

LIMIT CYCLES SPARKED BY MUTATION IN THE REPEATED PRISONER'S DILEMMA

DANIELLE F. P. TOUPO

*Center for Applied Mathematics, Cornell University
Ithaca, NY 14853, USA
dpt35@cornell.edu*

DAVID G. RAND

*Department of Psychology, Yale University
New Haven, CT 06511, USA
david.rand@yale.edu*

STEVEN H. STROGATZ

*Department of Mathematics, Cornell University
Ithaca, NY 14853, USA
strogatz@cornell.edu*

Received (to be inserted by publisher)

We explore a replicator-mutator model of the repeated Prisoner's Dilemma involving three strategies: always cooperate (ALLC), always defect (ALLD), and tit-for-tat (TFT). The dynamics resulting from single unidirectional mutations are considered, with detailed results presented for the mutations $TFT \rightarrow ALLC$ and $ALLD \rightarrow ALLC$. For certain combinations of parameters, given by the mutation rate μ and the complexity cost c of playing tit-for-tat, we find that the population settles into limit cycle oscillations, with the relative abundance of ALLC, ALLD, and TFT cycling periodically. Surprisingly, these oscillations can occur for unidirectional mutations between *any* two strategies. In each case, the limit cycles are created and destroyed by supercritical Hopf and homoclinic bifurcations, organized by a Bogdanov-Takens bifurcation. Our results suggest that stable oscillations are a robust aspect of a world of ALLC, ALLD, and costly TFT; the existence of cycles does not depend on the details of assumptions of how mutation is implemented.

Keywords: replicator-mutator model, evolutionary game theory, bifurcation analysis, limit cycle

1. Introduction

Cooperation, where individuals pay costs to benefit others, is a cornerstone of human civilization. By cooperating, people create value and thus increase the "size of the pie." This makes a group in which everyone cooperates better off than a group where everyone is selfish. Cooperation can be difficult to achieve, however, because creating that collective benefit is often individually costly. How, then, could the selfish process of natural selection give rise to such altruistic cooperation? Evolutionary game theorists have devoted a great deal of effort to answering this question [Axelrod, 1984; Boyd *et al.*, 2003; Boyd & Richerson, 1992; Fowler, 2005; Fu *et al.*, 2008; Fudenberg & Maskin, 1990; Hamilton, 1964; Hauert *et al.*, 2002; Helbing & Yu, 2009; Manapat *et al.*, 2012; May, 1987; McNamara *et al.*, 2008; Nakamaru *et al.*, 1997; Nowak, 2006; Nowak *et al.*, 2004; Panchanathan & Boyd, 2004; Rand & Nowak, 2012, 2013; Rand *et al.*, 2013; Szolnoki *et al.*, 2009; Szolnoki & Szabó, 2004; Tarnita *et al.*, 2009; Traulsen & Nowak, 2006; Trivers, 1971; Wedekind & Milinski, 2010].

The standard game-theoretic paradigm for studying cooperation is the Prisoner's Dilemma [Rapoport, 1965]. Two players simultaneously choose to either cooperate (C) or defect (D), and each receives a payoff depending on the two choices. Two cooperators both earn the reward of mutual aid R , while two defectors receive a punishment P . If one player cooperates while the other defects, the defector earns the temptation payoff T while the cooperator receives the sucker payoff S . If the relationship $T > R > P > S$ holds, the game is a Prisoner's Dilemma because mutual cooperation is better than mutual defection ($R > P$), but individually defectors always out-earn cooperators ($T > R$ if the partner cooperates, and $P > S$ if the partner defects). Thus the Prisoner's Dilemma captures the tension between individual and collective interests, the conundrum at the heart of cooperation.

To study the evolution of cooperation, evolutionary game theorists typically combine game theory with differential equations to create an evolutionary dynamic. The replicator equation is one of the most common such models [Hofbauer *et al.*, 1979]: strategies with above-average payoffs become more common over time while strategies with below-average payoffs become less common. As described above, defectors always out-earn cooperators. Thus in a simple world of pure cooperators versus pure defectors, evolution via the replicator equation always leads to the extinction of cooperation and a population made up solely of defectors.

How, then, do we explain the success of cooperation which is so evident in the world around us? Numerous mechanisms for the evolution of cooperation have been proposed [Nowak, 2006], and empirical evidence has been provided for their importance in human cooperation [Rand & Nowak, 2013]. Chief among these mechanisms is direct reciprocity, also known as reciprocal altruism [Axelrod, 1984; Nowak & Sigmund, 1992, 1993; Rand *et al.*, 2009; Trivers, 1971]: when agents interact repeatedly, evolution can favor cooperation. If I will only cooperate with you in the next period if you cooperate with me in the current period, cooperation can be the payoff-maximizing strategy (as long as the game continues to the next period with high enough probability).

Tit-for-tat (TFT) is the most well-known of these reciprocal strategies. TFT begins by cooperating, and then merely copies its opponent's move in the previous period. Thus cooperators receive cooperation and profit, whereas defectors receive defection and lose out. Moreover, a population in which everyone plays TFT is resistant to invasion by defectors. If a lone defector is introduced into such a population, that ALLD player receives a low payoff relative to the resident TFT players. As a result, selection disfavors the invader and TFT is evolutionarily stable: repeated interactions promote the evolution of cooperation.

Tit-for-tat, however, has an Achilles' heel [Nowak, 2006]: in some situations it can be invaded by kinder, gentler strategies [Boyd & Lorberbaum, 1987; Imhof *et al.*, 2005; van Veelen & García, 2010; van Veelen *et al.*, 2012]. To see this, imagine again a population where everyone plays TFT, except for one player who always cooperates (ALLC). Both the resident TFT players and the ALLC deviant cooperate in every round, and thus receive equal payoffs. This means that neutral drift (in the sense of population genetics) can allow ALLC to increase in frequency, a process called neutral invasion. A more serious weakness of the TFT players is that other factors can impose costs on them – costs which ALLC players avoid. For example, TFT is a more sophisticated strategy than ALLC, and hence may incur complexity costs. These costs arise because TFT needs to spend energy interpreting the other player's last move before it can respond appropriately, whereas ALLC is an unconditional strategy. On top of that, the vindictiveness of TFT can hurt it in noisy or error-prone environments. If players sometimes make mistakes and defect when they meant to cooperate, such a mistake would go unnoticed by ALLC, but would send two TFT players into a vendetta of retaliatory defections.

In the presence of such costs, the TFT residents are ultimately overtaken by the ALLC invaders. And there's the rub: once ALLC becomes sufficiently common at TFT's expense, it opens the door to invasion by nasty, uncooperative ALLD players, who can ruthlessly exploit ALLC and sweep through the population. Through this sort of scenario, TFT populations can eventually wind up succumbing to defection, suggesting that cooperation may be doomed even in repeated games.

Here we present a solution to this problem: we show that incorporating mutation into the replicator equation breaks ALLD's dominance over the evolutionary outcomes by giving rise to stable cycles involving substantial levels of cooperation. Surprisingly, it is not necessary for every strategy to mutate into every other strategy in order to get cyclical behavior, or even for ALLD to mutate into TFT. We find that adding

any single unidirectional pathway for mutation can lead to stable cycles. Thus we provide evidence that stable oscillations are a robust aspect of a world of ALLD, ALLC and costly TFT players, and that direct reciprocity can lead to substantial cooperation even in the face of invasion by unconditional cooperators.

1.1. Relation to previous work

To allow for mutation, the relevant mathematical setting needs to change from the replicator equation to the replicator-mutator equation [Nowak, 2006; Stadler & Schuster, 1992]. A number of previous authors have studied limit cycles in replicator-mutator equations, motivated by applications to language change [Mitchener & Nowak, 2004], autocatalytic chemical reaction networks [Stadler & Schuster, 1992], evolutionary games [Bladon *et al.*, 2010; Galla, 2011; Imhof *et al.*, 2005], and multi-agent decision making [Pais & Leonard, 2011; Pais *et al.*, 2012].

Our work is most closely related to that of Imhof *et al.* [2005]. They studied the evolutionary game dynamics of ALLD, ALLC, and TFT for finite populations, where stochastic effects become important. One of their most striking results was the observation of the phenomenon mentioned above – an evolutionary cycle that goes from ALLD to TFT to ALLC and back to ALLD again. Even more remarkably, they found that the cycle spends nearly all its time lingering in the vicinity of TFT, even though ALLD is a strict Nash equilibrium!

In the model considered by Imhof *et al.* [2005], the conditional strategy TFT was assumed to incur a complexity cost c , relative to the simpler unconditional strategies ALLC and ALLD. Mutation was assumed to be uniform: every strategy mutates to the other two with equal probability. More precisely, if $i \neq j$, strategy i mutates to strategy j with probability μ and stays the same with complementary probability $1 - 2\mu$.

For the deterministic case of infinitely large populations, Imhof *et al.* [2005] found that for certain combinations of parameters c and μ , the replicator-mutator equations have a stable limit cycle, corresponding to the evolutionary cycles observed in their simulations. They mentioned bifurcations associated with the birth and death of the limit cycle, but did not present a bifurcation analysis or a stability diagram to locate the bifurcation curves in the (μ, c) parameter space.

We were curious to learn more about the evolutionary cycle seen in the infinite population model. What bifurcations create and destroy this limit cycle? How does its bifurcation structure depend on the details of how the mutations are implemented, in a graph-theoretic sense? That is, if we think of the three strategies as the vertices of a triangle graph, with mutations occurring along the edges between them, the uniform mutation case studied by Imhof *et al.* [2005] amounts to a complete graph with equal weights μ on its six directed edges. What would happen, by contrast, in the opposite extreme case of unidirectional mutation along one of these six directed edges?

We found that for all six possible unidirectional mutations, a stable limit cycle exists in a certain part of (μ, c) space. The cycle always oscillates in the same rotational sense, moving from the neighborhood of ALLD to TFT to ALLC and back toward ALLD again, just as it does in the case of uniform mutation. Moreover, the commonalities extend to the types and locations of the bifurcations that create and destroy the cycle. The region in parameter space where stable limit cycles exist is always bounded on one side by a curve of supercritical Hopf bifurcations and on the other side by a curve of homoclinic bifurcations. These two curves meet tangentially at a Bogdanov-Takens point at one end of the stability region for the limit cycle. All of these statements are true of the uniform mutation case as well.

This paper is organized as follows. In the next section we review the formulation of the Prisoner's Dilemma and its associated replicator equation, followed by their generalization to the replicator-mutator equation. Then we focus on two special cases of unidirectional mutation, $\text{TFT} \rightarrow \text{ALLC}$ and $\text{ALLD} \rightarrow \text{ALLC}$, and summarize the results from the other four cases, as well as for the case of uniform mutation. The paper concludes with a conjecture and a brief discussion.

2. Model

2.1. Prisoner's Dilemma

Following Axelrod [1984], we fix the parameter values $T = 5, R = 3, P = 1, S = 0$. These satisfy the inequalities required for the game to qualify as a Prisoner's Dilemma: $T > R > P > S$ and $R > (T + S) / 2$. The final inequality implies that if the two players play many rounds with one another, it is better for both of them to cooperate all the time rather than engage in alternating bouts of getting suckered and suckering the other.

Now consider a repeated Prisoner's Dilemma among players using the strategies ALLC, ALLD, and TFT. Then, in the limit where the players meet infinitely often and there is no discounting of future interactions, their average payoffs are given by the payoff matrix given in Table 1 (the entries show the average payoff that the row player gets when playing against the specified column player):

Table 1. Payoff matrix of repeated prisoner's dilemma

	ALLD	TFT	ALLC
ALLD	P	P	T
TFT	P	R	R
ALLC	S	R	R

For example, an ALLC player gets suckered every time against an ALLD player, and therefore receives an average payoff of S in their depressing encounters. But when TFT plays ALLD, it gets suckered only on the first round, but after that it reciprocates each defection with defection, leading to an infinitely long string of mutual defections and hence an average payoff of P .

2.2. Replicator equation

Next we set the game in an evolutionary framework. If each strategy reproduces at a rate proportional to its relative fitness, the resulting dynamics can be approximated by the following set of ordinary differential equations, known as the replicator equations:

$$\begin{aligned}\dot{x} &= x(f_x - \phi) \\ \dot{y} &= y(f_y - \phi) \\ \dot{z} &= z(f_z - \phi).\end{aligned}\tag{1}$$

Here x, y and z denote the fractions of the population playing ALLD, TFT, and ALLC, respectively; f_i is the fitness of strategy i , defined as its expected payoff against the current mix of strategies; and

$$\phi = xf_x + yf_y + zf_z\tag{2}$$

is the average fitness in the whole population. By summing the differential equations for \dot{x}, \dot{y} , and \dot{z} , one can verify that $x + y + z = 1$ for all time, as required by the definition of x, y and z as relative frequencies.

The payoff matrix for the repeated Prisoner's Dilemma implies that the fitnesses are given by

$$\begin{aligned}f_x &= xP + yP + zT \\ f_y &= xP + yR + zR \\ f_z &= xS + yR + zR.\end{aligned}\tag{3}$$

For the parameter values $T = 5, R = 3, P = 1$, and $S = 0$ assumed above, and by replacing z with $1 - x - y$,

we find that the fitnesses reduce to

$$\begin{aligned} f_x &= 5 - 4x - 4y \\ f_y &= 3 - 2x \\ f_z &= 3 - 3x \end{aligned} \tag{4}$$

and the average fitness in the population becomes $\phi = 3 - x(1 + x + 3y)$.

The phase portrait for the replicator equation (1) with fitnesses (4) can be drawn in the (x, y) plane, by eliminating z via $z = 1 - x - y$. Unfortunately this way of presenting the phase portrait has certain disadvantages. It distorts the geometry of the trajectories and it arbitrarily highlights two of the strategies at the expense of the third. In the original (x, y, z) phase space, the phase portrait lives on the equilateral triangle defined by the face of the simplex $x + y + z = 1$, where $0 \leq x, y, z \leq 1$. Thus, a more appealing and symmetrical approach is to show the phase portrait as it actually appears on simplex. The following change of variables achieves this goal:

$$\begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \tag{5}$$

Equation (5) can be shown to be equivalent, up to a uniform scaling and a translation, to the change of variables given by Eq. 35 in Wesson & Rand [2013].

Figure 1 plots the resulting phase portrait. The first thing to note is that it contains a saddle point at $(x, y, z) = (1, 0, 0)$, corresponding to the entire population at ALLD. The fact that the ALLD corner is a saddle point, rather than a stable fixed point, reflects our simplifying assumptions that the game is infinitely repeated with no discounting of future interactions; in effect, the first round of the game is totally ignored. Without these assumptions, a standard result in repeated games [Nowak, 2006] is that there is bistability along the ALLD-TFT edge. In that case the ALLD corner would have a basin of attraction whose size depends on the payoffs.

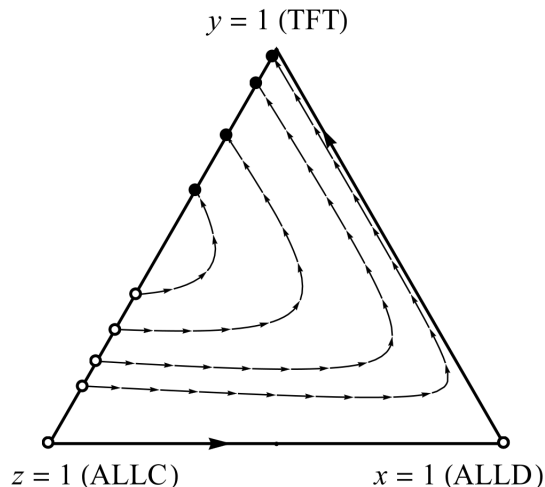


Fig. 1. Phase portrait of system (1) with $c = 0$ and $\mu = 0$.

Second, observe that Fig. 1 displays a line of neutrally stable fixed points along the side joining ALLC to TFT, where ALLD is absent. Neutral drift takes place along this side in finite populations. In the infinite population case shown here, the system almost always ends up in a nirvana of cooperation, with a mix of TFT and ALLC determined by the initial conditions, and with ALLD extinct.

However, the line of neutrally stable fixed points in this phase portrait is a structurally unstable feature. If the governing equations are perturbed by the addition of arbitrarily small terms, one expects that the line of fixed points will break and be replaced by something qualitatively different.

Table 2. Payoff matrix of replicator model with cost.

	ALLD	TFT	ALLC
ALLD	P	P	T
TFT	$P - c$	$R - c$	$R - c$
ALLC	S	R	R

Indeed, when we associate a small complexity cost c to playing TFT, the payoff matrix changes to that shown in Table 2:

The new fitnesses become

$$\begin{aligned} f_x &= 5 - 4x - 4y \\ f_y &= 3 - c - 2x \\ f_z &= 3 - 3x \end{aligned} \quad (6)$$

and

$$\phi = 3 - cy - x(1 + x + 3y). \quad (7)$$

In the corresponding phase portraits, shown in Fig. 2, the previous line of neutrally stable fixed points turns into an invariant line with the vector field flowing from a saddle point at TFT to another saddle point at ALLC.




Fig. 2. Phase portrait of system (1) with $c > 0$ and $\mu = 0$.

The structure of the rest of the phase portrait depends on the size of c , as shown in Fig. 2, but the long-term behavior of the system is the same in all three cases: ALLD takes over and cooperation dies out.

Notice that none of the phase portraits so far contain any limit cycles. This could have been anticipated from a general theorem that forbids limit cycles in any system of replicator equations involving $n = 3$ strategies, for any game and any payoff matrix [Hofbauer & Sigmund, 1998]. Periodic solutions can exist, but only within continuous families of neutrally stable cycles. Such periodic orbits are not isolated and hence do not qualify as limit cycles.

3. Replicator-Mutator Equation

Limit cycles do, however, become possible when we allow ALLD, ALLC, and TFT to mutate into one another. For simplicity, let us restrict attention to single unidirectional mutations only. Then there are six possibilities.

3.1. Example 1: TFT \rightarrow ALLC

In this case, we assume that after replication occurs, a player with strategy TFT mutates to ALLC with probability μ . Then the replicator-mutator system is

$$\begin{aligned}\dot{x} &= x(f_x - \phi) \\ \dot{y} &= y(f_y - \phi) - \mu y f_y.\end{aligned}\tag{8}$$

where x, y , and z denote the frequencies of ALLD, TFT, and ALLC, respectively. Note that as before, we have chosen to eliminate z from the equations by using the identity $z = 1 - x - y$. Likewise, we do not write the equation for \dot{z} explicitly, since it can be obtained from \dot{x} and \dot{y} if needed via $\dot{z} = -\dot{x} - \dot{y}$.

After insertion of (6) and (7) into (8), the replicator-mutator system becomes

$$\begin{aligned}\dot{x} &= x[(c-4)y + x(x+3y-3) + 2], \\ \dot{y} &= y[c(\mu+y-1) - 3\mu + x(2\mu+x+3y-1)].\end{aligned}\tag{9}$$

The right hand side is cubic in x and y . With the help of computer algebra, one can calculate the fixed points for the system, all but one of which lie on the boundary of the simplex (the equilateral triangle). Explicit formulas for these fixed points are presented in Appendix A. We have also calculated the curves in (μ, c) space at which the interior fixed point undergoes saddle-node and supercritical Hopf bifurcations; see Appendix A for details. As the parameters are varied, the stable limit cycles that are created by the supercritical Hopf bifurcation are later destroyed by a homoclinic bifurcation. The associated curve of homoclinic bifurcations was computed numerically with the help of the continuation package MATCONT [Govaerts & Kuznetsov, 2008].

Figure 3 plots the bifurcation curves in (μ, c) space. The saddle-node curve is shown in blue, the Hopf curve in red, and the homoclinic curve in green. These curves partition the parameter space into four regions, marked 1, 2, 3, and 4 on the figure.




Fig. 3. Stability diagram in (μ, c) space when TFT \rightarrow ALLC. Saddle-node curve, blue; supercritical Hopf curve, red; homoclinic curve, green; BT, Bogdanov-Takens point.

Figure 4 shows representative phase portraits for each of the four regions. In region 1, ALLD wipes out the other two strategies. (This makes sense intuitively. In region 1, TFT pays a high cost in fitness compared to the other two strategies. Furthermore, because μ is large in region 1, TFT mutates rapidly into ALLC, which in turn is clobbered by ALLD.) In region 2, a new pair of fixed points exist; they were created in a saddle-node bifurcation on the right-hand boundary of the equilateral triangle when the parameters crossed the saddle-node bifurcation curve. The unstable spiral seen in region 2 is the descendant of that node. Despite the existence of these two new fixed points, the system's long-term behavior remains the same as in region 1: almost all trajectories approach ALLD.

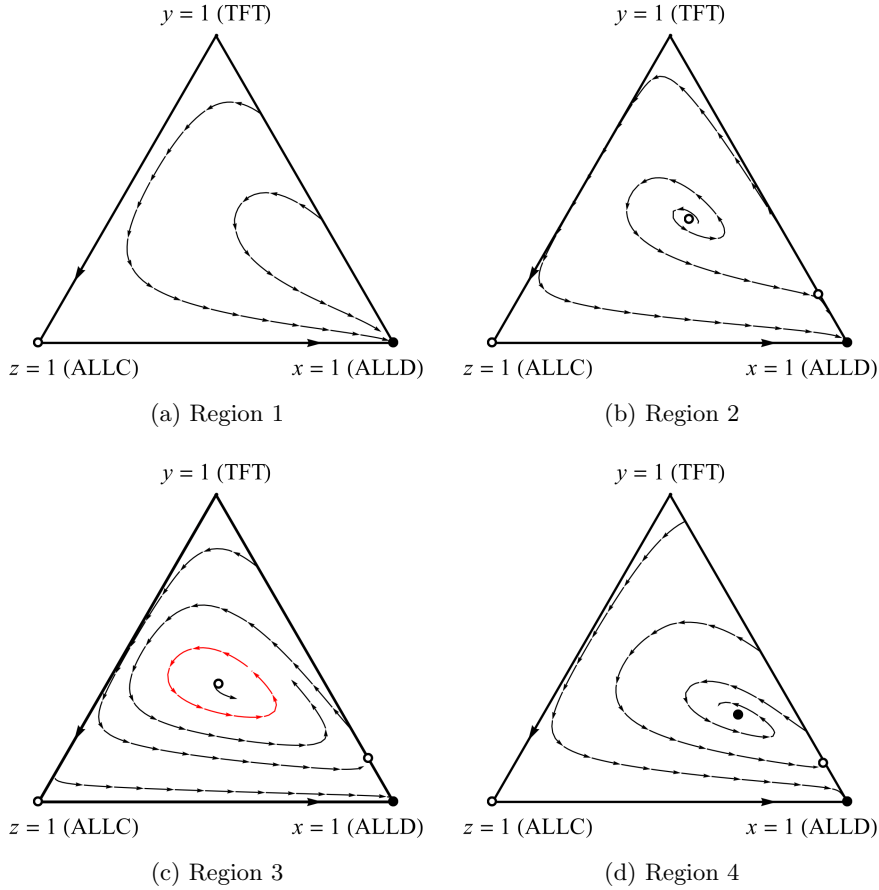


Fig. 4. Phase portrait of system (9) in different regions of (μ, c) space when $\text{TFT} \rightarrow \text{ALLC}$. The stable limit cycle is shown in red.

In region 3, however, a new attractor – a stable limit cycle, shown in red in Fig. 4(c) – coexists with ALLD. Where did this limit cycle come from? It emerged from a homoclinic orbit. When the parameters lie on the homoclinic bifurcation curve between region 2 and region 3, a homoclinic orbit starts and ends at the newly created saddle (the one close to the side between TFT and ALLD). This homoclinic orbit becomes a stable limit cycle when the parameters lie in region 3.

Finally, as we move from region 3 toward region 4, the limit cycle shrinks and ultimately becomes a point (a stable spiral) at the supercritical Hopf curve (the red curve in Fig. 3). When we move into the interior of region 4 the system becomes bistable, with the stable spiral sharing the state space with ALLD.

In biological terms, the population displays substantial levels of cooperation when it is on either the stable limit cycle of region 3 or the stable spiral of region 4.

3.2. *Example 2: ALLD \rightarrow ALLC*

Now we consider an alternative scenario of unidirectional mutation. Suppose that after replication occurs, each player with strategy ALLD mutates into ALLC with probability μ . Notice that this example shares one potentially important feature with Example 1 – in both cases, the target mutant being created is ALLC. On the other hand, this example differs from Example 1 in that the source of the mutant is now ALLD, not TFT. Which matters more: the commonality of the target or the non-commonality of its source?

It turns out that the commonality of the target is more important. To see this, observe that in the presence of ALLD \rightarrow ALLC mutations, the replicator-mutator system becomes

$$\begin{aligned}\dot{x} &= x(f_x - \phi) - \mu x f_x, \\ \dot{y} &= y(f_y - \phi)\end{aligned}\tag{10}$$

and $z = 1 - x - y$, where x , y , and z again denote the frequencies of ALLD, TFT, and ALLC, respectively. After insertion of (6) and (7) into (10), the replicator-mutator system yields

$$\begin{aligned}\dot{x} &= x[(c - 4)y - 5\mu + 4\mu(x + y) + x(x + 3y - 3) + 2], \\ \dot{y} &= y[c(y - 1) + x(x + 3y - 1)].\end{aligned}\tag{11}$$

Figure 5 plots the stability diagram in (μ, c) space. Note how much it resembles Fig. 3. In both cases, the parameter space divides into four regions bounded by curves of supercritical Hopf (red), saddle-node (blue) and homoclinic (green) bifurcations, all of which emerge from a Bogdanov-Takens point. The main difference is that the red curve of Hopf bifurcations goes through the origin in Fig. 3 and not in Fig. 5.




Fig. 5. Stability diagram in (μ, c) space when ALLD \rightarrow ALLC. Saddle-node curve, blue; supercritical Hopf curve, red; homoclinic curve, green; BT, Bogdanov-Takens point. Formulas for the Hopf and saddle-node bifurcation curves are given in Appendix B. The homoclinic curve was computed using MATCONT.

Figure 6 shows the phase portraits corresponding to the four regions. As in Example 1, defection dominates the long-term dynamics in regions 1 and 2, where a stable fixed point close to ALLD attracts almost all solutions. (Note that pure ALLD is no longer a fixed point, because of the assumed mutations from ALLD to ALLC. That is why the globally attracting fixed point lies between ALLD and ALLC. We will refer to this fixed point as “almost ALLD.”) In region 3, a stable limit cycle (shown in red) coexists with almost ALLD. In region 4, the limit cycle no longer exists; it has contracted to a stable spiral fixed point via the supercritical Hopf bifurcation between regions 3 and 4.

It is important to realize that the population exhibits large amounts of cooperation when it is in the stable spiral state or cycling periodically. This becomes clear when one looks at time series instead of phase portraits. Figure 7 shows the approach to a stable limit cycle in which ALLD is nearly absent except during brief spikes, whereas ALLC and TFT predominate for most of the cycle. Thus, the stable limit cycle signifies more than just an avoidance of all-out defection; it represents a state of significant cooperation.

3.3. Other single unidirectional mutations

Four other types of single unidirectional mutations are possible: ALLD \rightarrow TFT, ALLC \rightarrow TFT, ALLC \rightarrow ALLD, and TFT \rightarrow ALLD. Using the techniques above, we have analyzed these remaining cases completely. But rather than wade through the details, it seems clearer and more useful to summarize the main results, which are as follows.

For all four cases, the phase portraits and the bifurcation curves are more complicated than those presented in Examples 1 and 2. Nevertheless, all of them display stable limit cycles in some region of the parameter space. Specifically, the region in (μ, c) space where stable limit cycles exist is always bounded

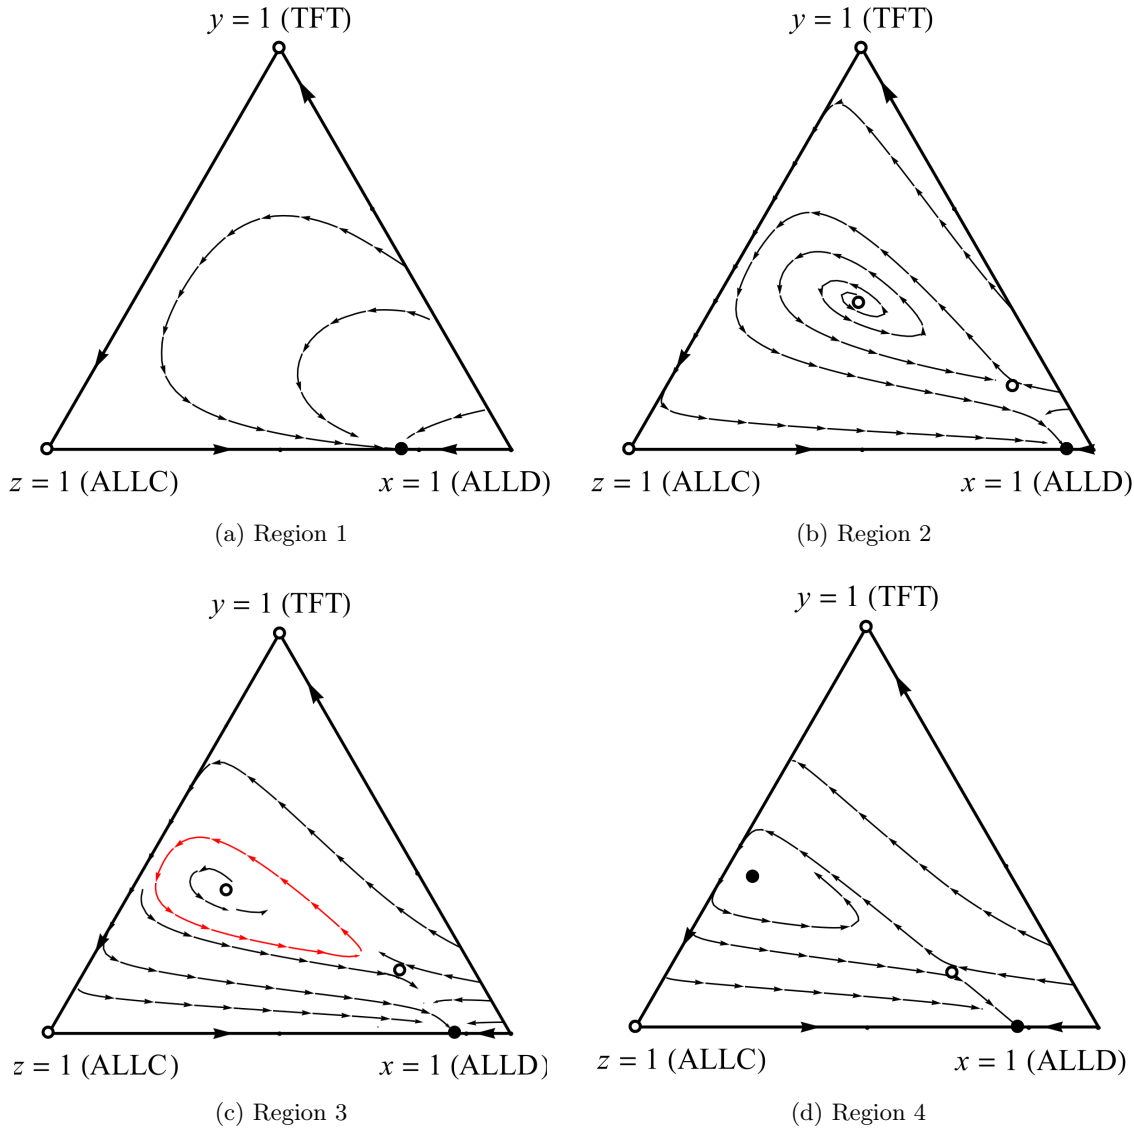


Fig. 6. Phase portrait of system (11) in different regions of (μ, c) space when ALLD \rightarrow ALLC. The stable limit cycle is shown in red.

on one side by a curve of supercritical Hopf bifurcations, and on the other side by a curve of homoclinic bifurcations. Both curves emanate from a Bogdanov-Takens point.

What we find most intriguing is that the homoclinic bifurcation curve – the counterpart of the green curve in Figs. 3 and 5 – always passes through the origin $(\mu, c) = (0, 0)$. Hence stable limit cycles always exist for arbitrarily small perturbations of the original system (1) with fitnesses (4), no matter how the mutation is implemented.

This is the sense in which limit cycles are “sparked by mutation” in the Prisoner’s Dilemma among ALLC, ALLD, and TFT. If just one of the strategies can mutate into just one of the others, it takes only an infinitesimal cost c and an infinitesimal mutation probability μ for the system to display self-sustained oscillations, with large amounts of cooperation during part of the cycle. In this way, the slightest bit of mutation allows the population to avoid a collapse into all-out defection.

3.4. *Uniform global mutation*

So far we have focused on a very restricted class of mutation pathways: single unidirectional mutations, in which exactly one strategy mutates into exactly one other. But we suspect that the sparking of limit

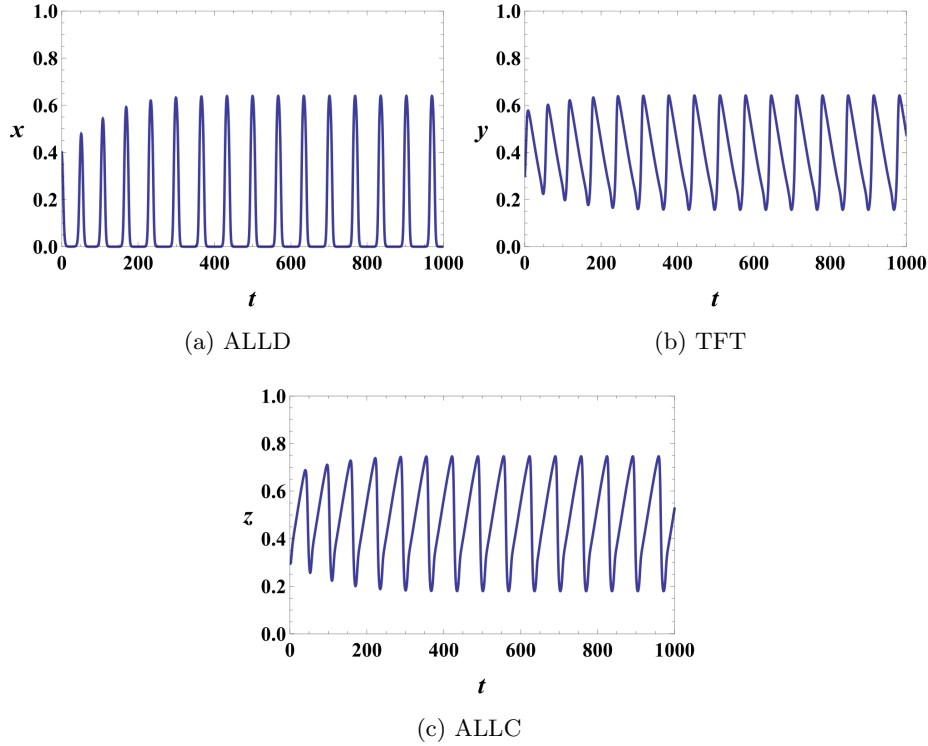


Fig. 7. Time series for a solution of system (11) as it approaches the stable limit cycle. Notice that the level of cooperation is high during much of the cycle, as shown by the high combined levels of ALLC and TFT, while ALLD remains low except for brief spikes. Parameter values: $(\mu, c) = (0.08, 0.04)$.

cycles by mutation is more general.

For example, consider the extreme opposite case of uniform global mutation, where each strategy mutates to the other two with probability μ , and hence stays the same with probability $1 - 2\mu$. The replicator-mutator equations for this case are given by

$$\begin{aligned}\dot{x} &= x [f_x (1 - 2\mu) - \phi] + \mu f_y y + \mu f_z z \\ \dot{y} &= y [f_y (1 - 2\mu) - \phi] + \mu f_x x + \mu f_z z\end{aligned}\quad (12)$$

where $z = 1 - x - y$. Substitution of (6) and (7) into (12) yields

$$\begin{aligned}\dot{x} &= \mu [x (11x + 9y - 16) - cy] + x [(c - 4)y + x(x + 3y - 3) + 2] + 3\mu, \\ \dot{y} &= cy(2\mu + y - 1) + 3\mu + x^2(y - \mu) + x(3y - 1)(\mu + y) - 9\mu y.\end{aligned}\quad (13)$$

Figure 8 plots the stability diagram in (μ, c) space, showing the curves where supercritical Hopf (red), saddle-node (blue) and homoclinic (green) bifurcations take place. The diagram splits into five regions. All the bifurcation curves of this system were computed using MATCONT [Govaerts & Kuznetsov, 2008].

Figure 9 shows the phase portraits corresponding to the five regions. As in the earlier examples 1 and 2, defection prevails in regions 1 and 2, where a stable fixed point close to ALLD attracts almost all solutions. In region 3, the possibility of cooperation reemerges: a stable limit cycle (shown in red) coexists with a stable node near ALLD. In region 4, the limit cycle becomes a global attractor; the stable node of region 3 is lost in a saddle-node bifurcation when the parameters move from region 3 to 4. The limit cycle no longer exists in region 5.

3.5. Conjecture

In every one of the examples considered so far, the region where stable limit cycles exist extends all the way down to the origin in parameter space. This means that limit cycles can be sparked by an arbitrarily small




Fig. 8. Stability diagram in (μ, c) space for the case of uniform global mutation. Saddle-node curve, blue; supercritical Hopf curve, red; homoclinic curve, green; BT, Bogdanov-Takens point; CP, cusp point. The bifurcation curves were computed using MATCONT.

mutation probability μ , if the complexity cost c of TFT is also suitably small. These results have been obtained for topologies at opposite ends of the graph-theoretic spectrum: single unidirectional mutation (in which mutation occurs along one directed edge) and uniform global mutation (in which mutation occurs in both directions along the edges of the complete graph).

We conjecture, therefore, that a similar sparking of stable limit cycles occurs for *any* pattern of mutation. To make this statement precise, we write down the general form of the replicator-mutator equation. Let the frequencies of the three strategies ALLD, TFT and ALLC be denoted as x_1, x_2, x_3 rather than x, y, z . Then the replicator-mutator dynamics are given by

$$\dot{x}_i = \left(\sum_{j=1}^3 x_j f_j Q_{ji} \right) - x_i \phi \quad (14)$$

for $i = 1, 2, 3$. Here f_i is the fitness of strategy i (obtained from Eq. (4), as before), $\phi = \sum_{j=1}^3 x_j f_j$ is the average fitness in the population, and Q_{ij} is the probability that players with strategy i mutate to playing strategy j . Since the Q_{ij} represent probabilities, they are non-negative and satisfy $\sum_j Q_{ij} = 1$; hence the matrix Q is row stochastic. In the limiting case where mutation does not occur, Q is the identity matrix and $Q_{ij} = \delta_{ij}$ where δ_{ij} is the Kronecker delta.

Now suppose that mutation does occur, and that it can be characterized by a mutation matrix M and a single parameter $0 < \mu \leq 1$ such that

$$Q_{ij} = \delta_{ij} - \mu M_{ij}. \quad (15)$$

Any matrix M that maintains the row-stochasticity and non-negativity of Q is admissible. This diversity of M is what we mean by any pattern of mutation.

Phrased in these terms, our conjecture is that for any fixed, admissible M , one can find an open set of parameters (μ, c) arbitrarily close to $(0, 0)$ such that stable limit cycles exist for the replicator-mutator system (14).

We have verified this conjecture informally for a handful of admissible mutation matrices M , but we have not done a comprehensive numerical study nor do we have analytical evidence to support the conjecture. So it could well be false. But we suspect it is true. A key step toward proving it would be to demonstrate that a curve of homoclinic bifurcations (the counterpart of the green curves in Figs. 3, 5, and 8) always emanates from the point $(\mu, c) = (0, 0)$, for any admissible M .

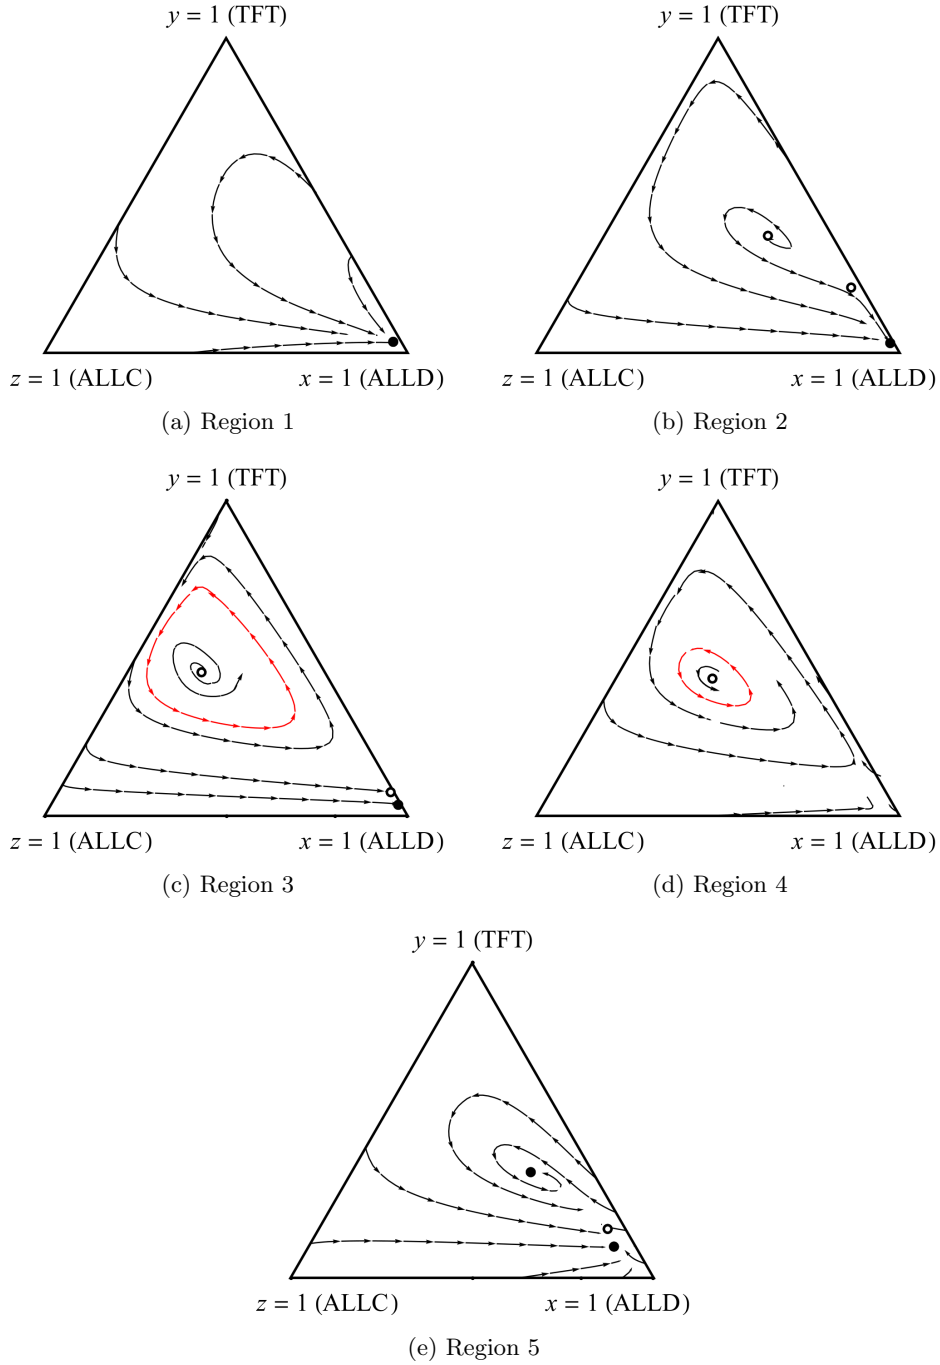


Fig. 9. Phase portrait of system (13) in different regions of (μ, c) parameter space. The stable limit cycle is shown in red.

4. Discussion

Our results show that a variety of different mutation structures give rise to evolutionary cycles of cooperation and defection in the repeated Prisoner's Dilemma. These cycles appear to be a robust feature of interactions in well-mixed populations of ALLC, ALLD and TFT, rather than being specific to particular assumptions regarding the mutation matrix. These cycles, as well as the stable spirals created by mutation in other parts of the (μ, c) parameter space, lead to substantial levels of cooperation. Thus the grim picture painted by the standard replicator equation without mutation, in which ALLD dominates even in repeated games, may be too pessimistic.

The replicator equation assumes that the population is well mixed. In some settings it is more realistic to regard the population as spatially structured, with players interacting only with a fixed subset of others on a regular lattice. Evolutionary cycles of cooperation and defection among ALLC, ALLD, and TFT can occur in this case too, even in the absence of mutation [Szolnoki *et al.*, 2009]. Similar cycles occur for populations playing a spatial version of the rock-paper-scissors game [Szolnoki & Szabó, 2004].

Although the replicator equation is typically thought of as describing genetic evolution, it can just as well describe a social learning dynamic in which people imitate the strategies of more successful others. In this context, mutation corresponds to experimentation or innovation, i.e., trying out a strategy other than the one which is performing well at the moment. Experimentation is a key element of human behavior [Rand *et al.*, 2013; Traulsen *et al.*, 2010], and our results suggest that it may also help to explain the success of cooperation in human societies.

5. Acknowledgments

Research supported in part by a Sloan-Colman fellowship to D.F.P.T. and a grant from the John Templeton Foundation through a subaward from the New Paths to Purpose Project at Chicago Booth to D.G.R. We thank Richard Rand and Elizabeth Wesson for helpful discussions.

References

- Axelrod, R. [1984] *The Evolution of Cooperation* (Basic Books, New York).
- Bladon, A. J., Galla, T., & McKane, A. J. [2010] “Evolutionary dynamics, intrinsic noise, and cycles of cooperation,” *Phys. Rev. E* **81**, 066122.
- Boyd, R. & Lorberbaum, J. P. [1987] “No pure strategy is evolutionarily stable in the repeated Prisoner’s Dilemma game,” *Nature* **327**, 58–59.
- Boyd, R. & Richerson, P. [1992] “Punishment allows the evolution of cooperation (or anything else) in sizable groups,” *Ethology and Sociobiology* **13**, 171–195.
- Boyd, R., Gintis, H., Bowles, S., & Richerson, P. J. [2003] “The evolution of altruistic punishment,” *Proc. Natl. Acad. Sci. USA* **100**, 3531–3535.
- Fowler, J. H. [2005] “Altruistic punishment and the origin of cooperation,” *Proc. Natl. Acad. Sci. USA* **102**, 7047–7049.
- Fu, F., Hauert, C., Nowak, M. A., & Wang, L. [2008] “Reputation-based partner choice promotes cooperation in social networks,” *Phys. Rev. E* **78**, 026117.
- Fudenberg, D., & Maskin, E. S. [1990] “Evolution and cooperation in noisy repeated games,” *Amer. Econ. Rev.* **80**, 274–279.
- Galla, T. [2011] “Cycles of cooperation and defection in imperfect learning,” *J. Statist. Mech: Theory and Experiment* **8**, P08007.
- Govaerts, W. & Kuznetsov, Yu. A. [2008] MATCONT, <http://www.matcont.ugent.be>.
- Hamilton, W. D. [1964] “The genetical evolution of social behaviour,” *J. Theor. Biol.* **7**, 1–16.
- Hauert, C., De Monte, S., Hofbauer, J., & Sigmund, K. [2002] “Volunteering as red queen mechanism for cooperation in public goods games,” *Science* **296**, 1129–1132.
- Helbing, D., & Yu, W. [2009] “The outbreak of cooperation among success-driven individuals under noisy conditions,” *Proc. Natl. Acad. Sci. USA* **106**, 3680–3685.
- Hofbauer, J., Schuster, P., & Sigmund, K. [1979] “A note on evolutionary stable strategies and game dynamics,” *J. Theor. Biol.* **81**, 609–612.
- Hofbauer, J. & Sigmund, K. [1998] *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, UK), p. 78.
- Imhof, L. A., Fudenberg, D., & Nowak, M. A. [2005] “Evolutionary cycles of cooperation and defection,” *Proc. Natl. Acad. Sci. USA* **102**, 10797–10800.
- Manapat, M. L., Nowak, M. A., & Rand, D. G. [2013] “Information, irrationality and the evolution of trust,” *J. Econ. Behav. Organ.* **90**, S57–S75.
- May, R. M. [1987] “More evolution of cooperation,” *Nature* **327**, 15–17.
- McNamara, J. M., Barta, Z., Fromhage, L., & Houston, A. I. [2008] “The coevolution of choosiness and cooperation,” *Nature* **451**, 189–192.
- Mitchener, W. G. & Nowak, M. A. [2004] “Chaos and language,” *Proc. R. Soc. Lond. B* **271**, 701–704.
- Nakamaru, M., Matsuda, H., & Iwasa, Y. [1997] “The evolution of cooperation in a lattice-structured population,” *J. Theor. Biol.* **184**, 65–81.
- Nowak, M. A. [2006] *Evolutionary Dynamics: Exploring the Equations of Life* (Harvard University Press, Cambridge, MA).
- Nowak, M. A., Sasaki, A., Taylor, C., & Fudenberg, D. [2004] “Emergence of cooperation and evolutionary stability in finite populations,” *Nature* **428**, 646–650.
- Nowak, M. A. & Sigmund, K. [1992] “Tit for tat in heterogeneous populations,” *Nature* **355**, 250–253.
- Nowak, M. A. & Sigmund, K. [1993] “A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game,” *Nature* **364**, 56–58.
- Pais, D. & Leonard, N. E. [2011] “Limit cycles in replicator-mutator network dynamics,” *Proc. 50th IEEE Conference on Decision and Control*, pp. 3922–3927.
- Pais, D., Caicedo-Nunez, C. H., & Leonard, N. E. [2012] “Hopf bifurcations and limit cycles in evolutionary network dynamics,” *SIAM J. Appl. Dyn. Syst.* **11**, 1754–1784.
- Panchanathan, K. & Boyd, R. [2004] “Indirect reciprocity can stabilize cooperation without the second-order free rider problem,” *Nature* **432**, 499–502.

- Rand, D. G., Ohtsuki, H., & Nowak, M. A. [2009] “Direct reciprocity with costly punishment: generous tit-for-tat prevails,” *J. Theor. Biol.* **256**, 45–57.
- Rand, D. G. & Nowak, M. A. [2012] “Evolutionary dynamics in finite populations can explain the full range of cooperative behaviors observed in the centipede game,” *J. Theor. Biol.* **300**, 212–221.
- Rand, D. G. & Nowak, M. A. [2013] “Human cooperation,” *Trends in Cognitive Sciences* **17**, 413–425.
- Rand, D. G., Tarnita, C. E., Ohtsuki, H., & Nowak, M. A. [2013] “Evolution of fairness in the one-shot anonymous ultimatum game,” *Proc. Natl. Acad. Sci. USA* **110**, 2581–2586.
- Rapoport, A. and Chammah, A. M. [1965] *Prisoner’s Dilemma: A Study in Conflict and Cooperation* (University of Michigan Press, Ann Arbor).
- Stadler, P. S. & Schuster, P. [1992] “Mutation in autocatalytic networks: An analysis based on perturbation theory,” *J. Math. Biol.* **30**, 597–631.
- Strogatz, S. H. [1994] *Nonlinear Dynamics and Chaos: with Applications to Physics, Biology, Chemistry, and Engineering* (Addison-Wesley, Reading, Massachusetts).
- Szolnoki, A., Perc, M., & Szabó, G. [2009] “Phase diagrams for three-strategy evolutionary prisoner’s dilemma games on regular graphs,” *Phys. Rev. E* **80**, 056104.
- Szolnoki, A. & Szabó, G. [2004] “Phase transitions for rock-scissors-paper game on different networks,” *Phys. Rev. E* **70**, 037102.
- Tarnita, C. E., Antal, T., Ohtsuki, H., & Nowak, M. A. [2009] “Evolutionary dynamics in set structured populations,” *Proc. Natl. Acad. Sci. USA* **106**, 8601–8604. doi: 10.1073/pnas.0903019106.
- Traulsen, A., & Nowak, M. A. [2006] “Evolution of cooperation by multilevel selection,” *Proc. Natl. Acad. Sci. USA* **103**, 10952–10955.
- Traulsen, A., Semmann, D., Sommerfeld, R. D., Krambeck, H.-J., & Milinski, M. [2010] “Human strategy updating in evolutionary games,” *Proc. Natl. Acad. Sci. USA* **107**, 2962–2966.
- Trivers, R. [1971] “The evolution of reciprocal altruism,” *Quart. Rev. Biol.* **46**, 35–57.
- van Veelen, M. & García, J. [2010] “In and out of equilibrium: evolution of strategies in repeated games with discounting,” *TI Discussion Papers*, (10-037/1).
- van Veelen, M., García, J., Rand, D. G., & Nowak, M. A. [2012] “Direct reciprocity in structured populations,” *Proc. Natl. Acad. Sci. USA* **109**, 9929–9934.
- Wedekind, C. & Milinski, M. [2000] “Cooperation through image scoring in humans,” *Science* **288**, 850–852.
- Wesson, E. N. & Rand, R. H. [2013] “Alternate models of replicator dynamics,” *J. Appl. Nonlin. Dyn.* **2**, 193–206.

Appendix A TFT \rightarrow ALLC

For Example 1 of the main text, in which TFT mutates into ALLC, the system (8) has at most 4 fixed points (x^*, y^*) in the simplex $x + y + z = 1$, with $0 \leq x, y, z \leq 1$:

$(x_1^*, y_1^*) = (0, 0)$ is a saddle,

$(x_2^*, y_2^*) = (1, 0)$ is a stable node,

$(x_3^*, y_3^*) = \left(\frac{A_1 - 5c\mu + c + 17\mu + 2}{12\mu + 4}, \frac{-(\mu + 1)A_1 - c\mu^2 + 8c\mu + c + \mu^2 - \mu + 2}{24\mu + 8} \right)$ is a saddle, and

$(x_4^*, y_4^*) = \left(\frac{-A_1 - 5c\mu + c + 17\mu + 2}{12\mu + 4}, \frac{(\mu + 1)A_1 - c\mu^2 + 8c\mu + c + \mu^2 - \mu + 2}{24\mu + 8} \right)$, could be a stable or unstable spiral (or node), depending on μ and c [Strogatz, 1994], where in the formulas above,

$$A_1 = \sqrt{c^2(\mu + 3)^2 - 2c((\mu - 11)\mu + 6) + (\mu - 28)\mu + 4}.$$

The fixed points (x_1^*, y_1^*) , (x_2^*, y_2^*) , and (x_3^*, y_3^*) are on the boundary of the simplex and (x_4^*, y_4^*) is inside the simplex. A saddle-node bifurcation occurs when (x_3^*, y_3^*) and (x_4^*, y_4^*) coalesce and disappear. The equation of the saddle-node bifurcation curve is

$$c = \frac{(\mu - 11)\mu - 4\sqrt{6}\sqrt{\mu(3\mu + 1)} + 6}{(\mu + 3)^2}. \quad (\text{A.1})$$

The fixed point (x_4^*, y_4^*) undergoes a supercritical Hopf bifurcation and switches from a stable spiral to an unstable spiral at certain parameter values. The equation of the Hopf bifurcation curve, which was computed analytically, is

$$c = -\frac{3 - 13\mu}{4(\mu - 1)} - \frac{1}{2} \sqrt{A_2 + \frac{A_4}{3\sqrt[3]{2}(\mu - 1)^2} + \frac{A_5}{3(\mu - 1)^2 A_4}} \quad (\text{A.2})$$

$$+ \frac{1}{2} \sqrt{A_2 - \frac{A_4}{3\sqrt[3]{2}(\mu - 1)^2} - \frac{A_5}{3(\mu - 1)^2 A_4} + \frac{A_6}{4\sqrt[4]{A_7 + \frac{A_4}{3\sqrt[3]{2}(\mu - 1)^2} + \frac{A_5}{3(\mu - 1)^2 A_4}}}} \quad (\text{A.3})$$

where

$$A_2 = \frac{(3 - 13\mu)^2}{4(\mu - 1)^2} + \frac{35\mu^2 - 34\mu - 8}{3(\mu^2 - 2\mu + 1)} - \frac{35\mu^2 - 34\mu - 8}{(\mu - 1)^2},$$

$$A_3 = (-5240604096\mu^{12} - 40578465024\mu^{11} + 200756188800\mu^{10} - 265354820640\mu^9 + 60401533248\mu^8 + 110419920576\mu^7 - 55059298560\mu^6 - 12327872544\mu^5 + 4987630080\mu^4 + 1779338880\mu^3 + 218439936\mu^2 - 110592\mu - 1880064)^{\frac{1}{2}},$$

$$A_4 = (204136\mu^6 - 578832\mu^5 + 416952\mu^4 + 51238\mu^3 - 79740\mu^2 - 12984\mu + A_3 - 1456)^{\frac{1}{3}},$$

$$A_5 = \sqrt[3]{2} (2272\mu^4 - 3160\mu^3 + 521\mu^2 + 316\mu + 100),$$

$$A_6 = -\frac{(3 - 13\mu)^3}{(\mu - 1)^3} + \frac{4(35\mu^2 - 34\mu - 8)(3 - 13\mu)}{(\mu - 1)^3} - \frac{8(49\mu^2 - 4\mu + 4)}{(\mu - 1)^2}, \text{ and}$$

$$A_7 = \frac{35\mu^2 - 34\mu - 8}{3(\mu^2 - 2\mu + 1)} - \frac{35\mu^2 - 34\mu - 8}{(\mu - 1)^2}.$$

Appendix B ALLD \rightarrow ALLC

For Example 2 of the main text, in which ALLD mutates into ALLC, system (10) has at most 5 fixed points (x^*, y^*) in the simplex $x + y + z = 1$, with $0 \leq x, y, z \leq 1$:

$(x_1^*, y_1^*) = (0, 0)$ is a saddle,

$(x_2^*, y_2^*) = (0, 1)$ is a saddle,

$(x_3^*, y_3^*) = \left(\frac{1}{2} \left(-4\mu - \sqrt{4\mu(4\mu - 1) + 1} + 3 \right), 0 \right)$ is a stable node,

$(x_4^*, y_4^*) = \left(-\frac{4c\mu + A_8 + c - 11\mu + 2}{16\mu - 4}, \frac{8c\mu^2 - 10c\mu + (2\mu - 1)A_8 + c + 18\mu^2 - 11\mu + 2}{8(\mu - 1)(4\mu - 1)} \right)$ is a saddle,

and

$(x_5^*, y_5^*) = \left(\frac{-4c\mu + A_8 - c + 11\mu - 2}{16\mu - 4}, \frac{8c\mu^2 - 10c\mu - (2\mu - 1)A_8 + c + 18\mu^2 - 11\mu + 2}{8(\mu - 1)(4\mu - 1)} \right)$ could be a stable or unstable spiral (or node), depending on μ and c . In the expressions above,

$$A_8 = \sqrt{(4c\mu + c - 11\mu + 2)^2 + 8c(4\mu - 1)(-c + \mu + 2)}.$$

The fixed points (x_1^*, y_1^*) , (x_2^*, y_2^*) , and (x_3^*, y_3^*) are on the boundary of the simplex, and (x_4^*, y_4^*) and (x_5^*, y_5^*) are in the interior. A saddle-node bifurcation occurs when (x_4^*, y_4^*) and (x_5^*, y_5^*) coalesce and disappear. The equation of the saddle-node bifurcation curve is

$$c = \frac{\mu(28\mu - 25) + 6 - 4\sqrt{6}\sqrt{(1 - \mu)\mu(3\mu - 2)(4\mu - 1)}}{(3 - 4\mu)^2}. \quad (\text{B.1})$$

The fixed point (x_5^*, y_5^*) undergoes a supercritical Hopf bifurcation and switches from a stable spiral to an unstable spiral at the curve given by

$$c = -\frac{(1 - i\sqrt{3})A_{10}}{6\sqrt[3]{2}} + \frac{(1 + i\sqrt{3})(48\mu^3 - 201\mu^2 + 114\mu - 33)}{3 \cdot 2^{2/3}A_{10}} + 2\mu + 1 \quad (\text{B.2})$$

where

$$A_9 = \sqrt{4(48\mu^3 - 201\mu^2 + 114\mu - 33)^3 + (-2592\mu^4 + 5670\mu^3 - 5022\mu^2 + 810\mu + 162)^2}$$

and

$$A_{10} = \sqrt[3]{-2592\mu^4 + 5670\mu^3 - 5022\mu^2 + A_9 + 810\mu + 162}.$$