

Some dynamics of signaling games

Simon Huttegger^a, Brian Skyrms^{a,1}, Pierre Tarrès^b, and Elliott Wagner^c

^aDepartment of Logic and Philosophy of Science, University of California, Irvine, CA 92697; ^bDepartment of Mathematics, University of Toulouse, 31000 Toulouse, France; and ^cDepartment of Philosophy, Kansas State University, Manhattan, KS 66506

Edited by John C. Avise, University of California, Irvine, CA, and approved April 4, 2014 (received for review February 7, 2014)

Information transfer is a basic feature of life that includes signaling within and between organisms. Owing to its interactive nature, signaling can be investigated by using game theory. Game theoretic models of signaling have a long tradition in biology, economics, and philosophy. For a long time the analyses of these games has mostly relied on using static equilibrium concepts such as Pareto optimal Nash equilibria or evolutionarily stable strategies. More recently signaling games of various types have been investigated with the help of game dynamics, which includes dynamical models of evolution and individual learning. A dynamical analysis leads to more nuanced conclusions as to the outcomes of signaling interactions. Here we explore different kinds of signaling games that range from interactions without conflicts of interest between the players to interactions where their interests are seriously misaligned. We consider these games within the context of evolutionary dynamics (both infinite and finite population models) and learning dynamics (reinforcement learning). Some results are specific features of a particular dynamical model, whereas others turn out to be quite robust across different models. This suggests that there are certain qualitative aspects that are common to many real-world signaling interactions.

costly signaling | replicator dynamics | Moran process

The flow of information is a central issue across the biological and social sciences. In both of those domains, entities have information that can be communicated, wholly or partly, to other entities by means of signals. Signaling games are abstractions that are useful for studying general aspects of such interactions. The simplest signaling games model interactions between two individuals: a sender and a receiver. The sender acquires private information about the state of the world and contingent on that information selects a signal to send to the receiver. The receiver observes the signal and contingent on the signal observed chooses an action. Payoffs for sender and receiver are functions of state of the world, action chosen, and (possibly) signal sent. Where payoffs only depend on state and act, interests of sender and receiver may be coincident, partially aligned, or totally opposed. Another layer of complexity is added when signals may be costly to send.

The baseline case is given by signaling when the interests of the sender and the receiver are fully aligned. This scenario was introduced by the philosopher David Lewis in 1969 to analyze conventional meaning (1). We thus call them Lewis signaling games. In the simplest case the action chosen by the receiver is appropriate for exactly one state of the world. If the appropriate action is chosen, then the sender and the receiver get the same payoff of, say, 1; otherwise, they get a payoff of 0.

By varying these payoffs, the sender's and the receiver's incentives may change quite radically (2). Sometimes the sender might have an incentive to not inform the receiver about which of several states is the true one because the action that would be best for the sender does not coincide with the action that would be best for the receiver in a certain state. An extreme form that will be discussed briefly below is arrived at when this is always the case so that the signaling interaction is essentially zero-sum.

Misaligned interests lead to the question of how reliable or honest signaling is possible in such cases (3–5). In costly signaling games certain cost structures for signals may facilitate

communication. Costly signaling games are studied in economics, starting with the Spence game (6), and in biology (e.g., ref. 4). The Spence game is a model of job market signaling. In a job market, employers would like to hire qualified job candidates, but the level of a job candidate's qualification is not directly observable. Instead, a job candidate can send a signal about her qualification to the employer. However, if signals are costless, then job candidates will choose to signal that they are highly qualified regardless of their qualification. If a signal is more costly for low-quality candidates, such that only job candidates of high quality can afford to send it, signaling can be honest in equilibrium. In this case employers get reliable information about a job candidate's quality.

The models of costly signaling in theoretical biology have a related structure. They model situations as diverse as predator–prey signaling, sexual signaling, or parent–offspring interactions (7). A simple example of the latter kind of situation is the so-called Sir Philip Sidney game, which was introduced by John Maynard Smith to capture the basic structure of costly signaling interactions (5). In the Sir Philip Sidney game there are two players, a child (sender) and a parent (receiver). The sender can be in one of two states, needy or not needy, and would like to be fed in either state. The receiver can choose between feeding the sender or abstaining from doing so. She would like to feed the sender provided that the sender is needy. Otherwise, the receiver would rather herself eat. This creates a partial conflict of interest between the sender and the receiver. If the sender is needy, then the interactions between the two players is similar to matching states and acts in the Lewis signaling game. However, if the sender is not needy, then the payoffs of sender and receiver diverge.

Now, it is assumed that a needy sender profits more from being fed than a sender that is not needy. In addition, the sender is allowed to send a costly signal. If the cost of the signal is sufficiently high, then there is again the possibility that in equilibrium the sender signals need honestly and the receiver feeds the sender upon receipt of the signal.

In these costly signaling games—the Spence game or the Sir Philip Sidney game—state-dependent signal costs to the sender or fixed signal costs with state-dependent benefits to the receiver realign the interests of sender and receiver.

Much of the analysis in the literature on signaling games has focused on the most mutually beneficial (Pareto optimal) equilibria of the games under consideration. This would lead one to expect perfect information transfer in partnership games, partial information transfer in games of partially aligned interests, and no information transfer in games of totally opposed interests.

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “In the Light of Evolution VIII: Darwinian Thinking in the Social Sciences,” held January 10–11, 2014, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, CA. The complete program and audio files of most presentations are available on the NAS website at www.nasonline.org/ILE-Darwinian-Thinking.

Author contributions: S.H., B.S., P.T., and E.W. designed research, performed research, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. Email: bskyrms@uci.edu.

Perfect information transfer might be restored in problematic cases by the right differential signaling costs.

However, we do not want to simply rely on faith that Pareto optimal equilibria will be reached. What is required is an investigation of an adaptive dynamic that may plausibly be operative. Many dynamic processes deserve consideration. Here we focus on some dynamics of evolution and of reinforcement learning, where sharp results are available.

Replicator Dynamics

The replicator dynamics is the fundamental dynamical model of evolutionary game theory (8). It describes evolutionary change in terms of the difference between a strategy's average payoff and the overall average payoff in the population. If the difference is positive, the strategy's share will increase; if it is negative, it will decrease. This is one way to capture a basic feature of any selection dynamics. Not surprisingly, the replicator dynamics can be derived from various first principles that describe selection more directly (9).

The two most common varieties of replicator dynamics are the one-population and the two-population replicator dynamics. The one-population replicator dynamics can be used for symmetric two-player games (i.e., two-player games where the player roles are indistinguishable). Let s_1, \dots, s_n denote the n pure strategies that are available to each player. Let $\pi_i(s_j, s_k)$ be the payoff that player i receives from choosing s_j when the other player is choosing s_k . Because the game is symmetric, $\pi_1(s_j, s_k) = \pi_2(s_j, s_k)$. This allows us to drop the indices referring to a player's payoff when considering symmetric games. Let $\mathbf{x} = (x_1, \dots, x_n)$ denote a mixed strategy. Then $\pi(s_i, \mathbf{x})$ is the expected payoff that a player gets when choosing s_i against \mathbf{x} in a random pairwise interaction:

$$\pi(s_i, \mathbf{x}) = \sum_j \pi(s_i, s_j)x_j.$$

Furthermore, $\pi(\mathbf{x}, \mathbf{x})$ is the expected payoff from choosing \mathbf{x} against itself:

$$\pi(\mathbf{x}, \mathbf{x}) = \sum_j \pi(s_j, \mathbf{x})x_j.$$

Suppose now that there is a population consisting of n types, one for each strategy. Then a mixed strategy \mathbf{x} describes the relative frequency of strategies in that population. The state space of the population is the $n - 1$ dimensional unit simplex. The population evolves according to the replicator dynamics if its instantaneous change is given by

$$\dot{x}_i = x_i(\pi(s_i, \mathbf{x}) - \pi(\mathbf{x}, \mathbf{x})) \quad \text{for } 1 \leq i \leq n. \quad [1]$$

Here, the payoff $\pi(s_i, \mathbf{x})$ is interpreted as the fitness of an i strategist in the population state \mathbf{x} , and $\pi(\mathbf{x}, \mathbf{x})$ is the average fitness of that population. Because these fitnesses are expected payoffs, the replicator dynamics requires the population to be essentially infinite.

The two-population replicator dynamics can be applied to asymmetric two-player games. Let s_1, \dots, s_n be player one's pure strategies and t_1, \dots, t_m player two's pure strategies. The mixed strategies $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_m)$ can be identified with the states of two populations, one corresponding to player one and the other to player two. The state space for an evolutionary dynamics of the two populations is the product of the $n - 1$ -dimensional unit simplex and the $m - 1$ -dimensional unit simplex. The two population replicator dynamics defined on this state space is given by

$$\dot{x}_i = x_i(\pi_1(s_i, \mathbf{y}) - \pi_1(\mathbf{x}, \mathbf{y})) \quad \text{for } 1 \leq i \leq n \quad \text{and} \quad [2a]$$

$$\dot{y}_j = y_j(\pi_2(t_j, \mathbf{x}) - \pi_2(\mathbf{y}, \mathbf{x})) \quad \text{for } 1 \leq j \leq m. \quad [2b]$$

Here, $\pi_1(s_i, \mathbf{y})$ is the fitness (expected payoff) of choosing strategy s_i against population state (mixed strategy) \mathbf{y} , and $\pi_1(\mathbf{x}, \mathbf{y})$ is

the average fitness in population one, likewise for $\pi_2(t_j, \mathbf{x})$ and $\pi_2(\mathbf{y}, \mathbf{x})$.

Both the two-population and the one-population replicator dynamics are driven by the difference between a strategy's fitness and the average fitness in its population. This captures the mean field effects of natural selection, but it disregards other factors such as mutation or drift. In many games these factors will only play a minor role compared with selection. However, as we shall see, the evolutionary dynamics of signaling games often crucially depends on these other factors. The reason is that the replicator dynamics of signaling games is generally not structurally stable (10). This means that small changes in the underlying dynamics can lead to qualitative changes in the solution trajectories.

This makes it important to study the effect of perturbations of the replicator dynamics. One plausible deterministic perturbation that has been studied is the selection mutation dynamics (11). We shall consider this dynamics in the context of two population models. If mutations within each population are uniform, the selection mutation dynamics is given by

$$\dot{x}_i = x_i(\pi_1(s_i, \mathbf{y}) - \pi_1(\mathbf{x}, \mathbf{y})) + \varepsilon(1 - nx_i) \quad \text{for } 1 \leq i \leq n \quad \text{and} \quad [3a]$$

$$\dot{y}_j = y_j(\pi_2(t_j, \mathbf{x}) - \pi_2(\mathbf{y}, \mathbf{x})) + \delta(1 - my_j) \quad \text{for } 1 \leq j \leq m. \quad [3b]$$

The nonnegative parameters ε and δ are the (uniform) mutation rates in population one and two, respectively. Instantaneously, every strategy in a population is equally likely to mutate into any other strategy at a presumably small rate. As ε and δ go to zero, the two-population selection mutation dynamics approaches the two-population replicator dynamics. If the replicator dynamics is structurally stable, there will be no essential difference between the replicator dynamics and the selection mutation dynamics as long as ε, δ are small. However, the introduction of deterministic mutation terms can significantly alter the replicator dynamics of signaling games.

Lewis Signaling Games. From the point of view of static game theory, the analysis of Lewis signaling games seems to be straightforward. If the number of signals, states, and acts coincides, then signaling systems are the only strict Nash equilibria. It can also be shown that they are the only evolutionarily stable states (12). However, other Nash equilibria, despite being non-strict, are neutrally stable states (a generalization of evolutionary stability) (13). This suggests that an analysis of the evolutionary dynamics will reveal a more fine-grained picture, as indeed it does.

Consider the one-population replicator dynamics first. The Lewis signaling game as given in the preceding section is not a symmetric game. It can, however, be symmetrized by assuming that a player assumes the roles of a sender and a receiver with equal probability and receives the corresponding expected payoffs (14). If there are n signals, states, and acts, the symmetrized signaling game will have n^{2n} strategies. This results in a formidable number of dimensions for the state space of the corresponding dynamics (1) even for relatively small n . A fairly complete analysis of this dynamical system is nevertheless possible because the Lewis signaling game exhibits certain symmetries. The assumption that both players get the same payoff in every outcome makes it a partnership game, a class of games for which it is known that the average payoff $\pi(\mathbf{x}, \mathbf{x})$ is a potential function (8). This implies that every solution trajectory converges to a rest point, which needs to be a Nash equilibrium. The stable rest points are the local maximizers of $\pi(\mathbf{x}, \mathbf{x})$.

It is clear that signaling systems are locally asymptotically stable because they are strict Nash equilibria. The question is whether there are any other locally asymptotically stable rest points. It can be proved that this is essentially not the case for

signaling games with two states, signals, and acts, where both states are equally probable (15): Every open set in the strategy simplex contains an \mathbf{x} such that the replicator dynamics with this initial condition converges to a signaling system. This rather special case does not generalize, however. If the states are not equi-probable in this signaling game, there is an open set of points whose trajectories do not converge to a signaling system. Instead, they converge to states where receivers always choose the act corresponding to the high-probability state (15).

If there are more than two signals, states, and acts, there always is an open set of points whose trajectories do not converge to a signaling system (13, 15). They typically converge to partial pooling equilibria (15, 16). By “partial pooling equilibria” we mean equilibria that share three features: (i) Some, but not all, signals are unequivocally used for states; (ii) some, but not all, acts are unequivocally chosen in response to a signal; and (iii) no signal is unused. The last feature makes it impossible for mutants who use a signaling system strategy to invade by exploiting an unused signal. A set of partial pooling equilibria P is not an attractor because for any neighborhood N of P there exists a solution trajectory that leaves N . Partial pooling equilibria are Liapunov-stable, though. As is shown in ref. 13, partial pooling equilibria coincide with those neutrally stable states that are not also evolutionarily stable states.

A result that holds for all Lewis signaling games concerns the instability of interior rest points. At an interior rest point, every strategy is present, creating a “tower of Babel” situation. It can be shown that any such rest point is linearly unstable. This implies that the unstable manifold of these rest points has a dimension of at least one. Hence, there is no open set of points whose trajectories converge to the set of interior rest points (15).

These results carry over to the case of two populations. The game is again a partnership game. It follows that $\pi_1(\mathbf{x}, \mathbf{y}) = \pi_2(\mathbf{y}, \mathbf{x})$ is a potential function of Eq. 2 and that every trajectory must converge. Signaling systems are asymptotically stable by virtue of being strict Nash equilibria. For signaling games with two signals, states, and acts where the states are equiprobable almost all trajectories converge to a signaling system. This result fails to hold when the states are not equiprobable (17). Furthermore, partial pooling equilibria are stable for Eq. 2 as in the one-population case.

The selection mutation dynamics (Eq. 3) of Lewis signaling games was studied computationally in ref. 16 and analytically in ref. 17. The main reason for studying selection mutation dynamics (Eq. 3) is that partial pooling equilibria are not isolated. They constitute linear manifolds of rest points. It is well known that this situation is not robust under perturbations. Introducing mutation terms as in Eq. 3 will destroy the linear manifolds of rest points and create a topologically different dynamics in that region of state space.

For other games this topic was studied in ref. 18. In ref. 17 it is shown that the function

$$\pi_1(\mathbf{x}, \mathbf{y}) + \varepsilon \sum_i \log x_i + \delta \sum_j \log y_j$$

is a potential function for the selection mutation dynamics of the Lewis signaling game. Hence, all trajectories converge. There are two additional general results. The first says that rest points of the perturbed dynamics (3) must be close to Nash equilibria of the signaling game. There are thus no “anomalous” rest points that are far away from any Nash equilibrium. Second, there is a unique rest point close to any signaling system that is asymptotically stable. Signaling systems remain evolutionarily significant.

There are no further general results. In ref. 17 the case of two states, two signals, and two acts is explored in more detail. If the states are equally probable, the overall conclusions are similar to

the results of the replicator dynamics. There are three rest points: Two are close to signaling systems (perturbed signaling systems) and the third is in the interior of the state space. The latter is linearly unstable whereas the perturbed signaling systems are linearly stable. So, although there are only finitely many rest points for the selection mutation dynamics (as opposed to the replicator dynamics) of this game, the basic conclusion is that in every open set in the state space there is an \mathbf{x} such that the selection mutation dynamics of Eq. 3 with this initial condition converges to a signaling system.

Things are more nuanced if the two states are not equiprobable, as shown in Fig. 1. In this case the dynamic behavior depends on the ratio of the mutation parameters δ/ε . If δ/ε is above a certain threshold, which includes the important case $\delta = \varepsilon$, then almost all trajectories converge to one of the signaling systems. If δ/ε is below the threshold, then there exists an asymptotically stable rest point where nearly all members of the receiver population choose the act that corresponds to the more probable signal. Thus, outcomes with basically no communication are robust under the introduction of mutation into the replicator Eqs. 2a and 2b.

It is very difficult to analyze the selection mutation dynamics (Eq. 3) of Lewis signaling games for the case of three states, acts,




Fig. 1. A bifurcation for a signaling game with two signals under the selection mutation dynamics. The game shows a fragment of the full state space with two sender strategies and three receiver strategies. Dots indicate rest points. Black dots indicate asymptotically stable rest points and white dots indicate unstable rest points. In *A* the ratio δ/ε is sufficiently large (including the case $\delta = \varepsilon$) so that orbits converge to one of the two signaling systems from almost all initial conditions. However, in *B*, one rest point close to the upper boundary of the state space is asymptotically stable; hence, some orbits converge to this rest point and not to the signaling systems.

and signals. The main reason is the rapidly increasing dimensionality of the state space. In a two-population model, there are 27 types of individuals in each population, resulting in the product of two simplexes that has 52 dimensions. By exploiting the underlying symmetry of the Lewis signaling game and by introducing certain simplifications, it is nonetheless possible to prove some results. The main results concern the existence of rest points close to Nash equilibria other than the signaling systems. Most notably, at least one rest point exists close to each component of partially pooling Nash equilibria. It can be shown that for all mutation parameters this rest point is linearly unstable.

Costly Signaling Games. The Spence game and the Sir Philip Sidney game were studied with the help of the replicator dynamics. Additionally, simplified versions of the Sir Philip Sidney game and related games were analyzed recently (19). For all these games, the replicator dynamics leads to very similar results, which can differ quite markedly from those obtained in finite population models (discussed below in *Finite Population Dynamics*).

The dynamics of Spence's model for job market signaling is explored in ref. 20 for a dynamic process of belief and strategy revision that is different from the dynamical models considered here, although the analysis leads to somewhat similar results. In ref. 21 the two-population replicator dynamics (2) of Spence's game is investigated. (Strictly speaking, it is a discretized variant of Spence's original game that has a continuum of strategies.) For all parameter settings there exists a pooling equilibrium where senders do not send a signal and receivers ignore the signal. If signaling cost is sufficiently high, then a separating equilibrium exists; the separating equilibrium can be viewed as the analog to a Lewis signaling system where signals are used for revealing information about the sender. If signaling cost is not high enough, a so-called hybrid equilibrium exists where senders mix between using signals reliably and unreliably, and where receivers sometimes respond to a signal and sometimes ignore it. As is pointed out in ref. 21, the hybrid equilibrium has been almost completely ignored in the large literature on costly signaling games, although it provides an interesting low-cost alternative to the standard separating equilibrium (19).

It can be shown that the pooling equilibrium in this version of Spence's game is always asymptotically stable for the replicator dynamics (Eq. 2). The same is true for the separating equilibrium, provided that it exists. The dynamic behavior of the hybrid equilibrium is particularly interesting. It is Liapunov-stable. More precisely, it is a spiraling center. It lies on a plane on the boundary of state space. With respect to the plane, the eigenvalues of the Jacobian matrix evaluated at the hybrid equilibrium have zero real part, which makes it into a center when we restrict the dynamics to the plane. In particular, on the plane the trajectories off the rest point are periodic cycles. With respect to the interior of the state space, the eigenvalues of the Jacobian evaluated at the hybrid equilibrium are negative. Hence, from the interior trajectories approach the hybrid equilibrium in a spiraling movement.

The same is true for other costly signaling models such as the well-known Sir Philip Sidney game (19, 22). In particular, hybrid equilibria exist and are dynamically stable for Eq. 2 in the same way as for the job market signaling game. This suggests that hybrid equilibria can be evolutionarily significant outcomes. Both refs. 21 and 22 present numerical simulations that reinforce this conclusion in terms of the relative sizes of basins of attraction. According to these simulations, the basin of attraction of hybrid equilibria is quite significant, whereas the basin of attraction for separating equilibria is surprisingly small.

A question that has only recently been investigated is whether the hybrid equilibrium continues to be dynamically stable under perturbations of the dynamics (Eq. 2). The question of structural

stability is important here as well because a spiraling center is not robust under perturbations. Small perturbations of the dynamics will push the eigenvalues with zero real part to having positive or negative real part. So it seems possible that by introducing mutation as in Eq. 3 the hybrid equilibrium might cease to be dynamically stable.

This is not so for sufficiently small mutation parameters. First, it follows from the implicit function theorem that there exists a unique rest point of Eq. 3 close to the hybrid equilibrium (19). Second, it can be proved that the real parts of all eigenvalues of the Jacobian matrix of Eq. 3 evaluated at this perturbed rest point are negative. Thus, the rest point corresponding to the hybrid equilibrium is not a spiraling center anymore but is asymptotically stable instead. This actually reinforces the qualitative point made above, namely, that hybrid equilibria should be considered as theoretically significant evolutionary outcomes.

Another question is whether the hybrid equilibrium is also empirically significant. One of the most robust findings in costly signaling experiments and field studies is that observed costs are generally too low to validate the hypothesis that high costs support signaling system equilibria (7). The hybrid equilibrium is in certain ways an attractive alternative to that hypothesis. It allows for partial information transfer at low costs. For this reason it was suggested that hybrid equilibria could be detected and distinguished from separating equilibria in real-world signaling interactions (19).

Opposed Interests. The possibility of signaling when interests conflict has also been studied in a rather extreme setting where the interests of senders and receivers are completely opposed (23). There is no signaling equilibrium possible in this case. However, in the two-population replicator dynamics (Eq. 2) there is information transfer off equilibrium to varying degrees because there exists a strange attractor in the interior of the state space. By information transfer we understand that signals have information in the sense of Kullback–Leibler entropy: Conditioning on the signal changes the probability of states so that on average information is gained (for details on applying information theory to signaling games see ref. 24). This result reinforces the diagnosis that a dynamical analysis is unavoidable if one wants to understand the evolutionary significance of signaling phenomena.

Finite Population Dynamics

Consider a small finite population of fixed size. Each step of the dynamics works as follows. First, everyone plays the base game with everyone else in a round-robin fashion. Each individual's fitness is given by a combination of her background fitness and her average payoff from the round robin tournament. Following ref. 25 we will take the fitness to be

$$1 - w + w \times u(s_i, \mathbf{x}),$$

where $w \in [0, 1]$ is a parameter that measures the intensity of selection and $u(s_i, \mathbf{x})$ is the expected payoff of strategy i against the population \mathbf{x} , just as in the case of the replicator dynamics.

The background fitness w is the same for everyone. If $w = 0$ the game's payoffs do not matter to an individual's fitness. If $w = 1$ the game's payoffs are all that matter. Next, one individual is selected at random to die (or to leave the group) and a new individual is born (or enters the group). The new individual adopts the strategy of an individual chosen from the population with probability proportional to its fitness. Successful strategies are more likely to be adopted and will therefore spread through the population. This dynamics, known as the frequency-dependent Moran process, is a Markov chain with the state being the number of individuals playing each strategy. Owing to the

absence of mutation or experimentation all monomorphic population compositions are absorbing states of this process.

Lewis Signaling Games. Pawlowitsch (26) studies symmetrized Lewis signaling games under this dynamics. Let N be the number of individuals all playing strategy s_j . Imagine that one spontaneously mutates to s_i . The probability that strategy s_i goes on to take over the entire population is given by the fixation probability

$$\rho_{s_j s_i} = \frac{1}{1 + \sum_{k=1}^{N-1} \prod_{l=1}^k \frac{g_l(s_i, s_j)}{f_l(s_i, s_j)}}$$

where $f_l(s_i, s_j)$ is the fitness of an s_i agent in a population of l individuals playing strategy s_i and $N-l$ individuals playing s_j and $g_l(s_i, s_j)$ is the fitness of an s_j individual in that same population. If the mutation is neutral then its probability of fixation is $1/N$. Pawlowitsch (26) uses this neutral threshold to assess the evolutionary stability of monomorphic population compositions. In Lewis signaling games under the Moran process with weak selection ($wN \ll N$) Pareto optimal strategies—that is, perfectly informative signaling strategies—are the only strategies for which there is no mutant type that has a fixation probability greater than this neutral threshold. For this reason Pawlowitsch argues that finite populations will choose an optimal language in Lewis signaling games. This highlights an important difference between the behaviors of infinite and finite populations in these games.

Costly Signaling Games. However, what about signaling games where interests conflict? To address this question we introduce mutation. Suppose that with small probability, ϵ , the new individual mutates or decides to experiment and chooses a strategy at random from all of the possible strategies—including those not represented in the population—with equal probability. The presence of this mutation makes the resulting Markov process ergodic. Fudenberg and Imhof (27) show that it is possible to use a so-called embedded Markov chain to calculate the proportion of time that the population spends in each state in the limiting case as ϵ goes to zero. The states in this embedded Markov chain are the monomorphic population compositions, and the transition probability from the monomorphic population in which all individuals play s_j to the monomorphic population in which all individuals play s_i is given by the probability that a type s_i mutant arises (ϵ) multiplied by the probability that this mutant fixes in the population ($\rho_{s_j s_i}$). The stationary distribution of this embedded chain gives the proportion of time that the population spends in each monomorphic state in the full Moran process as ϵ goes to zero. Intuitively, this is because when ϵ is very small the system will spend almost all of the time in a monomorphic state waiting for the next mutation event, and after an event the Moran process will return the population to a monomorphic state before the next mutant arises.

Consider the two-state, two-signal, two-act signaling game with payoffs given in Fig. 2. The receiver prefers the act high in the high state and low in the low state, whereas the sender always prefers the act high. This game is structurally similar to the Sir Philip Sidney game, but here we will assume that both signals are costless. In this game there is no Nash equilibrium in which the signals discriminate between the states. The only Nash equilibria are pooling, where the sender sends signals with probabilities independent of the states and the receiver acts low.

The Moran process is a one-population setting, so we will consider the symmetrized version of this game (14). We then let our round-robin phase match each pair both as (sender, receiver) and (receiver, sender). Suppose that selection is strong (i.e., $w=1$), and that the probability of state high is 0.4. Then in the

small mutation limit the process spends 57% of its time in states with perfect signaling and 19% of its time in Nash equilibria (28). In other words, this small population spends most of its time communicating perfectly even though such information transfer is not a Nash equilibrium of the underlying signaling game. This phenomenon is robust over a wide range of selection intensities and state probabilities, but as the population size is increased the proportion of time spent signaling diminishes. This is to be expected because as the population size tends to infinity the behavior of the Moran process tends toward the behavior of the payoff-adjusted replicator dynamic (29), which does not lead to information transfer in this game.

The importance of non-Nash play in the rare mutation limit is also evident in the case of cost-free preplay signaling. Two players play a base game, but before they do so each sends the other a cost-free signal from some set of available signals, with no preexisting meaning. The small population, rare mutation limit for a related dynamics is investigated in ref. 30 for the cases where the base game is (i) Stag Hunt and (ii) Prisoner's Dilemma. Like the Moran process, this related dynamics is composed of two steps. First, all individuals play the base game with each other in a round-robin fashion to establish fitness. Second, an individual is randomly selected to update her strategy by imitation. This agent randomly selects another individual and imitates that other individual with a probability that increases with an increase in the fitness difference between the two agents. In particular, this probability is given by the Fermi distribution from statistical physics so that the probability that an individual using strategy s_i will imitate an individual using strategy s_j is given by the function

$$\left[1 + e^{-\beta[\pi(s_i, x) - \pi(s_j, x)]}\right]^{-1},$$

where β should be interpreted as noise in the imitation process. For high values of β a small payoff difference translates into a high probability of imitation, whereas when β tends to zero selection is weak and the process is dominated by random drift (29).

In the case where the base game is the Prisoner's Dilemma, the only Nash equilibrium is noncooperation. However, if the population is small and the set of signals is large, the population may spend most of its time cooperating. If the base game is the Stag Hunt, where there are both cooperative and noncooperative equilibria, preplay signals enlarge the amount of time spent cooperating, with the more signals the better.

Why is it that dynamics in finite populations with rare mutations can favor informative signaling even when such behavior is not a Nash equilibrium? Consider again the game in Fig. 2. In a pooling equilibrium the sender's expected payoff is 0.48 and the receiver's is 0.6. Jointly separating, however, gets the sender an expected payoff of 0.88 and the receiver an expected payoff of 1. Signaling Pareto dominates pooling, and this means that a small population is likely to transition from a monomorphic pooling state to a monomorphic signaling state. Of course, if the receiver discriminates, then the sender can gain by always sending whichever signal induces act high. Such behavior will net the sender an expected payoff of 1. However, note that there is a smaller difference in payoff for the sender between this profile

	Act High	Act Low
State High	1, 1	0, 0
State Low	1, 0	.8, 1

Fig. 2. The payoff structure underlying a signaling game. If the state is high, then the sender's and the receiver's interest coincide. If the state is low, then the sender prefers the receiver to choose high, whereas the receiver would want to choose low.

and the separating profile than there is between the separating profile and the pooling equilibrium. Consequently, the probability of transitioning away from a monomorphic separating population to the monomorphic population in which the senders always induce the receivers to perform act high is less than the probability of transitioning from a monomorphic pooling population to a monomorphic separating population. For this reason, in the long run the population will spend more time signaling informatively than it will spend pooling. A similar story explains why finite population dynamics can favor informative signaling and cooperation in prisoner's dilemma games with cost-free preplay signaling (30).

Reinforcement Learning

In models from evolutionary game theory—such as the replicator dynamics—“learning” occurs globally because the size of populations with more fitness increases faster. However, in the reinforcement learning model it is the individuals' behavior that evolves iteratively: The players tend to put more weight on strategies that have enjoyed past success, as measured by the cumulative payoffs they have achieved. This linear response rule corresponds to Herrnstein's “matching law” (31).

Reinforcement learning is one of a variety of models of strategic learning in games, where players adapt their strategies with the aim to eventually maximize their payoffs: No-regret learning, fictitious play and its variants, and hypothesis testing, are other examples of such procedures analyzed in game theory (32).

However, reinforcement learning is a particularly attractive and simple model of players with bounded rationality. The amount of information used in the procedure is small: Players need only observe their realized payoffs and may not even be aware that they are playing a game with or against others. It accumulates inertia, because the relative increase in payoff decreases with time. On the one hand this might be expected from a learning procedure, although on the other hand it could be exploited by other players in certain games.

In Roth and Erev (33) and Erev and Roth (34) reinforcement learning is proposed and tested experimentally as a realistic model for the behavior of agents in games—see Harley (35) for a similar study in a biological context.

Formally, in a game played repeatedly by N players, each one having M strategies, each individual i is assumed at a time step n to have a propensity q_{jn}^i for each strategy j and plays the strategy with probability proportional to its propensity,

$$\frac{q_{jn}^i}{\sum_{k=1}^m q_{kn}^i}. \quad [4]$$

Each individual i is endowed with an initial vector of positive weights (q_{j0}^i) at time 0. At each iteration of the learning process, the strategy j taken by any player i results in a nonnegative payoff, $U_j(i)$, and the weights are updated by adding that payoff to the weight of the act taken,

$$q_{j,i+1}^i = q_{j,i}^i + U_j(i), \quad [5]$$

with the weights of strategies not taken remaining the same.

This process can be exemplified by an urn model. Each individual starts with an urn containing some balls of different colors, one for each potential strategy. Drawing a ball from the urn (and then replacing it) determines the choice of strategy. After receiving a payoff, the number of balls of the same color equal to the payoff achieved are added to the urn.

As balls pile up in the urn, jumps in probabilities become smaller and smaller in such a way that the stochastic process approximates a deterministic mean field dynamics, which is known as the adjusted or Maynard Smith version of the replicator

dynamics (36, 37). However, classical stochastic approximation theory (38–41) does not allow one to deduce much in general, because this ordinary differential equation takes place in an unbounded domain.

Beggs (36) shows that if all players apply this rule then iteratively strictly dominated strategies are eliminated, and the long-run average payoff of a player who applies it cannot be forced permanently below its minmax payoff. He also studies two-person constant-sum games, where precise results can be obtained. Hopkins and Posch (37) show convergence with probability zero toward unstable fixed points of the Maynard Smith replicator dynamics, even if they are on the boundary, which solves earlier questions raised in particular in ref. 42. This second result is, however, not relevant in signaling games, where the unstable fixed points are not isolated and consist of manifolds of finite dimension (43, 44).

Consider the simplest Lewis signaling game. The reinforcement learning model proposed in ref. 43 considers, in Eqs. 4 and 5, each state possibly transmitted by the sender as a player whose strategies are the signals and, similarly, each signal as a player whose strategies are the (guessed) states. Now a state i that “plays” signal j gets a payoff of 1 if, conversely, j “plays” i and Nature chooses i . In practice, at each time step, only the state chosen by Nature will play along with its chosen signal, so that this reinforcement procedure can be simply explained from a sender–receiver perspective.

Let us first consider the case of two states 1 and 2: Nature flips a fair coin and chooses one of them. The sender has an urn for state 1 and a different urn for state 2. Each has balls for signal A and signal B. The sender draws from the urn corresponding to the state and sends the indicated signal. The receiver has an urn for each signal, each containing balls for state 1 and for state 2. The receiver draws from the urn corresponding to the signal and guesses the indicated state. If correct, both sender and receiver are reinforced and each one adds a duplicate ball to the urn just exercised. If incorrect, there is no reinforcement and the urns are unchanged for the next iteration of the process.

There are now four interacting urns contributing to this reinforcement process. However, the dimensionality of the process can be reduced because of the symmetry resulting from the strong common interest assumption. Because the receiver is reinforced if and only if the sender is, the numbers of balls in the receiver's urns are determined by the numbers of balls in the sender's urns. Consider the four numbers of balls in the sender's urns: $1A$, $1B$, $2A$, $2B$. Normalizing these, $1A/(1A + 1B + 2A + 2B)$, etc., gives four quantities that live on a tetrahedron. The mean field dynamics may be written in terms of these. There is a Lyapunov function that rules out cycles. The stochastic process must then converge to one of the zeros of the mean field dynamics. These consist of the two signaling systems and a surface composed of




Fig. 3. Synonyms and bottlenecks. One state is sent to two different signals, which we can call synonyms. Two states are sent to only one signal, which we call an information bottleneck.

pooling equilibria. It is possible to show that the probability of converging to a pooling equilibrium is zero. Thus, reinforcement learning converges to a signaling system with probability one (43).

This result raises several questions: Does the same result hold for N states, N signals, and N acts? What happens if there are too few signals to identify all of the states or if there is an excess of signals—that is, N states, M signals, N acts? Nature now rolls a fair die to choose the state, the sender has N urns with balls of M colors, and the receiver has M urns with N colors.

This reinforcement process is analyzed in ref. 44. Common interest allows a reduction of dimensionality, as before. The sender is reinforced for sending signal m in state s just in case the receiver is reinforced for guessing state s when presented with signal m . Again, for a state i and a signal j , consider the number ij of balls of color j in the sender's urn i . As in the 2×2 case, the dynamics of the normalized vector of sender-receiver connections $ij/\sum ij$ is studied, as stochastic approximation of a non-continuous dynamics on the simplex.

The expected payoff can be shown to be a Lyapunov function for the mean field dynamics; convergence to the set of rest points is deduced, with a technically involved argument. This is required because of the discontinuities of the dynamics, and because not all Nash equilibria are rest points for this dynamics, contrary to what we have in the standard replicator dynamics.

The stability properties of the equilibria of this mean field dynamics can be linked to static equilibrium properties of the game, which are described by Pawlowitsch (13): The zeros of the gradient of the payoff, the linearly stable equilibria, as well as the asymptotically stable equilibria of the mean field dynamics, correspond, respectively, to the Nash equilibria, neutrally stable strategies, and evolutionarily stable strategies of the signaling game.

Finally, the following result can be stated in terms of a bipartite graph between states and signals, such that there is an edge between a state and a signal if and only if that signal is chosen infinitely often in that state. It is shown that any such graph with the following property, P, has a positive probability of being the limiting result of reinforcement learning:

P: (i) Every connected component contains a single state or a single signal and (ii) each vertex has an edge.

In the event where property P holds, if there is an edge between a state and a signal, then the limiting probability of sending that signal in that state is positive.

When $M = N$, property P is exemplified by signaling systems, where each state is mapped with probability one to a unique signal. However, even in this case it is also exemplified by configurations that contain both synonyms and information bottlenecks as in Fig. 3. Evolution of optimal signaling has positive probability, but so does the evolution of this kind of suboptimal equilibrium. The case of $M = N = 2$ is very special. This corresponds closely to the replicator dynamics of signaling games, where partially pooling equilibria (which contain synonyms and bottlenecks) can be reached by the replicator dynamics.

Discussion

The results on the replicator dynamics suggest that for large populations the emergence of signaling systems in Lewis signaling games with perfect common interest between sender and receiver is guaranteed only under special circumstances. The dynamics also converges to states with imperfect information transfer. Introducing mutation can have the effect of making the emergence of perfect signaling more likely, although this statement should be taken with a grain of salt because the precise outcomes may depend on the mutation rates.

When interests are diametrically opposed in signaling games, there is no information transmission in equilibrium. However, the equilibrium may never be reached. Instead senders and receivers may engage in a mad “Red Queen” chase, generating cycles or chaotic dynamics. In well-known costly signaling games, where interests are mixed, this Red Queen chase is a real possibility. Along the trajectories describing such a chase there are periods with significant information transfer from senders to receivers. Those interactions are undermined because of the underlying conflicts of interest, resulting in periods of low information transfer, from which a new period of higher information transfer can start.

Unlike large populations, a small population may spend most of its time efficiently signaling, even when the only Nash equilibrium does not support any information transfer. A small population engaged in “cheap talk” costless preplay signaling may spend most of its time cooperating even when the only Nash equilibrium does not support cooperation.

A similar difference between small and large populations may be at work in costly signaling games. In the small population, small mutation limit costless signaling is possible even in games with conflict of interest. In large populations this is not true, but there are alternatives to the costly signaling equilibria where signaling cost can be low while in equilibrium there is partial information transfer.

We encounter a similarly nuanced picture for models of individual learning. Herrnstein–Roth–Erev reinforcement learning leads to perfect signaling with probability one in Lewis signaling games only in the special case of two equiprobable states, two signals, and two acts. In more general Lewis signaling games, the situation is much more complicated. To our knowledge, nothing is known about reinforcement learning for games with conflicts of interest and costly signaling games. This would be a fruitful area for future research.

We conclude that the explanatory significance of signaling equilibria depends on the underlying dynamics. Signaling games have multiple Nash equilibria. One might hope that natural dynamics always selects a Pareto optimal Nash equilibrium, but this is not always so. On a closer examination of dynamics, in some cases, Nash equilibrium recedes in importance and other phenomena are to be expected.

ACKNOWLEDGMENTS. This work was supported by National Science Foundation Grant EF 1038456 (to S.H.).

- Lewis DK (1969) *Convention: A Philosophical Study* (Harvard Univ Press, Cambridge, MA).
- Crawford V, Sobel J (1982) Strategic information transmission. *Econometrica* 50(6):1431–1451.
- Zahavi A (1975) Mate selection—a selection for a handicap. *J Theor Biol* 53(1):205–214.
- Grafen A (1990) Biological signals as handicaps. *J Theor Biol* 144(4):517–546.
- Maynard Smith J (1991) Honest signalling: The Philip Sidney game. *Anim Behav* 42:1034–1035.
- Spence M (1973) Job market signaling. *Q J Econ* 87(3):355–374.
- Searcy WA, Nowicki S (2005) *The Evolution of Animal Communication* (Princeton Univ Press, Princeton).
- Hofbauer J, Sigmund K (1998) *Evolutionary Games and Population Dynamics* (Cambridge Univ Press, Cambridge, UK).
- Weibull J (1995) *Evolutionary Game Theory* (MIT Press, Cambridge, MA).
- Guckenheimer J, Holmes P (1983) *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields* (Springer, New York).
- Hofbauer J (1985) The selection mutation equation. *J Math Biol* 23(1):41–53.
- Wärneryd K (1993) Cheap talk, coordination and evolutionary stability. *Games Econ Behav* 5(4):532–546.
- Pawlowitsch C (2008) Why evolution does not always lead to an optimal signaling system. *Games Econ Behav* 63(1):203–226.
- Cressman R (2003) *Evolutionary Dynamics and Extensive Form Games* (MIT Press, Cambridge, MA).
- Huttenberger SM (2007) Evolution and the explanation of meaning. *Philos Sci* 74(1):1–27.
- Huttenberger SM, Skyrms B, Smead R, Zollman KJS (2010) Evolutionary dynamics of Lewis signaling games: Signaling systems vs. partial pooling. *Synthese* 172(1):177–191.
- Hofbauer J, Huttenberger SM (2008) Feasibility of communication in binary signaling games. *J Theor Biol* 254(4):843–849.

18. Binmore K, Samuelson L (1999) Evolutionary drift and equilibrium selection. *Rev Econ Stud* 66(2):363–394.
19. Zollman KJS, Bergstrom CT, Huttegger SM (2013) Between cheap and costly signals: the evolution of partially honest communication. *Proc Biol Sci* 280(1750):20121878.
20. Nöldeke G, Samuelson L (1997) A dynamic model of equilibrium selection in signaling games. *J Econ Theory* 73(1):118–156.
21. Wagner EO (2013) The dynamics of costly signaling. *Games* 4:161–183.
22. Huttegger SM, Zollman KJS (2010) Dynamic stability and basins of attraction in the Sir Philip Sidney game. *Proc Biol Sci* 277(1689):1915–1922.
23. Wagner EO (2012) Deterministic chaos and the evolution of meaning. *Br J Philos Sci* 63(3):547–575.
24. Skyrms B (2010) *Signals: Evolution, Learning, and Information* (Oxford Univ Press, Oxford).
25. Nowak MA, Sasaki A, Taylor C, Fudenberg D (2004) Emergence of cooperation and evolutionary stability in finite populations. *Nature* 428(6983):646–650.
26. Pawlowitsch C (2007) Finite populations choose an optimal language. *J Theor Biol* 249(3):606–616.
27. Fudenberg D, Imhof L (2006) Imitation processes with small mutations. *J Econ Theory* 131(1):251–262.
28. Wagner EO (2014) Semantic meaning in finite populations with conflicting interests. *Br J Philos Sci*, in press.
29. Traulsen A, Claussen JC, Hauert C (2005) Coevolutionary dynamics: From finite to infinite populations. *Phys Rev Lett* 95(23):238701.
30. Santos FC, Pacheco JM, Skyrms B (2011) Co-evolution of pre-play signaling and cooperation. *J Theor Biol* 274(1):30–35.
31. Herrnstein RJ (1970) On the law of effect. *J Exp Anal Behav* 13(2):243–266.
32. Young P (2005) *Strategic Learning and Its Limits* (Oxford Univ Press, Oxford).
33. Roth AE, Erev I (1995) Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games Econ Behav* 8:164–212.
34. Erev I, Roth AE (1998) Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *Am Econ Rev* 88(4):848–881.
35. Harley CB (1981) Learning the evolutionarily stable strategy. *J Theor Biol* 89(4):611–633.
36. Beggs AW (2005) On the convergence of reinforcement learning. *J Econ Theory* 122:1–36.
37. Hopkins E, Posch M (2005) Attainability of boundary points under reinforcement learning. *Games Econ Behav* 53(1):110–125.
38. Robbins H, Monro S (1951) A stochastic approximation method. *Ann Math Stat* 22:400–407.
39. Benaïm M, Hirsch MW (1996) Asymptotic pseudotrajectories and chain recurrent flows, with applications. *J Dyn Differ Equ* 8(1):141–176.
40. Benaïm M (1999) Dynamics of stochastic approximation algorithms. *Séminaire Probabilités XXXIII*:1–68.
41. Pemantle R (2007) A survey of random processes with reinforcement. *Prob Surveys* 4:1–79.
42. Laslier JF, Topol R, Walliser B (2001) A behavioral learning process in games. *Games Econ Behav* 37(2):340–366.
43. Argiento R, Pemantle R, Skyrms B, Volkov S (2009) Learning to signal: Analysis of a micro-level reinforcement model. *Stoch Proc Appl* 119(2):373–390.
44. Hu Y, Skyrms B, Tarrès P (2011) Reinforcement learning in a signaling game. arXiv:1103.5818v1.