen JJ (2000) Lifestyle evolution in symbiotic bacteria: Insights from genomics. *Trends Ecol Evol* 15:321–326.
60. Cook J, Rasplus J (2003) Mutualists with attitude: Coevolving fig wasps and figs. *Trends Ecol Evol* 18:241–248.
61. Caldera E, Currie C (2012) The population structure of antibiotic-producing bacterial symbionts of *Apterostigma dentigerum* ants: Impacts of coevolution and multipartite symbiosis. *Am Nat* 180:604–617.
62. Sniegowski P, Gerrish P, Lenski R (1997) Evolution of high mutation rates in experimental populations of *E. coli*. *Nature* 387:703–705.
63. Williams G (1957) Pleiotropy, natural selection, and the evolution of senescence. *Evolution* 11:398–411.
64. Hamilton W (1966) The moulding of senescence by natural selection. *J Theor Biol* 12:12–45.
65. Bull J (1994) Virulence. *Evolution* 48:1423–1437.
66. Nowak M, May R (1994) Superinfection and the evolution of parasite virulence. *Proc Biol Sci* 255:81–89.
67. Stearns S, Ackermann M, Doebeli M, Kaiser M (2000) Experimental evolution of aging, growth, and reproduction in fruitflies. *Proc Natl Acad Sci USA* 97:3309–3313.
68. Vérin M, Menu F, Rajon E (2016) The biased evolution of generation time. arXiv:1502.05508.
69. Wyatt G, Kiers E, Gardner A, West S (2016) Restricting mutualistic partners to enforce trade reliance. *Nat Comm* 7:10322.
70. Schenk H, Traulsen A, Gokhale C (2016) Chaotic provinces in the kingdom of the red queen. arXiv:1607.01564.
71. Rubin B, Moreau C (2016) Comparative genomics reveals convergent rates of evolution in ant–plant mutualisms. *Nat Commun* 7:12679.
72. Pal C, Maciá M, Oliver A, Schachar I, Buckling A (2007) Coevolution with viruses drives the evolution of bacterial mutation rates. *Nature* 450:1079–1081.
73. Burt A, Bell G (1987) Mammalian chiasma frequencies as a test of two theories of recombination. *Nature* 326:803–805.
74. Ohta T (1973) Slightly deleterious mutant substitutions in evolution. *Nature* 246: 96–98.
75. Ohta T (1992) The nearly neutral theory of molecular evolution. *Annu Rev Ecol Systemat* 23:263–286.
76. Lynch M (2007) *The Origins of Genome Architecture* (Sinauer, Sunderland, MA).
77. McCandlish D, Epstein C, Plotkin J (2014) The inevitability of unconditionally deleterious substitutions during adaptation. *Evolution* 68:1351–1364.
78. Gokhale C, Traulsen A (2010) Evolutionary games in the multiverse. *Proc Natl Acad Sci USA* 107:5500–5504.
79. Gokhale C, Traulsen A (2011) Strategy abundance in evolutionary many-player games with multiple strategies. *J Theor Biol* 283:180–191.
80. Gokhale C, Traulsen A (2012) Mutualism and evolutionary multiplayer games: Revisiting the red king. *Proc Biol Sci* 279:4611–4616.
81. Wu B, Traulsen A, Gokhale C (2013) Dynamic properties of evolutionary multi-player games in finite populations. *Games* 4:182–199.
82. Gokhale C, Traulsen A (2014) Evolutionary multiplayer games. *Dyn Games Appl* 4: 468–488.
83. McAvoy A, Hauert C (2016) Structure coefficients and strategy selection in multiplayer games. *J Math Biol* 72:203–238.
84. Wu B, Gokhale C, Wang L, Traulsen A (2012) How small are small mutation rates? *J Math Biol* 64:803–827.
85. Rand D, Tarnita C, Ohtsuki H, Nowak M (2013) Evolution of fairness in the one-shot anonymous Ultimatum Game. *Proc Natl Acad Sci USA* 110:2581–2586.
86. Vasconcelos V, Santos F, Santos F, Pacheco J (2017) Stochastic dynamics through hierarchically embedded Markov chains. *Phys Rev Lett* 118:058301.
87. Blume L (1993) The statistical mechanics of strategic interaction. *Games Econ Behav* 5:387–424.
88. Traulsen A, Nowak M, Pacheco J (2006) Stochastic dynamics of invasion and fixation. *Phys Rev E* 74:011909.
89. Traulsen A, Shoresh N, Nowak M (2008) Analytical results for individual and group selection of any intensity. *Bull Math Biol* 70:1410–1424.
90. Cooney D, Allen B, Veller C (2016) Assortment and the evolution of cooperation in a Moran process with exponential fitness. *J Theor Biol* 409:38–46.
91. Gokhale C, Papkou A, Traulsen A, Schulenburg H (2013) Lotka–Volterra dynamics kills the red queen: Population size fluctuations and associated stochasticity dramatically change host-parasite coevolution. *BMC Evol Biol* 13:254.
92. Song Y, Gokhale C, Papkou A, Schulenburg H, Traulsen A (2015) Host-parasite coevolution in populations of constant and variable size. *BMC Evol Biol* 15: 212.
93. Papkou A, Gokhale C, Traulsen A, Schulenburg H (2016) Host–parasite coevolution: Why changing population size matters. *Zoology* 119:330–338.
94. Frederickson M, et al. (2012) The direct and ecological costs of an ant-plant symbiosis. *Am Nat* 179:768–778.
95. Nash J, Jr (1950) The bargaining problem. *Econometrica* 18:155–162.
96. Constable G, McKane A, Rogers T (2013) Stochastic dynamics on slow manifolds. *J Phys Math Theor* 46:295002.
97. Constable G, Rogers T, McKane A, Tarnita C (2016) Demographic noise can reverse the direction of deterministic selection. *Proc Natl Acad Sci USA* 113: E4745–E4754.
98. Veller C, Muralidhar P, Constable G, Nowak M (2017) Drift-induced selection between male and female heterogamety. bioRxiv:141929.
99. Parsons T, Rogers T (2015) Dimension reduction via timescale separation in stochastic dynamical systems. arXiv:1510.07031.
100. Hilbe C, Nowak M, Sigmund K (2013) Evolution of extortion in Iterated Prisoner's Dilemma games. *Proc Natl Acad Sci USA* 110:6913–6918.

SI Appendix: The Red Queen and King in finite populations

Carl Veller[*,†] Laura Hayward[‡] Christian Hilbe[§] Martin A. Nowak[*,†,¶]

Contents

[*]Department of Organismic and Evolutionary Biology, Harvard University, Massachusetts, U.S.A.
[†]Program for Evolutionary Dynamics, Harvard University, Massachusetts, U.S.A.
[‡]Department of Mathematics, Columbia University, New York, U.S.A.
[§]IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria
[¶]Department of Mathematics, Harvard University, Massachusetts, U.S.A.

S1 Transition probabilities and time evolution of symbiotic dynamics

Here, we describe our full dynamical model of symbiosis evolution, calculate the probabilities of transitions from state to state, and use these to calculate the time evolution of the probability distribution over population states, given some initial distribution.

We suppose that the symbiosis is characterized by pairwise interactions according to the following asymmetric game:

		Player 2	
		C	D
Player 1	A	α_{AC}, β_{AC}	α_{AD}, β_{AD}
	B	α_{BC}, β_{BC}	α_{BD}, β_{BD}

(1)

Populations 1 and 2 (composed of player 1s and player 2s respectively) are of sizes N_1 and N_2, and have relative generation times g_1 and g_2. Each time-step, each individual receives its average payoff from interacting with a random member of the other population (each equally likely). If, in a given time-step, i members of population 1 play A, and j members of population 2 play C, then we may describe the population state in that time-step simply as (i, j), and the average payoff to individuals playing strategies A, B, C, and D are:

$$\pi_A(i,j) = \frac{j}{N_2}\alpha_{AC} + \frac{N_2 - j}{N_2}\alpha_{AD};$$
$$\pi_B(i,j) = \frac{j}{N_2}\alpha_{BC} + \frac{N_2 - j}{N_2}\alpha_{BD};$$
$$\pi_C(i,j) = \frac{i}{N_1}\beta_{AC} + \frac{N_1 - i}{N_1}\beta_{BC};$$
$$\pi_D(i,j) = \frac{i}{N_1}\beta_{AD} + \frac{N_1 - i}{N_1}\beta_{BD}.$$

If an individual in population l receives average payoff π, this translates to a positive fitness f via $f = 1 + w_l\pi$, so that $w_1, w_2 > 0$ calibrate the strength of selection in the two populations. The fitnesses of individuals playing the various strategies in population state (i, j) are $f_A(i,j) = 1 + w_1\pi_A(i,j)$, $f_B(i,j) = 1 + w_1\pi_B(i,j)$, $f_C(i,j) = 1 + w_2\pi_C(i,j)$, and $f_D(i,j) = 1 + w_2\pi_D(i,j)$

In each elementary time step, exactly one birth-death event occurs between the two populations. With probability $\frac{N_1/g_1}{N_1/g_1+N_2/g_2}$ this is in population 1, and with probability $\frac{N_2/g_2}{N_1/g_1+N_2/g_2}$ it is in population 2. If a birth-death event occurs in population l in a given time-step, then one individual in population l is chosen to reproduce, with probability proportional to fitness, and one individual is chosen to die, with each equally likely. The

same individual can be chosen to reproduce and die. The reproducing individual produces an offspring, which replaces the individual chosen to die. With probability $1-\varepsilon\mu_l$, the offspring inherits the strategy its parent plays, and with probability $\varepsilon\mu_l$ the offspring plays the other strategy instead (it 'mutates').

With populations of size N_1 and N_2, there are $n=(N_1+1)(N_2+1)$ possible populations states (i,j), which can be given some (arbitrary) enumeration $1,2,\ldots,n$. Let $P_{(i,j)\to(i',j')}$ be the probability that the system moves in one time-step from state (i,j) to step (i',j'). These one-step transition probabilities are:

$$P_{(i,j)\to(i+1,j)} = \frac{N_1/g_1}{N_1/g_1+N_2/g_2}\cdot\left[(1-\varepsilon\mu_1)\frac{if_A(i,j)}{if_A(i,j)+(N_1-i)f_B(i,j)}\frac{N_1-i}{N_1}+\varepsilon\mu_1\frac{N_1-i}{N_1}\right];$$
$$P_{(i,j)\to(i-1,j)} = \frac{N_1/g_1}{N_1/g_1+N_2/g_2}\cdot\left[(1-\varepsilon\mu_1)\frac{(N_1-i)f_B(i,j)}{if_A(i,j)+(N_1-i)f_B(i,j)}\frac{i}{N_1}+\varepsilon\mu_1\frac{i}{N_1}\right];$$
$$P_{(i,j)\to(i,j+1)} = \frac{N_2/g_2}{N_1/g_1+N_2/g_2}\cdot\left[(1-\varepsilon\mu_2)\frac{jf_C(i,j)}{jf_C(i,j)+(N_2-j)f_D(i,j)}\frac{N_2-j}{N_2}+\varepsilon\mu_2\frac{N_2-j}{N_2}\right];$$
$$P_{(i,j)\to(i,j-1)} = \frac{N_2/g_2}{N_1/g_1+N_2/g_2}\cdot\left[(1-\varepsilon\mu_2)\frac{(N_2-j)f_D(i,j)}{jf_C(i,j)+(N_2-j)f_D(i,j)}\frac{j}{N_2}+\varepsilon\mu_2\frac{j}{N_2}\right];$$

with $P_{(i,j)\to(i,j)} = 1-P_{(i,j)\to(i+1,j)}-P_{(i,j)\to(i-1,j)}-P_{(i,j)\to(i,j+1)}-P_{(i,j)\to(i,j-1)}$, and $P_{(i,j)\to(i',j')}=0$ for all other states (i',j'). We refer to the $n\times n$ matrix P that contains these one-step transition probabilities as the transition matrix of the complete Markov chain.

Given some initial probability distribution over the set of population states, $\mathbf{v}^0=\left(v^0_{(i,j)}\right)$, we can compute the distribution over population states after t time-steps by taking t successive powers of the transition matrix P:

$$\mathbf{v}^t=\mathbf{v}^0P^t. \tag{2}$$

Note that $\mathbf{v}^t$ is a stochastic vector of length n, respecting the particular enumeration $1,2,\ldots,n$ chosen for the n population states.

The stationary distribution, $\mathbf{v}=\lim_{t\to\infty}\mathbf{v}^t$, is the unique stochastic vector that solves, for each state (i,j),

$$v_{(i,j)}=\sum_{(i',j')}v_{(i',j')}P_{(i',j')\to(i,j)}. \tag{3}$$

S1.1 Mutualistic symbioses

The payoff matrix of the mutualism game [game (2) in the Main Text] is

		Player 2	
		C	D
Player 1	A	$2,1$	$0,0$
	B	k,k	$1,2$

(4)

Among the two coordination matchings, (A, C) and (B, D), population 1 prefers (A, C) and population 2 prefers (B, D). The other matchings give the members of populations 1 and 2 equal payoff. Therefore, given a population state (i, j), a measure of the success of population 1 in that state is simply the average proportion of (A, C) matchings in that state minus the average proportion of (B, D) matchings: $\frac{ij}{N_1 N_2} - \frac{(N_1-i)(N_2-j)}{N_1 N_2}$. Given any distribution over population states $\mathbf{p}$, we may calculate the expected value of population 1's success, $\sum_{(i,j)} p_{(i,j)} \cdot \left(\frac{ij}{N_1 N_2} - \frac{(N_1-i)(N_2-j)}{N_1 N_2}\right)$.

Therefore, given an initial distribution over population states, $\mathbf{v}^0 = \left(v^0_{(i,j)}\right)$, making use of Eq. (2), the expected value of population 1's relative success after t time-steps is:

$$\sum_{(i,j)} v^t_{(i,j)} \cdot \left(\frac{ij}{N_1 N_2} - \frac{(N_1 - i)(N_2 - j)}{N_1 N_2}\right). \tag{5}$$

This is the basis of Figs. 2 and 3 in the *Main Text*, and Figs. S1, S6B,C, and S5B,C,E,F in this *SI Appendix*. For each possible starting point, these figures begin with a degenerate initial distribution that is one at that point and zero elsewhere, and calculate according to Eq. (5) the average relative success of population 1 after a short-run and long-run number of time-steps. The figures display the results for all possible initial starting points.

S1.2 Antagonistic symbioses

The payoff matrix of the antagonism game [game (1) in the Main Text] is

		Player 2	
		C	D
Player 1	A	$1, 0$	$0, 1$
	B	$0, 1$	$1, 0$

(6)

Among the four matchings, population 1 prefers (A, C) and (B, D) while population 2 prefers (A, D) and $B, C)$. Therefore, given a population state (i, j), a measure of the success of population 1 in that state is the average proportion of (A, C) and (B, D) matchings in that state minus the average proportion of (A, D) and $B, C)$ matchings: $\frac{ij+(N_1-i)(N_2-j)}{N_1 N_2} - \frac{i(N_2-j)+(N_1-i)j}{N_1 N_2}$. Given any distribution over population states $\mathbf{p}$, we may calculate the expected value of population 1's success, $\sum_{(i,j)} p_{(i,j)} \cdot \left(\frac{ij+(N_1-i)(N_2-j)}{N_1 N_2} - \frac{i(N_2-j)+(N_1-i)j}{N_1 N_2}\right)$.

Given an initial distribution over population states, $\mathbf{v}^0 = \left(v^0_{(i,j)}\right)$, making use of Eq. (2), the expected value of population 1's relative success after t time-steps is:

$$\sum_{(i,j)} v^t_{(i,j)} \cdot \left(\frac{ij + (N_1 - i)(N_2 - j)}{N_1 N_2} - \frac{i(N_2 - j) + (N_1 - i)j}{N_1 N_2}\right). \tag{7}$$

This is the basis of Fig. 1 in the *Main Text*, and S6A, and S5A,D here.

S2 Numerical analysis of the complete Markov chain

The stochastic process for the complete Markov chain defined by Eq. (2) does not make any restrictions on the evolutionary parameters. It applies for any choice of generation times, mutation rates, strengths of selection, and population sizes. But computation of the time evolution of the distribution over population states requires us to calculate powers of a transition matrix of dimension $n \times n$, where $n = (N_1 + 1)(N_2 + 1)$, and so analytical results for arbitrary parameter values quickly become infeasible. However, provided that the two populations are of moderate size, we can still use Eq. (2) to obtain exact numerical results, which we report in what follows.

S2.1 Short-run dynamics versus long-run dynamics

For the complete Markov chain, we can distinguish two timescales. In the short run, the fate of the two populations largely depends on the populations' initial composition. For example, when mutualistic symbioses start close to one of the two equilibria, (A,C) or (B,D), they can be expected to further approach that equilibrium in the short run. But as time goes by, mutations and random movements against the gradient of selection can lead populations to leave one equilibrium, and to move near the other one. In the long run, the process's dependence on initial conditions disappears: as the transition matrix P in Eq. (2) is primitive, the eventual success of a population is independent of where the populations have started.

Figs. 1-3 in the *Main Text*, and Figs. S1, S6 show how differences in generation times, mutation rates, selection strength and population sizes affect population 1's success [as measured by Eqs. (5) and (7)] for the two different timescales. In mutualistic symbioses where k is large, we observe that populations with a longer generation time, lower selection strength and smaller population size have an advantage in the short run (i.e., the respective population is favored for a majority of initial conditions; see *Main Text* Fig. 2). In the long run, however, differences in generation time become largely irrelevant, and the weaker strength of selection and the smaller population size turn out to put population 1 at a disadvantage. These results are reversed when k is small (*Main Text* Fig. 3), in which case weaker selection is a disadvantage in the short run, but an advantage in the long run.

In antagonistic symbioses, the effect of the evolutionary parameters is unambiguous. Here, the population with shorter generation time, higher mutation rate, higher selection strength or larger population size is favored for all time scales.

S2.2 Robustness of results with respect to our modeling assumptions

In the Main Text, we predominantly considered in our long-run analysis the case where mutations are rare (i.e., where the parameter ε is very small). In addition, we assumed the fitness of an individual to be a *linear* function of its expected payoff, though other functions are also admissible. To show that these assumptions are not responsible for our

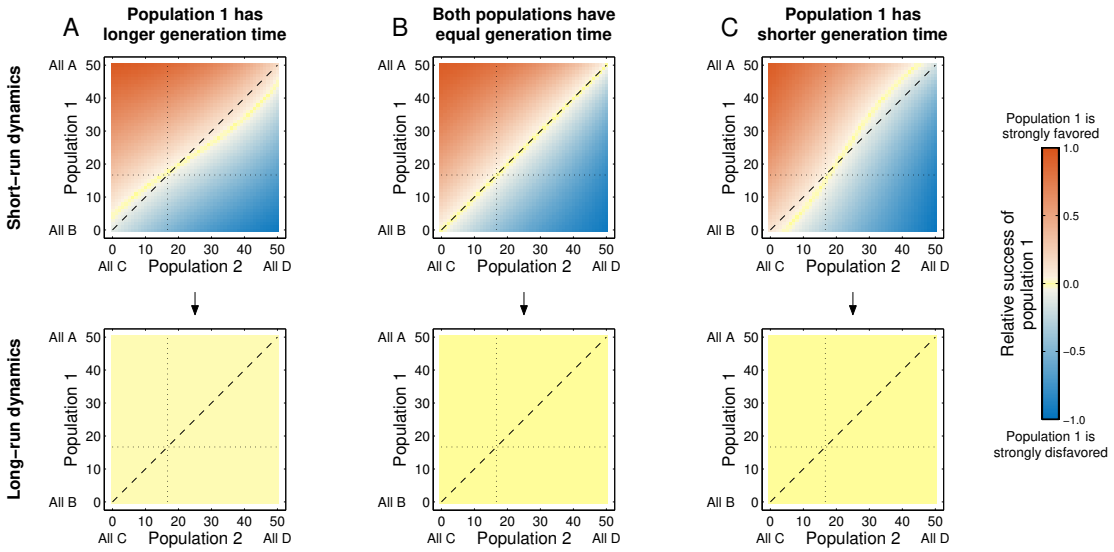

Figure S1: Mutualism dynamics when $k = 3/2$, and population 1 has (A) a longer generation time than population 2, (B) an equal generation time to population 2's, (C) a shorter generation time than population 2. For comparison with Fig. 2 in Bergstrom and Lachmann [1]. Each panel shows the numerically computed dynamics, assuming that both populations coincide in their other evolutionary parameters. The upper panels show population 1's success, as measured by Eq. (5), after 50 generations for every possible initial population (i.e., for each point in the squared state space). We observe a Red King effect: when population 1 has a longer generation time, the red area covers more than 50% of the square; the reverse is true when population 1 has a shorter generation time. These short-run results are similar to Bergstrom and Lachmann's, and indeed, the upper panels of this figure closely resemble their Fig. 2. The lower panels show the success of population 1 after 50,000 generations. By this time, selection-mutation equilibrium has been reached, so that the starting point no longer influences the dynamics: all lower panels have a uniform color. The populations are equally successful in the long run, no matter their relative generation times. Parameters: $k = 3/2$, $g_1 = g_2 = 1$, $N_1 = N_2 = 50$, $w_1 = w_2 = 0.05$, $g_1 = g_2 = 1$, $\mu_1 = \mu_2 = 1$, and $\varepsilon = 0.001$. (A) $g_1 = 10$; (C) $g_2 = 10$. A 'generation' is defined as $N_1 + N_2$ elementary updating events of the Moran process.

qualitative results, we have calculated the invariant distribution of the complete Markov chain for various values of ε (Figs. S2–S4), and using an exponential fitness function instead of the linear one (Fig. S5).

Figs. S2–S4 visualize the invariant distribution of the Markov chain for three different values of the baseline mutation rate, varying the parameter ε from $\varepsilon = 0.0001$ (top panels) to $\varepsilon = 0.01$ (bottom panels). When mutations are rare, the two populations are almost always situated at a boundary of the state space, and the states at the four corners are most abundant (as indicated by the red color). That is, in this regime, the two populations are usually monomorphic. Occasionally, a mutation introduces a new strategy into one of the two populations; this mutation typically goes extinct or fixes in that population before the next mutation occurs in either population. In this case, the weak-mutation methodology (as introduced in detail in Section S4) provides an excellent approximation of the dynamics of the complete Markov chain. We note that for mutualistic symbioses, the value of k in the payoff matrix controls the path that leads from one equilibrium to

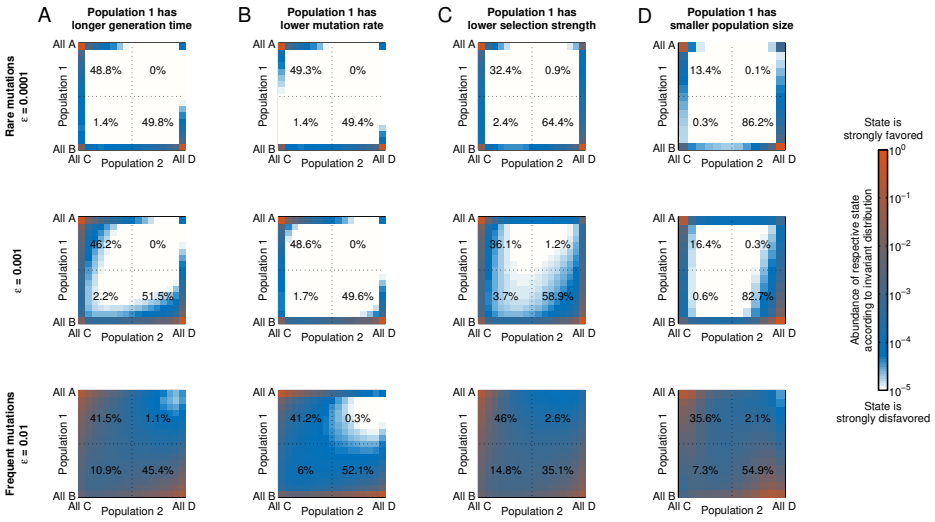

Figure S2: Effect of baseline mutation rate on the long-run abundance of strategies in mutualistic symbioses with $k = 3/2$. All panels show how abundant each population state is in the long run, according to the invariant distribution of the complete Markov chain. The numbers in each quadrant correspond to the fraction of time the respective quadrant is visited. Increases in the baseline mutation rate do not qualitatively change the weak-mutation conclusions for this game, except that differences in generation time and mutation rates lead to weak Red Queen effect [(A) and (B)]. With respect to differences in selection strength we find that a strong Red Queen effect for small mutation rates can reverse to a weak Red King effect when mutation rates are very large [(C), bottom panel]. For better visibility, we have increased the baseline strength of selection compared to the previous figures: $w_1 = w_2 = 0.2$ [and $w_1 = 0.05$, $w_2 = 0.5$ in (C)]. All other parameters are the same as before.

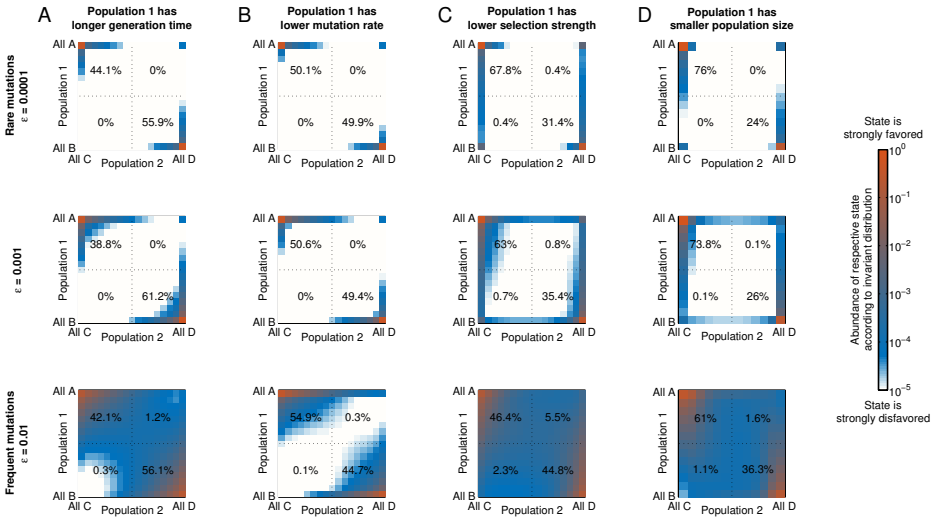

Figure S3: Effect of baseline mutation rate on the long-run abundance of strategies in mutualistic symbioses with $k = 1/2$. As in Fig. S2, the figure illustrates the invariant distribution of the complete Markov chain. Differences in selection strength and population size lead to a comparably strong Red King effect, whereas differences in mutation rates yield a weak Red King effect. Differences in generation time can result in a notable Red Queen effect. All parameters are the same as in the previous figure, except k.

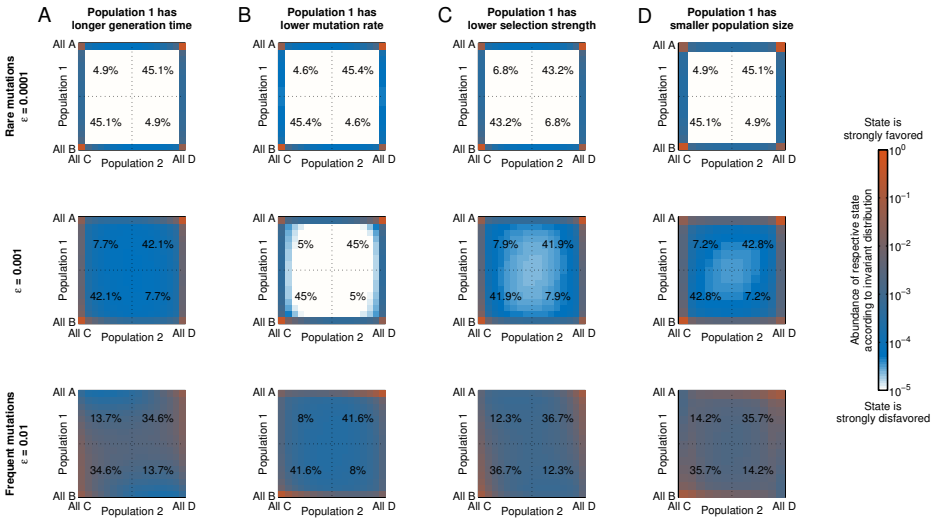

Figure S4: Effect of baseline mutation rate on the long-run abundance of strategies in antagonistic symbioses. In all panels, the invariant distribution puts less weight on states that are beneficial for population 1, i.e., states close to the (A,C) and (B,D) corners. This is as expected under Red Queen dynamics, as population 1 has a larger generation time (A), a lower mutation rate (B), weaker selection (C), or a smaller population size (D). All evolutionary parameters are the same as in Figs. S2 and S3.

another. For $k = 3/2$, transitions from (A,C) to (B,D) and back typically pass through the non-coordination state (B,C) (Fig. S2), whereas for $k = 1/2$, these transitions typically pass through the other non-coordination state (A,D) (Fig. S3).

As the mutation parameter ε increases, states in the interior of the state space occur more often, and the weak-mutation approximation is no longer an excellent quantitative match. Nevertheless, Figs. S2–S4 indicate that the qualitative predictions from the weak-mutation case are largely robust to changes in ε, in the sense that populations that are favored under the weak-mutation regime also tend to be favored as the baseline mutation rate is increased.

Similarly, we observe only a small quantitative change, and no qualitative change, when we assume that the fitness of an individual is defined as $f = \exp(w\pi)$, instead of the linear specification $f = 1 + w \cdot \pi$ used in the Main Text (Fig. S5).

S3 Mixing time of the Markov chain

We have distinguished the behavior of the dynamics in the 'short run' and the 'long run'. Here, we provide a more precise picture of how many generations it takes to be in the 'long run', for the two games we have studied. Since the object of interest in the long-run dynamics is the stationary distribution of the co-evolutionary process, the 'long run' should be defined as the number of generations required for the distribution over population states

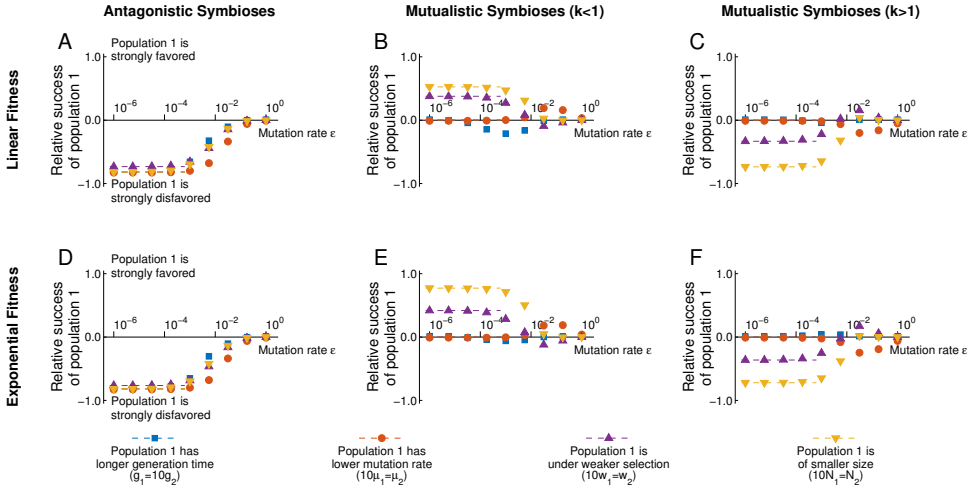

Figure S5: Impact of baseline mutation rate on the long-run relative success of population 1 in antagonistic and mutualistic symbioses. For various parameter combinations and two different fitness specifications, each panel shows the relative success of population 1, either using the payoff matrix for antagonistic symbioses (A), for mutualistic symbioses with $k = 1/2$ (B), or for mutualistic symbioses with $k = 3/2$ (C). We consider differences in generation time (blue), mutation rate (red), the strength of selection (purple), and population size (yellow). Dots show the numerically computed measures of population 1's success, as defined in (5) and (7) respectively. Dashed lines indicate the respective analytical solution in the limit of rare mutations. In antagonistic symbioses, all four dimensions put population 1 at a disadvantage when mutations become sufficiently rare. In mutualistic symbioses, only population size and selection strength have a notable effect; this effect is positive for $k = 1/2$ and negative for $k = 3/2$. The analytical solution provides a good approximation when $\varepsilon\mu \leq 10^{-4}$. Parameters are the same as in Figs. S2–S4.

to be close to the stationary distribution, i.e., the 'mixing time' of the process.

For small populations, such as those for which we have displayed results in our various figures, the time evolution of the probability distribution can be computed exactly, and so we can directly see how rapidly it converges to being stationary (Fig. S6). We see that, for the parameters we have studied in the Main Text figures and in Figs. S1-6 here (i.e., all with small population sizes, of order 50), the long run in antagonistic symbioses can be thought to be after about 10^2-10^3 generations (Fig. S6A), while in mutualistic symbioses, the long run is after about 10^3-10^4 generations (Fig. S6B,C).

For larger population sizes, direct computation of the time evolution of the evolutionary process is not feasible, since it involves taking successive powers of very large matrices. Moreover, numerical estimation of the mixing time is computationally very difficult, and in fact not feasible for our purposes, which are to see how the mixing time changes as we vary key parameters like population size. Therefore, we resort to approximate analytical arguments. We shall focus on the case where mutations are rare in both populations, so that there are seldom several mutations segregating in the populations.

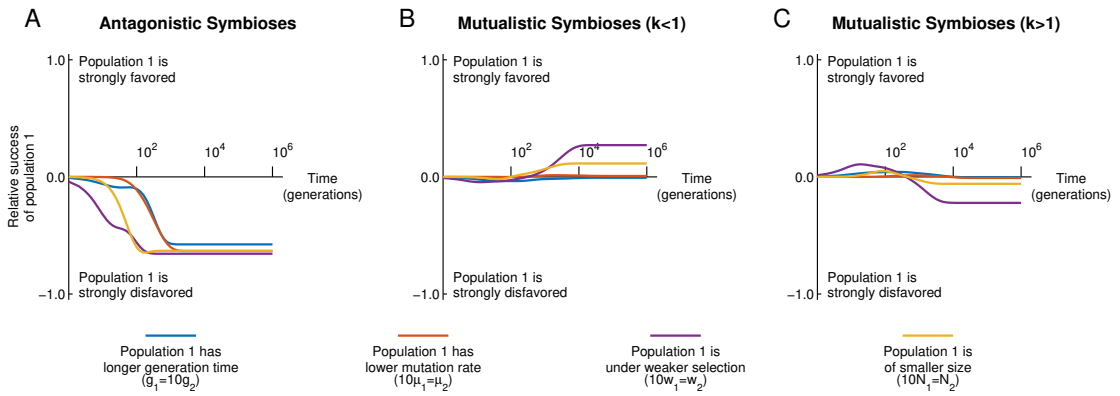

Figure S6: Dynamics of mutualistic and antagonistic symbioses over time. For various parameter combinations, the three plots show the relative success of population 1, using the payoff matrix for antagonistic symbioses (A), for mutualistic symbioses with $k = 1/2$ (B), and for mutualistic symbioses with $k = 3/2$ (C). The panels illustrate the temporal effects of differences in generation time (blue), mutation rate (red), the strength of selection (purple), and population sizes (yellow). The curves show the expected evolution of population 1's relative success, as measured by the quantities (5) for mutualistic symbioses and (7) for antagonistic symbioses. (A) In the long run, antagonistic symbioses disfavor the population with longer generation time, lower mutation rate, weaker selection strength and smaller population size. (B) In mutualistic symbioses with $k = 1/2$, the population with longer generation time and weaker selection strength is disfavored in the intermediate run. However, in the long run, differences in generation time have a negligible impact on the success of a population, and the effect of weaker selection becomes positive. The other two dimensions have either no substantial effect (differences in mutation rates) or they have a positive effect as well (smaller population size). (C) When $k = 3/2$, the effects of weaker selection and smaller population size are reversed from the case of $k = 1/2$. Parameters are the same as in Figs. 1-3 in the *Main Text*.

S3.1 Antagonistic symbioses

In antagonistic symbioses, when the mutation rate is small, the probability distribution over time is influenced predominantly by those transitions from pure state to pure state that are driven by positive selection (i.e., the 'arms race'). Given some initial starting point, the dynamics quickly move towards one of the pure states, after which the cycle from pure state to pure state begins. Having moved to some new pure state, how quickly the dynamics subsequently move to a new pure state is a random variable, but after sufficiently many such transitions, we expect the proportion of time that has been spent in each state to be close to that in the stationary distribution.

Therefore, the mixing time should be proportional to the time it takes to move from one pure state to another. Starting from a pure state, this requires a mutation to *substitute*, i.e., to arrive and fix, in the presently disfavored population. If the mutation rate is low enough, then the rate-limiting part of this process is the arrival time of a mutation destined to fix, and not the time it takes the mutation to fix (which occurs relatively quickly). Mutations arrive in this population i at rate $N_i \varepsilon \mu_i$, and have selective advantage w_i $(= [1 + w_i(1)] - [1 + w_i(0)])$.

We may distinguish two cases: weak selection and strong selection. Weak selection applies when $N_i w_i < 1$, in which case the fixation probability of beneficial mutations is

approximately $1/N_i$ (i.e., they are nearly neutral). So, under weak selection, the substitution rate of beneficial mutations in population i is approximately $N_i\varepsilon\mu_i/N_i = \varepsilon\mu_i$, and so the mixing time of the process will be approximately proportional to $1/\min_i(\varepsilon\mu_i)$. That is, the mixing time should decrease as the mutation rates of the populations increase. This approximation should be good as long as the waiting time for a mutation destined to fix in a population, $1/\varepsilon\mu_i$, is significantly longer than the time it takes that mutation to fix, which in the case of weak selection is about $2N_i$ generations, on average. So, our above analytical reasoning should be valid when is $N_i\varepsilon\mu_i \ll 1$.

Strong selection applies when $N_i w_i > 1$, and if w_i is small, the fixation probability of beneficial mutations with selective advantage w_i is approximately w_i [3]. Under strong selection, therefore, the substitution rate of beneficial mutations in population i is approximately $N_i\varepsilon\mu_i w_i$, and so the mixing time of the process will be proportional to $1/\min_i(N_i\varepsilon\mu_i w_i)$. That is, the mixing time should decrease as the sizes, mutation rates, and selection strengths of the populations increase. Again, this approximation will be good if the waiting time for a mutation destined to fix in a population, here $1/(N_i\varepsilon\mu_i w_i)$, is much larger than the time it takes to fix, which is $\log(N_i)/w_i$ on average. So, we require $N_i \log(N_i)\varepsilon\mu_i \ll 1$.

S3.2 Mutualistic symbioses

For the evolutionary process to mix in the case of the mutualism game also requires that each pure state be visited sufficiently often, the equilibria (A, C) and (B, D), as well as the non-equilibria (A, D) and (B, C).

In the case of weak selection in both populations, $N_1 w_1, N_2 w_2 < 1$, substitutions that drive the population state from one pure state to the other occur at rate $\varepsilon\mu_i$ if in population i, and so the mixing time of the process should be proportional to $1/\min_i(\varepsilon\mu_i)$, as in the case of antagonistic symbioses with weak selection.

In the case of strong selection in at least one population, the mixing time of the evolutionary process will be determined by how long it takes to substitute *against* selection in that population (or both), i.e., to substitute out of equilibrium. This is because substitutions into equilibrium, under positive selection, occur on a much faster timescale than substitutions out of equilibrium when the mutation rates in the two directions are equal. Another way of saying this is that the evolutionary process will spend most of the time in equilibrium states.

For example, consider the substitution rate from the equilibrium state (A, C) to the non-equilibrium state (A, D). This involves the arrival in an all-C population 2 of a D mutant, and the subsequent fixation of the D mutant. D mutants arrive at rate $N_2\varepsilon\mu_2$. They receive payoff 0 versus the incumbent strategy C's payoff of 1, and are therefore at relative selective disadvantage $-w_2/(1 + w_2)$. Under strong selection, the fixation probability of such a mutant is approximately $\frac{w_2}{1+w_2}\exp\left(-N_2\frac{w_2}{1+w_2}\right)$ [5], and so the substitution rate from (A, C) to (A, D) is approximately $N_2\varepsilon\mu_2\frac{w_2}{1+w_2}\exp\left(-N_2\frac{w_2}{1+w_2}\right)$. The substitu-

tion rates from (A,C) to (B,C), from (B,D) to (B,C), and from (B,D) to (A,D) can be estimated similarly as $N_1\varepsilon\mu_1\frac{w_1(2-k)}{1+2w_1}\exp\left(-N_1\frac{w_1(2-k)}{1+2w_1}\right)$, $N_2\varepsilon\mu_2\frac{w_2(2-k)}{1+2w_2}\exp\left(-N_2\frac{w_2(2-k)}{1+2w_2}\right)$, and $N_1\varepsilon\mu_1\frac{w_1}{1+w_1}\exp\left(-N_1\frac{w_1}{1+w_1}\right)$ respectively.

The mixing time under strong selection should be proportional to the inverse of the smallest of these four substitution rates. In each case, the properties of the mixing time will be similar. Suppose, for instance, that the (A,C) to (A,D) substitution is slowest. Then the mixing time of the process is proportional to

$$\frac{\exp\left(N_2\frac{w_2}{1+w_2}\right)}{N_2\frac{w_2}{1+w_2}\varepsilon\mu_2}.$$

For large $N_2\frac{w_2}{1+w_2}$, the exponential term in the numerator dominates, and the mixing time becomes very large. In such cases, the stationary distribution is unlikely to be approached on realistic timescales. For smaller values of $N_2\frac{w_2}{1+w_2}$, mixing will occur on realistic timescales.

S4 Weak-mutation methodology

A general description of the weak-mutation methodology that we have employed can be found in ref. [9]. The inclusion of generation rates is the only difference between this and the methodology set out there. Here, we shall give the method in full for the particular case of 2-player, 2-strategy games that we have studied. We alter the notation a little from the previous section, so that the payoff to a member of population i playing strategy X who encounters a member of the other population playing strategy Y is denoted $\pi_i(X,Y)$

In the weak mutation limit, $\varepsilon\to 0$, the evolutionary dynamics converge to an embedded dynamics over just the pure states (in which each population is monomorphic). These pure states can be labelled (A,C), (A,D), (B,C), and (B,D), which we shall, for notational convenience, enumerate as pure states 1, 2, 3, and 4. We write $\rho_i(x,y)$ for the fixation probability of a single payoff-x mutant in population i otherwise pure for a payoff-y strategy. In the Moran process that we have predominantly used, and which is described in the previous section,

$$\rho_i(x,y)=\frac{1-\left(\frac{1+w_ix}{1+w_iy}\right)^{-1}}{1-\left(\frac{1+w_ix}{1+w_iy}\right)^{-N_i}}.$$

The dynamics can also easily be extended to other processes. For example, it could be that population i experiences a full-population Wright-Fisher update with probability proportional to g_i. In this case, applying to the above, we write $s(x,y)=1-\frac{1+w_ix}{1+w_iy}$ and

make use of Kimura's diffusion approximation

$$\rho_i(x,y) = \frac{1-e^{-s(x,y)}}{1-e^{-N_i s(x,y)}}.$$

In fact, we may make use of any process that satisfies the properties that, without mutations, (a) absent strategies would remain absent, and (b) any strategy that is present but not fixed has positive probability of increasing its representation in the next generation. Given such a process, and therefore fixation probability functions $\rho_1(\cdot,\cdot)$ and $\rho_2(\cdot,\cdot)$ for the two populations, the one-step transition probabilities in the embedded dynamics are:

$$\begin{aligned}
&P_{1\to 2} = \frac{N_2\mu_2}{g_2}\rho_2\Big(\pi_2(D,A),\pi_2(C,A)\Big); \quad P_{2\to 1} = \frac{N_2\mu_2}{g_2}\rho_2\Big(\pi_2(C,A),\pi_2(D,A)\Big);\\
&P_{1\to 3} = \frac{N_1\mu_1}{g_1}\rho_1\Big(\pi_1(B,C),\pi_1(A,C)\Big); \quad P_{3\to 1} = \frac{N_1\mu_1}{g_1}\rho_1\Big(\pi_1(A,C),\pi_1(B,C)\Big);\\
&P_{2\to 4} = \frac{N_1\mu_1}{g_1}\rho_1\Big(\pi_1(B,D),\pi_1(A,D)\Big); \quad P_{4\to 2} = \frac{N_1\mu_1}{g_1}\rho_1\Big(\pi_1(A,D),\pi_1(B,D)\Big);\\
&P_{3\to 4} = \frac{N_2\mu_2}{g_2}\rho_2\Big(\pi_2(D,B),\pi_2(C,B)\Big); \quad P_{4\to 3} = \frac{N_2\mu_2}{g_2}\rho_2\Big(\pi_2(C,B),\pi_2(D,B)\Big);\\
&P_{1\to 4} = P_{4\to 1} = P_{2\to 3} = P_{3\to 2} = 0;\\
&P_{1\to 1} = 1 - P_{1\to 2} - P_{1\to 3}; \;\; P_{2,2} = 1 - P_{2\to 1} - P_{2\to 4};\\
&P_{3\to 3} = 1 - P_{3\to 1} - P_{3\to 4}; \;\; P_{4\to 4} = 1 - P_{4\to 2} - P_{4\to 3}.
\end{aligned}$$

The stationary distribution over the states, $\lambda = [\lambda_1, \lambda_2, \lambda_3, \lambda_4]$, is the stochastic vector that solves

$$\lambda_i = \sum_{j=1}^{4} \lambda_j P_{j\to i} \quad \text{for } i = 1,2,3,4.$$

S5 Antagonistic symbioses, weak-mutation limit

We consider here the antagonistic symbiosis payoff matrix (game [1] in the *Main Text*).

Imposing no restrictions on the rate parameters μ_1, μ_2, g_1, g_2, N_1, N_2, w_1, and w_2, the stationary distribution, calculated according to the method described in the previous section, is

$$\lambda = [s_1, 1, 1, s_1]/\bar{\lambda},$$

where $\bar{\lambda} = 2 + 2s_1$, and

$$s_1 = \frac{\frac{N_1\mu_1}{g_1}\rho_1(1,0) + \frac{N_2\mu_2}{g_2}\rho_2(0,1)}{\frac{N_1\mu_1}{g_1}\rho_1(0,1) + \frac{N_2\mu_2}{g_2}\rho_2(1,0)}.$$

The long-run success of population 1 relative to population 2 is proportional to $\lambda_1 - \lambda_4 = (s_1 - 1)/\bar{\lambda}$, which is increasing in s_1.

Holding the rate parameters of population 2 constant, and focusing on the rate pa-

rameters of population 1, a sufficient condition for the success of population 1 to increase is therefore that the term $\frac{N_1\mu_1}{g_1}\rho_1(1,0)$ increases and the term $\frac{N_1\mu_1}{g_1}\rho_1(0,1)$ decreases. The former term is the substitution rate of beneficial mutations in population 1, while the latter is the substitution rate of deleterious mutations.

S5.1 Selection strength

A higher selection strength w_1 increases the fixation probability of beneficial mutations $\rho_1(1,0)$ and decreases the fixation probability of deleterious mutations $\rho_1(0,1)$ for any reasonable evolutionary process (this is, after all, a plausible *definition* of selection strength), but has no effect on the arrival rate of these mutations. It therefore increases the substitution rate of beneficial mutations and decreases the substitution rate of deleterious mutations, and therefore increases s_1 and thus the relative success of population 1.

S5.2 Population size

It follows that, if the evolutionary process is such that beneficial mutations have a faster and deleterious mutations a slower substitution rate in larger populations (loosely, if natural selection acts more efficiently in larger populations), then a larger population size for population 1 is associated with a larger value of s_1, i.e., greater long-run success. It is folk knowledge that this property holds of the Wright-Fisher and Moran processes, though we have been unable to find direct proofs in the literature, and so provide them here.

We should note that the usual statement of the result that beneficial mutations substitute at a higher rate in larger populations assumes that the beneficial mutations in question confer a selective advantage $s \ll 1/N_e$ while $N_e s \gg 1$, so that their fixation probability in isolation can be approximated by s (in a haploid population) or $2s$ (in a diploid population, where s is the haploid fitness contribution). It follows directly that the substitution rate, $\mu N s$ (haploid) or $4\mu N s$ (diploid), is increasing in N.

In contrast, the results below hold for *all* selection coefficients s (including $s < 0$) and population sizes N.

Moran: Suppose that mutants are of relative fitness $r \neq 1$. Then, if independent (no interference of any sort), they arrive and fix at rate

$$\mu N \frac{1 - r^{-1}}{1 - r^{-N}}.$$

The proportional change of this quantity with respect to change in N is

$$\frac{\partial}{\partial N} \ln\left(\mu N \frac{1 - 1/r}{1 - 1/r^N}\right) = \frac{1}{N} - \frac{\ln r}{r^N - 1} \overset{\text{sign}}{=} \operatorname{sgn}(r-1)[r^N - 1 - N \ln r] \overset{\text{sign}}{=} \operatorname{sgn}(r-1)$$

where the second step is the result of multiplying through by $N(r^N - 1)$, which is positive if $r > 1$ and negative if $r < 1$, and the last step follows from $x > 1 + \ln(x)$ for all $x \neq 1$.

Therefore,

$$\operatorname{sgn}\left[\frac{\partial}{\partial N}\left(\mu N\frac{1-r^{-1}}{1-r^{-N}}\right)\right]=\operatorname{sgn}(r-1),$$

and so beneficial mutations ($r>1$) arrive and fix at a higher rate in a larger population, while detrimental mutations ($r<1$) arrive and fix at a lower rate in a larger population.

Wright-Fisher (diffusion approximation): For a Wright-Fisher process, in the diffusion limit, writing $s=r-1$, the arrival-fixation rate varies with N according to

$$\begin{aligned}\frac{\partial}{\partial N}\left(\mu N\frac{1-e^{-s}}{1-e^{-Ns}}\right)&=\mu\left(\frac{1-e^{-s}}{1-e^{-Ns}}-N\frac{1-e^{-s}}{(1-e^{-Ns})^2}se^{-Ns}\right)\\&\overset{\text{sign}}{=}\operatorname{sgn}(s)\left[1-e^{-Ns}-Nse^{-Ns}\right]=\operatorname{sgn}(s)\left[1-\frac{1+Ns}{e^{Ns}}\right].\end{aligned}$$

But $e^x>1+x$ for all $x\neq 0$, and so

$$\operatorname{sgn}\left[\frac{\partial}{\partial N}\left(\mu N\frac{1-e^{-s}}{1-e^{-Ns}}\right)\right]=\operatorname{sgn}(s).$$

Though they do not bear on the problem considered in our *Main Text*, it is nonetheless worth noting some implications of the above results. One major implication concerns the rate at which populations of different size are expected to adapt; that is, the average rate at which fitness increases in populations of differing size. Again, the result below is folk knowledge in population genetics, but we have been unable to find a general proof.

Assume that mutations arise at rate μ per replication, and sufficiently infrequently at the population level that their fate, extinction or fixation, is almost always determined before the arrival of the next mutation in the population—this is the commonly-assumed 'sequential fixations' model [6]. We study a haploid population of size N, where the next mutation that appears in the population is drawn from some fitness effect distribution $f(s)$, where s is the fitness difference between current members of the population and the potential mutant. We assume this distribution to be atomless in what follows, but this is not necessary in general. Writing $\rho(s,N)$ for the fixation probability of a mutant of fitness effect s in a population of size N (as written for the Moran and Wright-Fisher processes above), the current rate of fitness increase of the population can then be written

$$R_{\text{fitness}}=N\mu\int_{-\infty}^{\infty}sf(s)\rho(s,N)ds.$$

This changes with N according to

$$\begin{aligned}\frac{\partial}{\partial N}R_{\text{fitness}} &= \frac{\partial}{\partial N}\left[N\mu\int_{-\infty}^{\infty} sf(s)\rho(s,N)ds\right] \\ &= \frac{\partial}{\partial N}\left[-\int_{-\infty}^{0} -|s|f(s)\mu N\rho(s,N)ds + \int_{0}^{\infty} sf(s)\mu N\rho(s,N)ds\right] \\ &= \int_{-\infty}^{0} -|s|f(s)\mu\left[\frac{\partial}{\partial N}N\rho(s,N)\right]ds + \int_{0}^{\infty} sf(s)\mu\left[\frac{\partial}{\partial N}N\rho(s,N)\right]ds.\end{aligned}$$

From what we showed earlier, we note that if the evolutionary process is a Wright-Fisher or a Moran process, then both integrands in this last line are positive, and so $\frac{\partial}{\partial N}R_{\text{fitness}} > 0$.

S5.3 Mutation rate and generation time

Finally, write $r_1 = (\mu_1/g_1)/(\mu_2/g_2)$, so that $s_1 = \frac{N_1 r_1 \rho_1(1,0)+N_2\rho_2(0,1)}{N_1 r_1 \rho_1(0,1)+N_2\rho_2(1,0)}$. Then

$$\frac{\partial s_1}{\partial r_1} = N_1 N_2 \frac{\rho_1(1,0)\rho_2(1,0) - \rho_1(0,1)\rho_2(0,1)}{[N_1 r_1\rho_1(0,1) + N_2\rho_2(1,0)]^2},$$

which is positive if $\rho_1(1,0) > \rho_1(0,1)$ and $\rho_2(1,0) > \rho_2(0,1)$, i.e., if, in both populations, beneficial mutations have higher fixation probability than deleterious mutations. Again, this is true of all reasonable evolutionary processes.

S5.4 Moran process

In the special case of a Moran process operating in each population, and assuming $N_1 = N_2 = N$ and $w_1 = w_2 = w$, we have $\rho_1 \equiv \rho_2 = \rho$, and $\frac{\rho(1,0)}{\rho(0,1)} = (1+w)^{N-1} =: \gamma$. Then

$$s_1 = \frac{Nr_1\rho(1,0) + N\rho(0,1)}{Nr_1\rho(0,1) + N\rho(1,0)} = \frac{r_1\frac{\rho(1,0)}{\rho(0,1)} + 1}{r_1 + \frac{\rho(1,0)}{\rho(0,1)}} = \frac{r_1\gamma + 1}{r_1 + \gamma}.$$

This is the basis of Eq. [3] in the Main Text.

S6 Mutualistic symbioses, weak-mutation limit

The mutualism payoff matrix (game [2] in the *Main Text*) is

		Player 2	
		C	D
Player 1	A	$2,1$	$0,0$
	B	k,k	$1,2$

The stationary distribution, calculated according to the method above, is

$$\lambda = \left[\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3, \tilde{\lambda}_4\right] / \bar{\lambda},$$

where $\bar{\lambda}$ is a normalization constant and

$$\begin{aligned}
\tilde{\lambda}_1 &= \rho_1(2,k)\left[\rho_2(k,2)\rho_1(1,0) + \rho_2(1,0)\rho_1(0,1)\right] + r_1\rho_2(1,0)\left[\rho_2(2,k)\rho_1(0,1) + \rho_2(k,2)\rho_1(2,k)\right] \\
\tilde{\lambda}_2 &= \rho_1(0,1)\left[\rho_2(2,k)\rho_1(k,2) + \rho_2(0,1)\rho_1(2,k)\right] + r_1\rho_2(0,1)\left[\rho_2(2,k)\rho_1(0,1) + \rho_2(k,2)\rho_1(2,k)\right] \\
\tilde{\lambda}_3 &= \rho_1(k,2)\left[\rho_2(k,2)\rho_1(1,0) + \rho_2(1,0)\rho_1(0,1)\right] + r_1\rho_2(k,2)\left[\rho_2(0,1)\rho_1(1,0) + \rho_2(1,0)\rho_1(k,2)\right] \\
\tilde{\lambda}_4 &= \rho_1(1,0)\left[\rho_2(2,k)\rho_1(k,2) + \rho_2(0,1)\rho_1(2,k)\right] + r_1\rho_2(2,k)\left[\rho_2(0,1)\rho_1(1,0) + \rho_2(1,0)\rho_1(k,2)\right]
\end{aligned}$$

Here, $r_1 = (\mu_1 N_1/g_1)/(\mu_2 N_2/g_2)$ (this definition is slightly different from that we employed in the antagonistic symbiosis section above). Notice that the dependence on r_1 is not necessarily as simple as it immediately appears from the expressions for the $\tilde{\lambda}_i$, because the normalizing constant $\bar{\lambda}$ also depends on r_1.

S6.1 Mutation rate and generation time

If we fix $N_1 = N_2 = N$ and $w_1 = w_2 = w$, then the fixation probability functions of the two populations coincide: $\rho_1 \equiv \rho_2 =: \rho$. The stationary distribution takes on the simple form

$$\lambda = \left[1, \frac{\rho(0,1)}{\rho(1,0)}, \frac{\rho(k,2)}{\rho(2,k)}, 1\right] / \bar{\lambda},$$

where $\bar{\lambda}$ normalizes that λ sums to one. The stationary distribution is therefore independent of the mutations rates and generation times of the two populations.

In the case of the Moran process, on which we focused in the *Main Text*, the stationary distribution simplifies to Eq. [5] in the Main Text:

$$\lambda = \left[1, \left(\frac{1}{1+w}\right)^{N-1}, \left(\frac{1+kw}{1+2w}\right)^{N-1}, 1\right] / \bar{\lambda}.$$

Why do the mutation rate and generation time have no effect on the stationary distribution when mutation is weak? This result is a special case of the following general statement: For any 4-state Markov chain, if

$$\begin{aligned}
P_{1\to2}/P_{2\to1} = P_{2\to4}/P_{4\to2} &= A, \\
P_{1\to3}/P_{3\to1} = P_{4\to3}/P_{3\to4} &= B, \\
\text{and } \; P_{1\to4} = P_{4\to1} = P_{2\to3} = P_{3\to2} &= 0,
\end{aligned}$$

then the stationary distribution is simply $\lambda = [1, A, B, 1]/\bar{\lambda}$, where $\bar{\lambda} = 2 + A + B$.

In the mutualism game, when $N_1 = N_2 = N$ and $w_1 = w_2 = w$, the fixation probability functions $\rho_1(\cdot,\cdot)$ and $\rho_2(\cdot,\cdot)$ are identical (not just for the Moran process, but for any

evolutionary process whose fixation probability function depends only on population size and selection strength). Let $\rho(\cdot,\cdot) = \rho_1(\cdot,\cdot) = \rho_2(\cdot,\cdot)$. Then

$$\frac{P_{1\to2}}{P_{2\to1}} = \frac{\frac{N_2\mu_2}{g_2}\rho_2(0,1)}{\frac{N_2\mu_2}{g_2}\rho_2(1,0)} = \frac{\rho(0,1)}{\rho(1,0)} = \frac{\frac{N_1\mu_1}{g_1}\rho_1(0,1)}{\frac{N_1\mu_1}{g_1}\rho_1(1,0)} = \frac{P_{2\to4}}{P_{4\to2}}.$$

So $P_{1\to2}/P_{2\to1} = P_{2\to4}/P_{4\to2} = A$, where $A = \rho(0,1)/\rho(1,0)$ is independent of the mutation rates and generation rates of the two populations. Similarly, $P_{1\to3}/P_{3\to1} = P_{4\to3}/P_{3\to4} = B$, where $B = \rho(k,2)/\rho(2,k)$ too depends only on N and w.

So the stationary distribution does not depend on μ_1, μ_2, g_1, or g_2. This result stems from consideration of probabilities of transitions against selection (e.g., $P_{1\to2} > 0$); it cannot be found using methods that rule out such transitions, such as deterministic dynamics like the replicator equation.

S6.2 Selection strength

In order to explore the effect of selection strength on how the populations fare (in the case of a Moran process) we look at the effect of this parameter on $\lambda_{(A,C)}$ and $\lambda_{(B,D)}$, i.e., λ_1 and λ_4. Using the expression given above for the stationary distribution, we calculate that

$$\begin{aligned}\lambda_{(A,C)} - \lambda_{(B,D)} \overset{\text{sign}}{=} &\left(\frac{1}{\rho_2(2,k)} + r_1\frac{1}{\rho_1(2,k)}\right)\left(\left(\frac{1}{1+w_1}\right)^{N_1-1} - \left(\frac{1}{1+w_2}\right)^{N_2-1}\right)\\ &- \left(\frac{1}{\rho_2(1,0)} + r_1\frac{1}{\rho_1(1,0)}\right)\left(\left(\frac{1+kw_1}{1+2w_1}\right)^{N_1-1} - \left(\frac{1+kw_2}{1+2w_2}\right)^{N_2-1}\right) =: g(w_1,w_2).\end{aligned}$$

Fix $N_1 = N_2 = N$ and $w_2 = w$, and assume a slight increase in w_1 from this value ($w_1 \to w + \Delta w$). Define

$$c_1(x) := \frac{1 - \left(\frac{1+kx}{1+2x}\right)^{-1}}{1 - \left(\frac{1+kx}{1+2x}\right)^{-N}};\quad g_1(x) := \left(\frac{1}{1+x}\right)^{N-1};$$

$$c_2(x) := \frac{1 - \left(\frac{1}{1+x}\right)^{-1}}{1 - \left(\frac{1}{1+x}\right)^{-N}};\quad g_2(x) := \left(\frac{1+kx}{1+2x}\right)^{N-1},$$

and note that $c_1(w) = 1/\rho_2(2,k)$ and $c_2(w) = 1/\rho_2(1,0)$. Then

$$g(w_1,w) = \left[c_1(w) + r_1c_1(w_1)\right]\left(g_1(w_1) - g_1(w)\right) - \left[c_2(w) + r_1c_2(w)\right]\left(g_2(w_1) - g_2(w)\right).$$

Let $w_1 = w + \Delta w$. Taylor expanding c_1, c_2, g_1 and g_2 around $\Delta w = 0$,

$$g(w + \Delta w, w) = \underbrace{(1 + r_1)\left[c_1(w)g_1'(w) - c_2(w)g_2'(w)\right]}_{\text{First order term}} \Delta w + o(\Delta w)$$

Hence, the sign of $\lambda_{(A,C)} - \lambda_{(B,D)}$ for Δw sufficiently small (and positive) will be determined by the sign of the first order term:

$$\lambda_{(A,C)} - \lambda_{(B,D)} \overset{\text{sign}}{=} \frac{1}{\rho_2(2,k)} g_1'(w) - \frac{1}{\rho_2(1,0)} g_2'(w)$$

When $w = w^* := \frac{1-k}{k}$, $\rho_2(2,k) = \rho_2(1,0)$ and $g_1(w) = g_2(w)$. This suggests that a useful reparameterization is $w = \frac{1-k+\eta}{k}$. Then $\eta > 0$ corresponds to $w > w^*$ and $\eta < 0$ corresponds to $w < w^*$. Since w is constrained to be positive, we must have that $\eta > k-1$. With this reparameterization we have

$$\begin{aligned}
\tilde{c}_1(\eta) &:= c_1\left(\frac{1-k+\eta}{k}\right) = \frac{1 - k^{-N}\left(1 - \frac{\eta}{2-k+2\eta}\right)^{-N}}{1 - k^{-1}\left(1 - \frac{\eta}{2-k+2\eta}\right)^{-1}}; \\
\tilde{g}_1(\eta) &:= g_1\left(\frac{1-k+\eta}{k}\right) = k^{N-1}\left(1 - \frac{\eta}{1+\eta}\right)^{N-1}; \\
\tilde{c}_2(\eta) &:= c_2\left(\frac{1-k+\eta}{k}\right) = \frac{1 - k^{-N}\left(1 - \frac{\eta}{1+\eta}\right)^{-N}}{1 - k^{-1}\left(1 - \frac{\eta}{1+\eta}\right)^{-1}}; \\
\tilde{g}_2(\eta) &:= g_2\left(\frac{1-k+\eta}{k}\right) = k^{N-1}\left(1 - \frac{\eta}{2-k+2\eta}\right)^{N-1}.
\end{aligned}$$

Define $v_1(\eta) := k\left(1 - \frac{\eta}{2-k+2\eta}\right)$ and $v_2(\eta) := k\left(1 - \frac{\eta}{1+\eta}\right)$. Then

$$\begin{aligned}
\lambda_{(A,C)} - \lambda_{(B,D)} &\overset{\text{sign}}{=} \frac{1 - v_2^{-1}(\eta)}{1 - v_2^{-N}(\eta)} \frac{d}{d\eta}\left(v_2^{N-1}(\eta)\right) - \frac{1 - v_1^{-1}(\eta)}{1 - v_1^{-N}(\eta)} \frac{d}{d\eta}\left(v_1^{N-1}(\eta)\right) \\
&\overset{\text{sign}}{=} \frac{1 - v_2^{-1}(\eta)}{1 - v_2^{-N}(\eta)} v_2^{N-2}(\eta) v_2'(\eta) - \frac{1 - v_1^{-1}(\eta)}{1 - v_1^{-N}(\eta)} v_1^{N-2}(\eta) v_1'(\eta)
\end{aligned}$$

It can easily be verified that $v_1'(\eta) < 0$ and $v_2'(\eta) < 0$. Writing $f(x) := \frac{1-x^{-1}}{1-x^{-N}} x^{N-2}$,

$$\lambda_{(A,C)} - \lambda_{(B,D)} \overset{\text{sign}}{=} f(v_2(\eta)) v_2'(\eta) - f(v_1(\eta)) v_1'(\eta)$$

Let

$$h_1(\eta) := \frac{\eta}{2-k+2\eta} \qquad \left[\Rightarrow h_1'(\eta) := \frac{2-k}{(2-k+2\eta)^2} = \frac{(2-k)}{\eta^2}\left(1-\frac{v_1(\eta)}{k}\right)^2\right]$$

$$h_2(\eta) := \frac{\eta}{1+\eta} \qquad \left[\Rightarrow h_2'(\eta) := \frac{1}{(1+\eta)^2} = \frac{1}{\eta^2}\left(1-\frac{v_2(\eta)}{k}\right)^2\right].$$

Clearly $h_1'(\eta) = -v_1'(\eta)/k$ and $h_2'(\eta) = -v_2'(\eta)/k$, so

$$\begin{aligned}\lambda_{(A,C)} - \lambda_{(B,D)} &\overset{\text{sign}}{=} f(v_1(\eta))h_1'(\eta) - f(v_2(\eta))h_2'(\eta) \\ &\overset{\text{sign}}{=} (2-k)f(v_1(\eta))\left(1-\frac{v_1(\eta)}{k}\right) - f(v_2(\eta))\left(1-\frac{v_2(\eta)}{k}\right)\end{aligned}$$

Setting $t_k(x) := \left(1-\frac{x}{k}\right)^2 f(x)$,

$$\lambda_{(A,C)} - \lambda_{(B,D)} \overset{\text{sign}}{=} (2-k)t_k(v_1(\eta)) - t_k(v_2(\eta)).$$

We note that for any fixed values of x_1 and x_2 — for example, $x_1 = v_1(\eta)$ and $x_2 = v_2(\eta)$ — and any k, if $x_1 < x_2$, then

$$\begin{aligned}\frac{t_k(x_1)}{t_k(x_2)} &= \left(\frac{x_1}{x_2}\right)^{N-2}\left(\frac{1-\frac{1}{x_2}^N}{1-\frac{1}{x_1}^N}\cdot\frac{1-\frac{1}{x_1}}{1-\frac{1}{x_2}}\cdot\left[\frac{1-\frac{x_1}{k}}{1-\frac{x_2}{k}}\right]^2\right) \\ &= \left(\frac{x_1}{x_2}\right)^{2N-2}\left(\frac{1-x_2^N}{1-x_1^N}\cdot\frac{1-\frac{1}{x_1}}{1-\frac{1}{x_2}}\cdot\left[\frac{1-\frac{x_1}{k}}{1-\frac{x_2}{k}}\right]^2\right) \xrightarrow[N\to\infty]{} 0\end{aligned}$$

It can also be verified that $v_1(\eta) < v_2(\eta)$ when $\eta < 0$, while $v_1(\eta) > v_2(\eta)$ when $\eta > 0$. Therefore, for large enough N:

$$\begin{aligned}&\text{if } \eta > 0, \;\; (2-k)t_k(v_1(\eta)) > t_k(v_2(\eta)) \qquad \left[\Rightarrow \lambda_{(A,C)} - \lambda_{(B,D)} > 0\right]; \\ &\text{if } \eta < 0, \;\; (2-k)t_k(v_1(\eta)) < t_k(v_2(\eta)) \qquad \left[\Rightarrow \lambda_{(A,C)} - \lambda_{(B,D)} < 0\right].\end{aligned}$$

In the special case $\eta = 0$ (that is, $w = w^*$) we have that $t_k(v_1(\eta)) = t_k(v_2(\eta))$ so

$$\lambda_{(A,C)} - \lambda_{(B,D)} \overset{\text{sign}}{=} 1-k$$

In summary, in the limit of large N, when $w_1 = w + \Delta w > w = w_2$,

1. If $k \in (0,1)$, then
$$\lambda_{(B,D)} < \lambda_{(A,C)} \qquad \text{if } w \geq \frac{1-k}{k}$$
That is, the increase in the strength of selection in population 1 is beneficial.

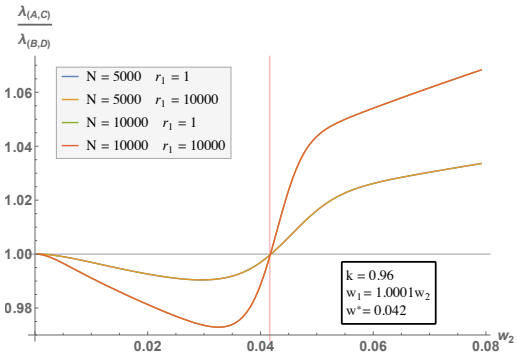

(a) Increasing the strength of selection when $k \in (0, 1)$.

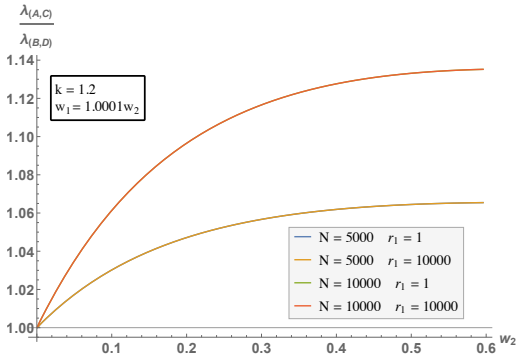

(b) Increasing the strength of selection when $k \in [1, 2)$.

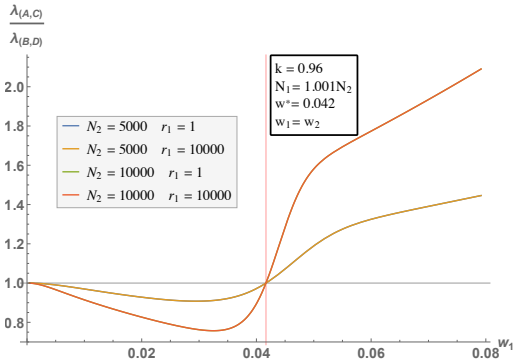

(c) Increasing the population size when $k \in (0, 1)$.

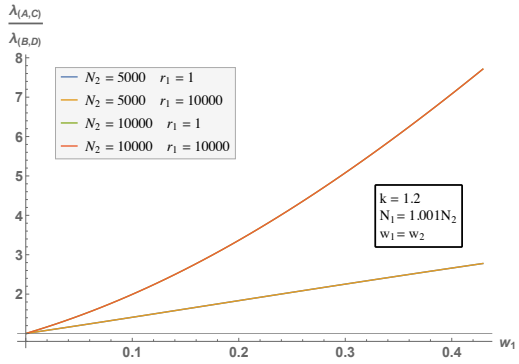

(d) Increasing the population size when $k \in [1, 2)$.

Figure S7: The effect of a slight increase in the population size or selection strength of population 1 for different values of k. In each case, the blue and green lines are invisible because they coincide with the yellow and red lines respectively. This highlights the insensitivity of effects to large changes in r_1.

$$\lambda_{(B,D)} > \lambda_{(A,C)} \qquad \text{if } w < \frac{1-k}{k}$$

That is, the increase in the strength of selection in population 1 is detrimental.

2. If $k \in [1, 2)$:

$$\lambda_{(B,D)} < \lambda_{(A,C)}$$

That is, the increase in the strength of selection in population 1 is beneficial.

To see a graphical representation of this behaviour refer to Figures S7a and S7b.

Numerical calculations suggest that for a large enough population size a generalization of the above findings hold, beyond the marginal ($w_1 = w_2 + \Delta w$) effects we have studied analytically above. That is, for fixed w_2 we have

1. If $k \in (0, 1)$, then

   (a) If $w_2 > w^*$

   $$\lambda_{(B,D)} < \lambda_{(A,C)} \qquad \text{if } w_1 > w_2 \text{ or } w_1 < w^*$$

   That is, if w_2 is above the threshold then population 1 can do better either by having a selection strength below the threshold or by having a selection strength greater than that in population 2.

   $$\lambda_{(B,D)} > \lambda_{(A,C)} \qquad \text{if } w^* < w < w_2$$

   That is, having stronger selection in population 1 is detrimental.

   (b) If $w_2 < w^*$

   $$\lambda_{(B,D)} < \lambda_{(A,C)} \qquad \text{if } w_1 < w_2$$

   That is, having weaker selection in population 1 is beneficial.

   $$\lambda_{(B,D)} > \lambda_{(A,C)} \qquad \text{if } w_1 > w_2$$

   That is, having stronger selection in population 1 is detrimental.

2. If $k \in [1, 2)$:

   $$\lambda_{(B,D)} < \lambda_{(A,C)} \qquad \text{if } w_1 > w_2$$

   That is, having stronger selection in population 1 is beneficial.

   $$\lambda_{(B,D)} > \lambda_{(A,C)} \qquad \text{if } w_1 < w_2$$

   That is, having weaker selection in population 1 is detrimental.

A graphical illustration of these results can be seen Figure S8.

S6.3 Population size

To investigate the effect of population size we fix $w_1 = w_2 = w$ and $N_2 = N = lN_0$, and assume a slight increase in N_1 ($N_1 \to lN_0 + \Delta N$). Or rather, $l_1 \to l + \Delta l$ so that $N_1 \to N_0(l + \Delta l)$. Using the same methods and definitions as in the investigation of the

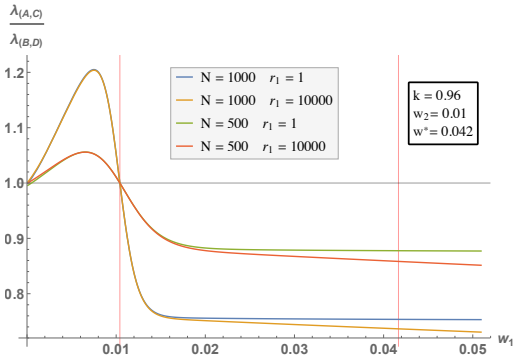

(a) Varying the strength of selection in population 1 when $k \in (0, 1)$ and $w_2 < w^*$.

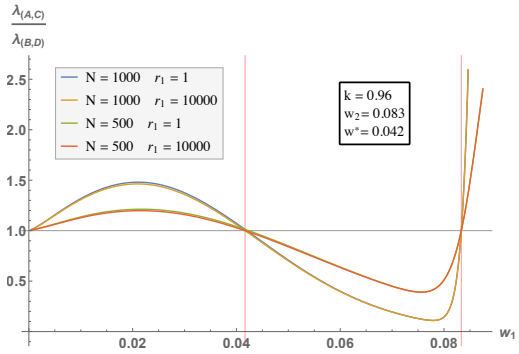

(b) Varying the strength of selection in population 1 when $k \in (0, 1)$ and $w_2 > w^*$.

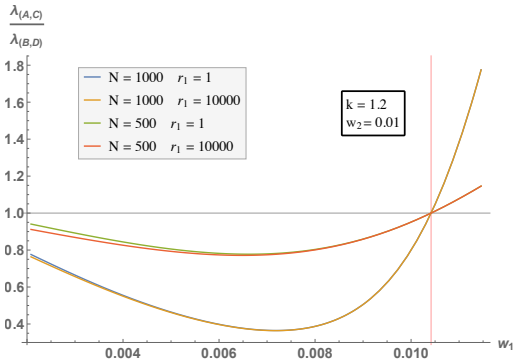

(c) Varying the strength of selection in population 1 when $k \in [1, 2)$.

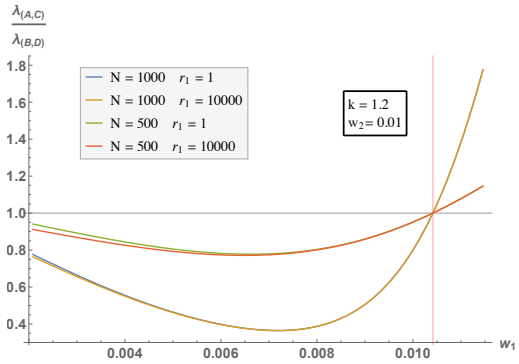

(d) Varying the strength of selection in population 1 when $k \in [1, 2)$.

Figure S8: The effect of changing the strength of selection in population 1, with selection strength in population 2 held constant, for different values of k. The two populations are of equal size. Again, note the insensitivity of effects to large changes in r_1.

effect of selection strength, we arrive at

$$\begin{aligned}
\lambda_{(A,C)} - \lambda_{(B,D)} &\overset{\text{sign}}{=} \frac{(1 - v_2^{-1}(\eta))}{1 - v_2^{-lN_0}(\eta)} \frac{d}{dl}\left(v_2^{lN_0-1}(\eta)\right) - \frac{(1 - v_1^{-1}(\eta))}{1 - v_1^{-lN_0}(\eta)} \frac{d}{dl}\left(v_1^{lN_0-1}(\eta)\right) \\
&\overset{\text{sign}}{=} \frac{(1 - v_2^{-1}(\eta))}{1 - v_2^{-lN_0}(\eta)} v_2^{lN_0-1}(\eta) \ln\left(v_2(\eta)\right) - \frac{(1 - v_1^{-1}(\eta))}{1 - v_1^{-lN_0}(\eta)} v_1^{lN_0-1}(\eta) \ln\left(v_1(\eta)\right) \\
&\overset{\text{sign}}{=} \frac{(1 - v_2^{-1}(\eta))}{1 - v_2^{-N}(\eta)} v_2^{N-1}(\eta) \ln\left(v_2(\eta)\right) - \frac{(1 - v_1^{-1}(\eta))}{1 - v_1^{-N}(\eta)} v_1^{N-1}(\eta) \ln\left(v_1(\eta)\right)
\end{aligned}$$

Setting $h(x) := \frac{1-x^{-1}}{1-x^{-N}} x^{N-1} \ln(x)$,

$$\lambda_{(A,C)} - \lambda_{(B,D)} \overset{\text{sign}}{=} h(v_2(\eta)) - h(v_1(\eta)).$$

For $0 < x < 1$

$$h^{'}(x) \overset{\text{sign}}{=} x^N(1-x)(1-x^N) - \ln(x)(2-x)(1-x^N) + N\ln(x)\left(\frac{1-x}{1-x^N}\right),$$

which is decreasing on any fixed interval $[x_1, x_2]$, where $0 < x_1 < x_2 < 1$, for N large enough. We use once again that $v_1(\eta) < v_2(\eta)$ when $\eta < 0$, while $v_1(\eta) > v_2(\eta)$ when $\eta > 0$.

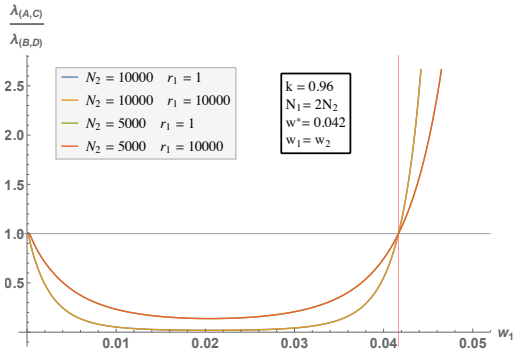

(a) A twofold difference in population sizes when $k \in (0, 1)$.

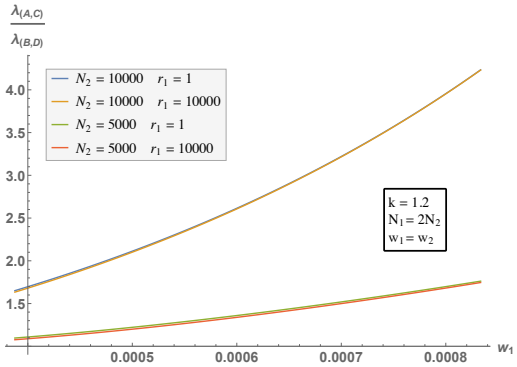

(b) A twofold difference in population sizes when $k \in [1, 2)$.

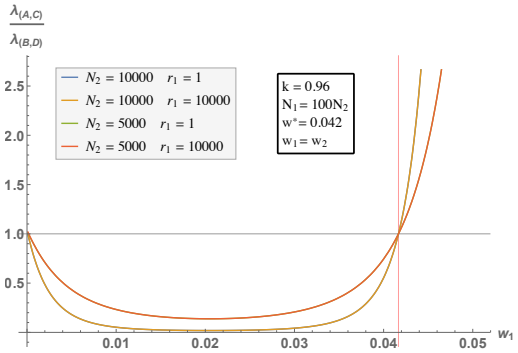

(c) A hundredfold difference in population sizes when $k \in (0, 1)$.

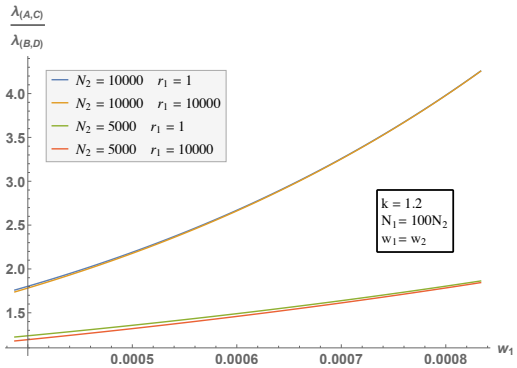

(d) A hundredfold difference in population sizes when $k \in (0, 1)$.

Figure S9: The effect of large differences in population size for different values of k. Note the insensitivity of effects to r_1.

In summary we find that if N is sufficiently large, and $N_1 = N + \Delta N > N = N_2$,

1. If $k \in (0, 1)$

$$\lambda_{(B,D)} < \lambda_{(A,C)} \qquad \text{if } w > \frac{1-k}{k}$$

That is, the increase in population 1's size is beneficial.

$$\lambda_{(B,D)} > \lambda_{(A,C)} \qquad \text{if } w < \frac{1-k}{k}$$

That is, the increase in population 1's size is detrimental.

2. If $k \in [1, 2)$:

$$\lambda_{(B,D)} < \lambda_{(A,C)}$$

That is, the increase in population 1's size is beneficial.

To see a graphical representation of this behaviour refer to Figures S7c and S7d.

As in the case of selection strength, numerical calculations suggest that for a large enough population size a generalization of the above findings hold. That is, for fixed $w_1 = w_2 = w$, the above summary holds for any $N_1 > N_2$ (N_2 large enough).

A graphical illustration of these results can be seen Figure S9.

S6.4 Summary of results

A summary of our weak-mutation results for the mutualism game is found in Fig. S10.

S7 Mutualistic symbioses, weak-selection limit

We have shown that, when mutation rates are very small, they (and generation times) have no effect on the stationary distribution of the evolutionary process. To explore how this result changes when mutation rates are allowed to be appreciably large, we resort to an alternative simplification, allowing mutation rates to be large, but forcing selection to be weak.

Since we are interested in the effect of mutation rates in the two populations, we set their sizes (N), selection strengths, and generation times equal. We make use of the two-population Moran process studied elsewhere in this paper, which is identical to that studied by Ohtsuki [8] when selection is weak. (In fact, Ohtsuki uses an exponential translation of payoffs to fitnesses, $f = \exp(w\pi)$, while we have used a linear translation, $f = 1 + w\pi$. These translations coincide in the weak-selection limit, $w \to 0$, for which Ohtsuki derives results, and which we consider here. This can easily be seen by Taylor expanding the exponential translation around $w = 0$. Therefore, Ohtsuki's weak-selection results hold also for the linear fitness translation.)

The evolutionary process is an irreducible, aperiodic Markov chain over the state space S of all possible population states. Ohtsuki [8] shows that, when the selection strength w is small ($w \to 0$), the stationary distribution of this process, $\lambda(s), s \in S$, simplifies significantly. Denote by $\langle x \rangle_w$ the expectation of x, taken with respect to the stationary distribution; i.e., $\langle x \rangle_w = \sum_{s \in S} x(s)\lambda(s)$. Then, for example, the long run frequency of members of population 1 who play strategy A is $\langle p_A \rangle_w = \sum_{s \in S} p_A(s)\lambda(s)$, where $p_A(s)$ is

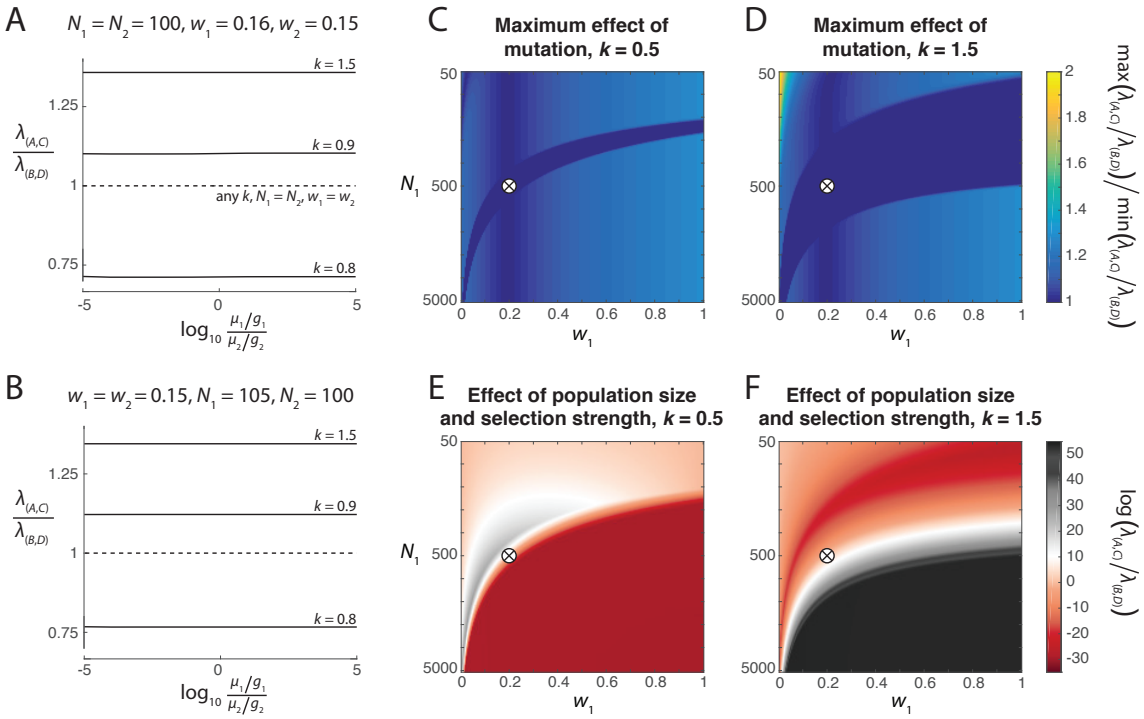

Figure S10: Long-run evolutionary dynamics of mutualisms, when mutations are rare ($\varepsilon \to 0$). (*A*) Setting the populations' sizes equal, we increase population 1's selection strength slightly above population 2's value of $w_2 = 0.15$. Then population 1 does better ($\lambda_{(A,C)} > \lambda_{(B,D)}$) when $k = 0.9$ [because then $w_2 = 0.15 > (1-k)/k = 0.11$], and when $k > 1.5$, but worse when $k = 0.8$ [because then $w_2 = 0.15 < (1-k)/k = 0.25$]. The populations do equally well when $N_1 = N_2$ and $w_1 = w_2$, no matter their relative mutation rates or generation times. (*B*) Setting the populations' selection strengths equal, we increase population 1's size slightly above population 2's. Again, population 1 benefits from its greater rate of evolution when $k = 1.5$ and $k = 0.9$, but not when $k = 0.8$. (*C, D*) We fix population 2's size at $N_2 = 500$ and selection strength at $w_2 = 0.2$ (marked by the cross), set equal population 1 and 2's generation times, and vary population 1's size N_1 and selection strength w_1. For each (N_1, w_1) combination, we vary the relative rate of mutation of the two populations across ten orders of magnitude, and calculate the ratio of the largest and smallest values of $\lambda_{(A,C)}/\lambda_{(B,D)}$, i.e., the maximum effect of mutation on the long-run relative success of the two populations. (*E, F*) For the same parameters as in (*C, D*), we set equal the population's mutation rates, and plot the natural logarithm of $\lambda_{(A,C)}/\lambda_{(B,D)}$ for each (N_1, w_1) combination.

the proportion of population 1 playing strategy A in population state s. For population 2, we denote strategy frequencies by q_C and q_D.

The quantities we are interested in are the long-run frequencies of (A, C) and (B, D) pairings/interactions. If the former frequency is larger than the latter, then we say that population 1 is more successful (recall that all other pairings yield the same payoff for the population 1 and 2 interactants). These quantities, $\langle p_A q_C \rangle_w$ and $\langle p_B q_D \rangle_w$, can be

calculated from Eq. (29) in ref. [8]. Writing $\hat{\mu}_1 = (N-1)\varepsilon\mu_1$ and $\hat{\mu}_2 = (N-1)\varepsilon\mu_2$,

$$\langle p_A q_C \rangle_w = \frac{1}{4}\left\{1 + w\frac{N-1}{(\varepsilon\mu_1+\varepsilon\mu_2)(1+\hat{\mu}_1)(1+\hat{\mu}_2)}\left[\varepsilon\mu_1(1-\varepsilon\mu_1)\frac{2-k}{2} + \varepsilon\mu_2(1-\varepsilon\mu_2)\frac{1}{2} + \frac{1-k}{4}(\varepsilon\mu_2-\varepsilon\mu_1)[1+\hat{\mu}_1+\hat{\mu}_2]\right]\right\};$$

$$\langle p_B q_D \rangle_w = \frac{1}{4}\left\{1 + w\frac{N-1}{(\varepsilon\mu_1+\varepsilon\mu_2)(1+\hat{\mu}_1)(1+\hat{\mu}_2)}\left[\varepsilon\mu_1(1-\varepsilon\mu_1)\frac{1}{2} + \varepsilon\mu_2(1-\varepsilon\mu_2)\frac{2-k}{2} - \frac{1-k}{4}(\varepsilon\mu_2-\varepsilon\mu_1)[1+\hat{\mu}_1+\hat{\mu}_2]\right]\right\}.$$

From these expressions, we calculate the long run advantage to population 1,

$$\begin{aligned}
\langle p_A q_C \rangle_w - \langle p_B q_D \rangle_w &= \frac{w}{4}\frac{N-1}{(\varepsilon\mu_1+\varepsilon\mu_2)(1+\hat{\mu}_1)(1+\hat{\mu}_2)}\left[\varepsilon\mu_1(1-\varepsilon\mu_1)\left(\frac{2-k}{2}-\frac{1}{2}\right)\right. \\
&\qquad \left. +\varepsilon\mu_2(1-\varepsilon\mu_2)\left(\frac{1}{2}-\frac{2-k}{2}\right) + \frac{1-k}{2}(\varepsilon\mu_2-\varepsilon\mu_1)(1+\hat{\mu}_1+\hat{\mu}_2)\right] \\
&= \frac{w}{4}\frac{N-1}{(\varepsilon\mu_1+\varepsilon\mu_2)(1+\hat{\mu}_1)(1+\hat{\mu}_2)}\left[\varepsilon\mu_1(1-\varepsilon\mu_1)\frac{1-k}{2} - \varepsilon\mu_2(1-\varepsilon\mu_2)\frac{1-k}{2}\right. \\
&\qquad \left. +\frac{1-k}{2}(\varepsilon\mu_2-\varepsilon\mu_1)[1+(N-1)(\varepsilon\mu_1+\varepsilon\mu_2)]\right] \\
&= \frac{w}{4}\frac{N-1}{(\varepsilon\mu_1+\varepsilon\mu_2)(1+\hat{\mu}_1)(1+\hat{\mu}_2)}\frac{1-k}{2}\left[N(\varepsilon\mu_2-\varepsilon\mu_1)(\varepsilon\mu_1+\varepsilon\mu_2)\right] \\
&= \frac{w}{8}\frac{N(N-1)}{(1+\hat{\mu}_1)(1+\hat{\mu}_2)}(1-k)\varepsilon(\mu_2-\mu_1) \\
&= A(1-k)\varepsilon(\mu_2-\mu_1),
\end{aligned}$$

where $A = \frac{w}{8}\frac{N(N-1)}{(1+\hat{\mu}_1)(1+\hat{\mu}_2)} > 0$. Therefore, if $k < 1$, population 1 does better ($\langle p_A q_C \rangle_w - \langle p_B q_D \rangle_w > 0$) when it has the smaller mutation rate, $\mu_1 < \mu_2$. On the other hand, when $k > 1$, population 1 does better when it has the larger mutation rate, $\mu_1 > \mu_2$. So slower evolution, in terms of mutation rates, is favored when $k < 1$, and faster evolution is favored when $k > 1$.

When mutation rates are low ($\varepsilon \approx 0$), changes in relative mutation rates have very little effect on $\langle p_A q_C \rangle_w - \langle p_B q_D \rangle_w$, consistent with our weak-mutation results above.

S8 Mutualistic symbioses with continuous strategy spaces

So far, we have considered discrete-strategy games, where members of each population choose between two strategy options (A or B for population 1, and C or D for population 2). In the discrete mutualism game that we have studied, this discreteness results in there being two stable equilibrium outcomes, one preferred by population 1, and the other preferred by population 2. In the introductory section of our Main Text, we have dis-

cussed examples of mutualistic interactions where this strategy discreteness clearly applies. However, in many mutualisms, strategies are of a more continuous nature; for example, *how much* energy does one species devote to a certain mutually beneficial task, instead of reserving this energy for uses beneficial only to itself? In this section, we provide a preliminary setup and analysis of how stochastic evolutionary dynamics apply to such continuous-strategy mutualisms.

As before, we seek a game that describes interactions between two individuals drawn from two different populations. The game we shall study is a continuous version of the Nash bargaining game, which nests the classical 'divide-the-dollar' version [7] as a special case. The two individuals choose some activity level, $x \in [0,1]$ for individuals from population 1 and $y \in [0,1]$ for individuals from population 2. Depending on the parameterization, these activity levels may be thought of as contributions to a public good, for example, or the amount extracted from a joint resource. Payoffs are given by

$$\pi_1(x,y) = \begin{cases} \alpha x + (1-\alpha)y & \text{if } x+y \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \pi_2(x,y) = \begin{cases} (1-\alpha)x + \alpha y & \text{if } x+y \leq 1 \\ 0 & \text{otherwise} \end{cases} . \tag{8}$$

The parameter α measures to what extent individuals benefit from their own activity levels. The assumption $\alpha \in (0,1)$ ensures that the interaction can be interpreted as mutualistic: provided that the sum of both players' activities is at most 1, increasing one individual's activity increases both individuals' payoffs. As in many other studied bargaining games, the restriction that payoffs are positive only if $x+y \leq 1$ permits a simple characterization of the equilibrium set: any combination (x,y) with $x+y = 1$ is an equilibrium of the game. If $\alpha < 1/2$, an equilibrium favors population 1 if $x < y$ (i.e., the equilibrium favors the population that shows the lower activity level). Conversely, if $\alpha > 1/2$, an equilibrium favors population 1 if $x > y$. In the limiting case $\alpha \to 1$, the interaction reduces to 'divide-the-dollar': there is a resource worth a total of one unit in payoffs, and the individuals' activities represent what fraction of this resource they demand for themselves (with the convention that when the summed demands exceed the whole resource, both individuals get nothing).

To model the evolutionary dynamics, we consider a weak-mutation scenario similar to that we have employed in the discrete-strategy case—mutations are assumed to be rare, so that populations are almost always monomorphic for some strategy. The key difference between the previous analyses and that here is that, while in the discrete 2-strategy case, a mutation simply means switching to the other strategy, in the continuous case we need to specify the distribution of possible mutations. We employ a 'local mutations' model [4].

From any state (x,y) in which populations 1 and 2 are monomorphic for strategies x and y respectively, a mutation occurs in one of the populations; as before, the probability that this happens in population l is proportional to $N_l \mu_l / g_l$. Mutations are 'local', that is, of small effect: if the mutation appears in population 1 when it is fixed for strategy x, then the mutant's strategy is $x' = x + u$, where u is taken from the uniform distribution

on $[-\delta, \delta]$ (if x' happens not to be in the unit interval, another mutant strategy is drawn). The fitness of a resident x-player and of the mutant x'-player are $f_1 = 1 + w_1\pi_1(x, y)$ and $f_1' = 1 + w_1\pi_1(x', y)$, respectively. If, instead, the mutant appears in population 2, then the respective fitnesses f_2 and f_2' are calculated analogously (with the same maximum mutation effect δ applying to both populations, though their mutation *rates* can differ). As in the model with only two strategies, the mutant strategy fixes in its population l with probability $(1 - f_l/f_l')/[1 - (f_l/f_l')^{N_l}]$. This elementary updating process then iterates: another mutation occurs in one of the two population, and it again either fixes or goes extinct. Initially, individuals from both populations are assumed to show no activity $x=y=0$. By iterating this process for a sufficient timespan (for the simulations shown in the following we introduce 10^9 mutant strategies in total), we approximate the distribution of strategies in the mutation-selection equilibrium.

The evolutionary dynamics of this model with continuous activity levels is different from the evolutionary dynamics of the discrete games studied in the Main Text, for at least three reasons: First, given our assumptions that mutant strategies are close to the strategies from which they mutated and that the payoff functions are continuous when total activity does not exceed 1, the difference between the fitness of a mutant strategy and of the strategy from which it mutated will usually be small, so that we are in a weak-selection regime [10]. Therefore, even mutants of reduced fitness may substitute at an appreciable rate. Second, while the discrete mutualism game had two distinct equilibria, in the continuous model there is a continuous path of equilibria (the line $x+y=1$); with local mutation and the associated weak selection, we may expect to see stochastic movement along this path. Third, in the discrete mutualism game, the Red King effect arose in a setting where it required *both* populations to evolve to reach an equilibrium—differences in their evolutionary rates then influenced the basins of attraction of the two equilibria. In the continuous strategy section that we are considering here, mutation and selection in one of the two populations would be sufficient to reach an equilibrium. For example, if population 1 has a very much shorter generation time than population 2, we may expect that, starting from the origin $(x, y) = (0, 0)$, the dynamics quickly move towards a state where individuals in population 1 show a much higher activity, $(x, y) \approx (1, 0)$. Whether we observe a Red King effect then depends on whether states close to $(1, 0)$ favor population 1 or population 2 (i.e., whether α is larger or smaller than $1/2$).

Figures S11 and S12 display simulation runs for two different values of α ($\alpha = 0.75$ and $\alpha = 0.25$). As for the discrete-strategy mutualism game studied in the Main Text, we consider how the dynamics are affected when the two populations differ in each of four different evolutionary rate parameters: generation time (first column in Figs S11 and S12), mutation rate (second column), selection strength (third column), and population size (fourth column). For all parameters, we assume that it is population 2 that evolves at a faster rate, either because it has a shorter generation time, a higher mutation rate, is under stronger selection, or is larger. We again distinguish between the short-run dynamics (here defined as the time it takes the populations to evolve to, or close to, one of the states on

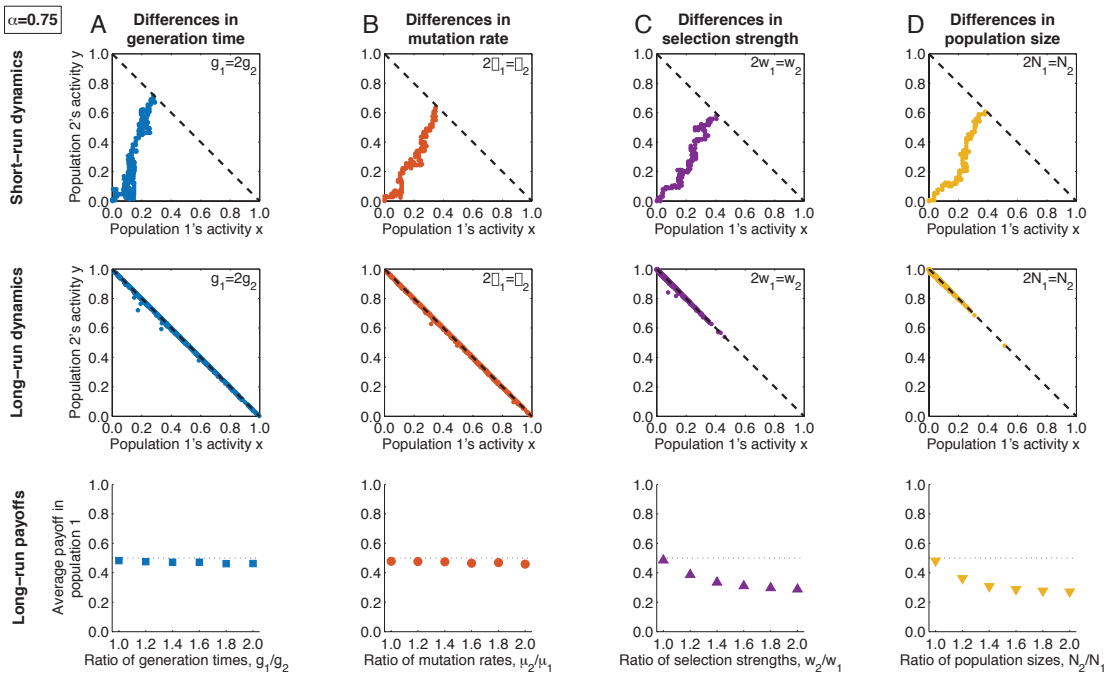

Figure S11: Short-run and long-run dynamics for continuous-strategy mutualisms, when α is large. Population 1 evolves slower owing to (A) a longer generation time, (B) a lower mutation rate, (C) weaker selection, and (D) a smaller population size. The top panels illustrate the short-run dynamics, showing typical evolutionary trajectories that the populations take from the initial state $(x, y) = (0, 0)$ to somewhere on or near the line of equilibria, $x+y = 1$ (the populations are considered near this line if $x + y \geq 0.98$). The middle panels show the long-run positions of the two populations over 10^9 elementary updating events. For clarity, these panels depict only those strategy combinations corresponding to the 10,000 most successful strategy combinations over time. The bottom panels depict the long-run average payoff of individuals in population 1, depending on how much the two populations differ in the relevant evolutionary rate parameter. Differences in generation time and mutation rate have little effect in the long run, but differences in selection strength and population size give rise to a strong Red Queen effect: the slower-evolving population 1 gets a lower share of the total payoff. This is similar to the results we obtained for the discrete-strategy mutualism game in the Main Text. Baseline parameters: $\alpha = 0.75$; $g_1 = g_2 = 1$, $\mu_1 = \mu_2 = 1$, $w_1 = w_2 = 0.5$ and $N_1 = N_2 = 100$; mutant strategies are at most a distance $\delta = 0.05$ from the strategy from which they mutated.

the equilibrium line) and the long-run dynamics (which corresponds to the 10^9 elementary updating events that we have simulated the evolutionary process for).

Independent of the value of α, and independent of the evolutionary rate parameter that is varied, we observe that the short-run dynamics are dominated by the population that evolves at a faster rate (see upper panels in Figs. S11 and S12). That is, when population 2 evolves faster, by the time the two populations reach the line of equilibria $x + y = 1$, we typically observe that $y > x$. The long-run dynamics depend on the evolutionary parameter that is varied (middle panels in Figs. S11 and S12). Differences

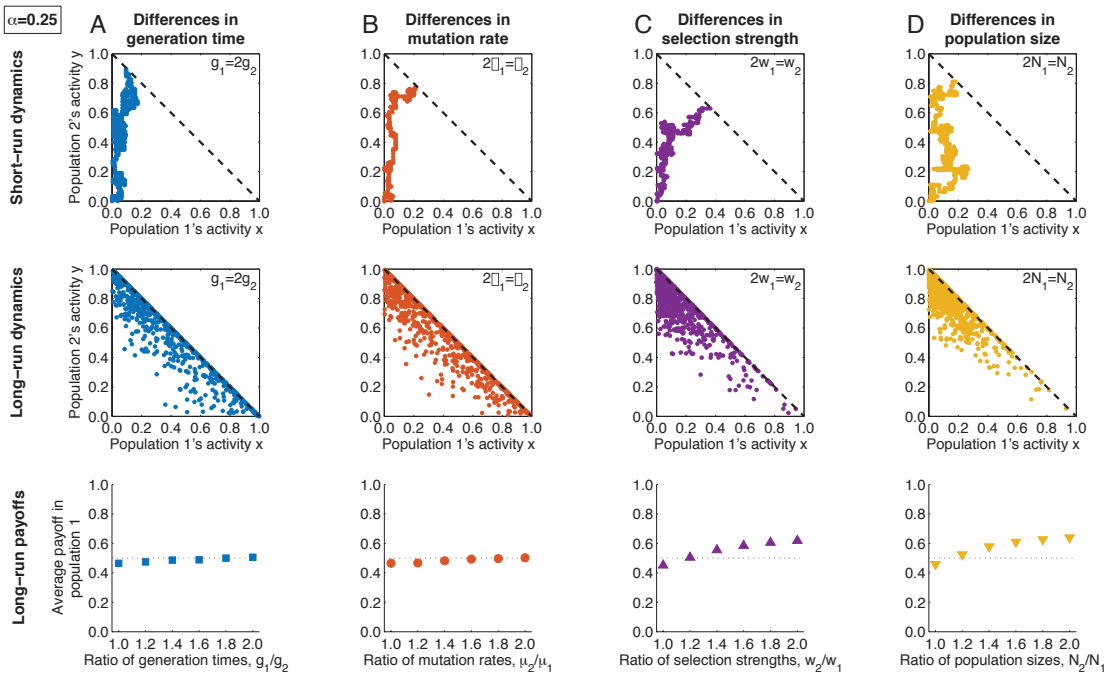

Figure S12: Short-run and long-run dynamics for mutualisms with a continuous trait space when α is small. The setup of the simulations is the same as in Fig. S11, except that here, $\alpha = 0.25$ (i.e., changes in own activity level usually have little influence on own payoff, but have a strong influence on payoffs of individuals in the other population). In this case, we observe a Red King effect when the populations differ in their selection strength or in their population size. Again, differences in generation time and mutation rate have little effect in the long run, similar to our results for the discrete-strategy mutualism game.

in generation time and mutation rate have little influence on the activity levels of the two populations, but differences in selection strength and population size create a bias towards higher activity levels in the faster-evolving population. To understand this, consider an arbitrary population state (x, y) on the equilibrium line (so that $x + y = 1$), and suppose that the two populations differ only in their selection strength: $w_2 > w_1$. In this case, the most likely evolutionary trajectory out of the present equilibrium state (x, y) involves the fixation in population 1 of a mutant with slightly reduced activity level $x' < x$ — such a mutant in population 1 is of reduced fitness, but is more likely to fix than an equivalent mutant in population 2 because selection is stronger in population 2 (and note that, in any population, a mutant with *increased* activity level is strongly selected against because it typically receives payoff 0). Then, again because $w_2 > w_1$, the next evolutionary step most likely involves population 2 slightly increasing its activity level y. Therefore, stochastic fluctuations off the equilibrium line tend to be in the direction of reduced x, and tend to return to the equilibrium line in the direction of increased y. In total, these two effects make the two populations move to the upper left corner, as depicted in the middle panels of Figs. S11 and S12. This is an example of drift-induced selection along an

equilibrium line that would otherwise be stationary under deterministic dynamics; drift-induced selection is a phenomenon of recent and growing interest in stochastic evolutionary dynamics (e.g., [2]).

In Fig. S11, we consider the case $\alpha = 0.75$, where the population that on average shows a higher activity level obtains higher average payoffs. We observe a Red Queen effect in this scenario when populations differ in their selection strengths or population sizes: the faster-evolving population 2 gets a higher share of the total payoff. In contrast, Fig. S12 shows the case $\alpha = 0.25$, where the population with the lower average activity level obtains higher average payoffs. Now, the slower-evolving population 1 outperforms population 2 when its slower evolution is due to reduced selection strength or smaller population size—a Red King effect. Again, differences in mutation rate and generation time have little effect. In addition, we note that for $\alpha = 0.25$, the two populations are more dispersed across the strategy space. Intuitively, because α is smaller, changes of own strategy within a population are under weaker selection, so that deviations below the equilibrium line are selected against less effectively.

References

[1] C. T. Bergstrom and M. Lachmann. The Red King effect: when the slowest runner wins the coevolutionary race. *Proceedings of the National Academy of Sciences*, 100(2):593–598, 2003.

[2] G. W. A. Constable, T. Rogers, A. J. McKane, and C. E. Tarnita. Demographic noise can reverse the direction of deterministic selection. *Proceedings of the National Academy of Sciences*, 113(32):E4745–E4754, 2016.

[3] J. B. S. Haldane. A mathematical theory of natural and artificial selection, part V: selection and mutation. *Mathematical Proceedings of the Cambridge Philosophical Society*, 23(07):838–844, 1927.

[4] L. A. Imhof and M. A. Nowak. Stochastic evolutionary dynamics of direct reciprocity. *Proceedings of the Royal Society of London B: Biological Sciences*, 277(1680):463–468, 2010.

[5] D. M. McCandlish, C. L. Epstein, and J. B. Plotkin. Formal properties of the probability of fixation: Identities, inequalities and approximations. *Theoretical population biology*, 99:98–113, 2015.

[6] D. M. McCandlish and A. Stoltzfus. Modeling evolution using the probability of fixation: history and implications. *The Quarterly review of biology*, 89(3):225–252, 2014.

[7] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

[8] H. Ohtsuki. Stochastic evolutionary dynamics of bimatrix games. *Journal of theoretical biology*, 264(1):136–142, 2010.

[9] C. Veller and L. K. Hayward. Finite-population evolution with rare mutations in asymmetric games. *Journal of Economic Theory*, 162:93–113, 2016.

[10] G. Wild and A. Traulsen. The different limits of weak selection and the evolutionary dynamics of finite populations. *Journal of Theoretical Biology*, 247(2):382–390, 2007.