Deception and the Evolution of Plasticity

Rory Smead*

Recent models using simple signaling games provide a theoretical setting for investigating the evolutionary connection between signaling and behavioral plasticity. These models have shown that plasticity is typically eliminated in common-interest signaling games. In many real cases of signaling, however, interests do not align. Here, I present a model of the evolution of plasticity in signaling games and consider games of common, opposed, and partially aligned interests. I find that the setting of partial common interest is most conducive to the evolution of plasticity. Plastic individuals succeed by learning when to deceive others and when to trust others' signals.

1. Introduction. Cognition is generally costly to maintain, can lead to errors, and is not always necessary. This presents an evolutionary puzzle: Why did cognition evolve? One intuitive answer is that cognition enables adaptation to a variable or complex environment (Godfrey-Smith 1996). Other organisms are an important source of environmental complexity, and many of our cognitive processes are directed toward one another, which suggests that social interaction may be particularly important for the evolution of cognition. This is the general notion behind the social intelligence hypothesis (Humphrey 1976): cognition evolved to help us navigate our social world. We might go further, however, and investigate what types of social interaction are important for the evolution of cognition. Communication is an obvious candidate. Cognition seems to enable better use of signals and information—both sending and receiving. Here, I address this possibility by exploring some evolutionary models of the simplest instances of language and cognition: signaling games and behavioral plasticity.

Plasticity is nicely exemplified by rudimentary behavioral learning. In recent years, evolutionary game theory has been used to examine the evolution of learning in social settings (see, e.g., Harley 1981; Maynard Smith 1982;

*To contact the author, please write to: 371 Holmes Hall, 360 Huntington Ave., Northeastern University, Boston, MA 02115; e-mail: r.smead@neu.edu.

Philosophy of Science, 81 (December 2014) pp. 852–865. 0031-8248/2014/8105-0012$10.00
Copyright 2014 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

Dubois, Morand-Ferron, and Giraldeau 2010; Katsnelson et al. 2012; Smead 2012). In the context of signaling, game theoretic models have shown that plasticity does not generally persist in common-interest signaling games: plasticity can help a population reach efficient signaling systems, but it is ultimately eliminated by natural selection in favor of nonplastic behavior (Zollman and Smead 2010). While common-interest signaling may occur frequently in nature, many signaling games are not common interest. When interests of signaling parties diverge, deception becomes possible and may be a probable evolutionary outcome (Skyrms 2010). This article presents a set of models for the evolution of a simple reinforcement learning rule in signaling games with varying degrees of common interest. As we will see, the best case for the evolution of plasticity involves partial common interest, where plastic types can gain an advantage by learning when to deceive and when to trust others.

In what follows we will make use of two-population evolutionary models, which are most naturally interpreted as representing interactions between species. However, I will argue that these models need not be interpreted in this way and may also represent interactions among conspecifics or even sets of interacting cells within an organism. Furthermore, although common-interest signaling is not generally conducive to the evolution of plasticity, the two-population model reveals possibilities not seen in one-population settings.

2. The Model. The Lewis sender-receiver game has become a standard model for studying the evolution of signaling (Lewis 1969; Skyrms 1996). In this game, there are two players: the sender and the receiver. The sender observes some state (s) of the world and sends a message (m) to the receiver. The receiver observes the message, but not the state, and must choose an action (a). The payoffs to the sender and the receiver depend on the state-act combination.

For tractability, I focus only on two-state, two-message, two-act signaling games with equiprobable states. A sender's strategy tells her which message to send in which state. A receiver's strategy tells her which act to do in response to each possible message. The possible strategies are listed below.

S1. Send m_1 in s_1 and m_2 in s_2.
S2. Send m_2 in s_1 and m_1 in s_2.
S3. Always send m_1.
S4. Always send m_2.
R1. Do a_1 when m_1 and a_2 in m_2.
R2. Do a_2 when m_1 and a_1 in m_2.
R3. Always do a_1.
R4. Always do a_2.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

Strategies S1, S2, R1, and R2 are referred to as "separating" strategies; the others are "pooling" strategies. Let $\pi_s(i, j)$ be the sender's payoff when strategy i plays strategy j, and $\pi_r(i, j)$ is the receiver's payoff.

2.1. Evolution in Two Populations. I use the two-population replicator dynamic to model evolution (Taylor and Jonker 1978; Hofbauer and Sigmund 1998). Individuals in a sender population are randomly paired with individuals in a receiver population to play a signaling game. Populations are continuous, and states are represented as vectors $X = (x_1, \ldots, x_N)$, where x_i tracks the frequency of type i in the population and there are N types. Types in the population correspond to strategies in the signaling game. Let X represent the sender population and Y the receiver population. The fitness for each type can be calculated as $F(i, Y) = \sum_j y_j \pi_s(i, j)$ and $F(i, X) = \sum_j x_j \pi_r(j, i)$ for senders and receivers, respectively. Evolution in each population occurs according to a set of differential equations: $dx_i/dt = x_i(F(i, Y) - \sum_j x_j F(j, Y))$ for the sender population and similarly for the receiver population.

It will be convenient at several points to employ computer simulations of a discrete-time version of the two-population replicator dynamic. The discrete-time dynamic uses the difference equation $x_i^{t+1} = x_i^t(F(i, Y)/\sum_j x_j^t F(j, Y))$ and shares many properties with the continuous time version. Where simulation results are reported, results are aggregated from 1,000 independent simulations with an initial population chosen at random from the population space with a uniform distribution. All simulations were run for 10^6 generations unless noted otherwise.

Several of the equilibria identified below are only weakly stable (Lyapunov stable) in the replicator dynamic and, as a result, may be partially influenced by drift and the stochastic nature of selection. For this reason, it will be helpful to also include mutation in the model. A standard approach is to assume that mutation occurs uniformly across strategies with probability μ. In the discrete-time replicator dynamic this amounts to $x_i^{t+1} = (1 - \mu)\, x_i^t(F(i, Y)/\sum_j x_j^t F(j, Y)) + \mu/N$.

2.2. Including Plasticity. An additional plastic type L will also be included in the sender and the receiver populations. It is assumed that interactions between pairs of individuals are repeated and that L is able to recognize individuals, giving this type an opportunity to adapt within each interaction.[1] More specifically, plastic individuals will be Roth-Erev rein-

1. The procedure described here follows that of Zollman and Smead (2010). The assumption that learners are able to recognize individuals (or at least their types), although strong, enables a more tractable analysis. Furthermore, Smead (2014) argues that plas-

https://doi.org/10.1086/677685 Published online by Cambridge University Press

forcement learners (Herrnstein 1970; Roth and Erev 1995; Beggs 2005). These learners begin each repeated interaction by selecting a strategy stochastically according to an initial set of weights (1.0 for each strategy), playing that strategy, then observing the payoff and updating the weights by an amount equal to the payoff received. This process is repeated for a fixed number of rounds (10,000 in this case). At each iteration, the probability of choosing a strategy is equal to the cumulative payoffs received when choosing that strategy (plus the initial weight) divided by the total payoffs received to that point (plus all initial weights).

Evolutionary analysis proceeds in two phases. The first phase involves constructing a *type game* that includes L in addition to all pure-strategy types. For contests between non-L types, payoffs of the type game are identical to payoffs of the original game. In contests with L types, payoffs are estimated using 10^5 independent numerical simulations of learning in a repeated interaction and calculating the average payoff per round of interaction. The second phase of analysis is to apply the replicator dynamic to the type game and assess whether and when plasticity will evolve and persist.

Plasticity is frequently assumed to carry some cost. Roth-Erev reinforcement learning already carries the cost of inefficiency from possible errors and time to learn optimal behavior. Since we are only using the first 10,000 learning trials to estimate payoffs, this inefficiency will be captured in the type game. However, there may also be metabolic or other fitness costs associated with the ability to learn. For this reason, we will consider the evolutionary prospects for plasticity both with and without an exogenous cost (c).

3. Plasticity with Common-Interest Signaling. To represent perfect common interest, suppose both players receive a payoff of 1 when a_1 is chosen in s_1 and when a_2 is chosen in s_2, but receive 0 otherwise. Evolution in common-interest signaling games can be complex and does not always lead to perfect communication (Huttegger and Zollman 2011). In the simple case we are considering, however, evolution always leads to a signaling system: either (S1, R1) or (S2, R2).

Zollman and Smead (2010) have investigated the evolution of plasticity in settings of common-interest signaling using one-population models. They find that learners can initially increase in frequency within a population but invariably get eliminated by selection. Plasticity is useful when there is variation in the signaling strategies used by members of the population. Once plasticity becomes common, however, it is no longer worth

ticity without this ability will generally be disadvantageous in settings of social interaction.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

the cost of error—a strategy that signals in a fixed way will be accommodated by plastic individuals and perfectly coordinate with themselves.

It is known that the results of two-population models and one-population models do not always coincide, and this turns out to be the case here. Including learners in, say, the sender population does not change the stability of either signaling system. If the receiver population always plays a separating strategy, there is no benefit to behavioral flexibility, and the corresponding (fixed) signaling strategy does best. However, if the receiver population includes enough variation, it pays to be plastic. Type L can adopt S1 when interacting with R1's and S2 when interacting with R2's.[2] If there is a large degree of initial variation in the receiver population, this can lead to the evolutionary fixation of learning in the sender population even when there is a small exogenous cost to learning (see fig. 1). This also holds true if L is included in both populations. Without mutation, fixation of plasticity in one of the two populations occurred in 26.6% of simulations when $c = 0$ and 19.4% when $c = .05$. Thus, the results of Zollman and Smead (2010) do not generalize to the two-population case. Variation in the signaling behavior in one population can lead to the evolution of plasticity in the other population.

There are some important caveats to this result, however. First, the population states that include plastic individuals are not strongly stable. Once learners dominate one population, there is no fitness difference between the separating strategies in the other population because L will accommodate each equally well. If evolutionary drift leads to a predominance of either separating type, the corresponding type in the other population will outperform the plastic L. Second, it is not possible to maintain plasticity in both the sender and receiver populations simultaneously. The reason is that only one of the two players in each interaction needs to be plastic in order to ensure a high degree of signaling success. If one population is dominated by plastic individuals, fixed separating strategies will win out in the other population.

Additionally, simulations that include mutation show that plasticity does not persist. With $\mu = .001$ and $\mu = .01$, all simulated populations reached approximate signaling systems with a minimal frequency of L. Since mutation rates are positive, evolution never fixates on the signaling systems but gets very close (e.g., $\approx .993$ for $\mu = .001$).

However, if we examine simulated populations in the relatively short run with low mutation, we can see that plasticity occasionally persists for some time before being driven out. With no cost and $\mu = .001$, plasticity made up

2. Estimates from simulations show that L receives $\approx .967$ when interacting with separating strategies, .5 when interacting with pooling strategies, and $\approx .765$ when interacting with other learners.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

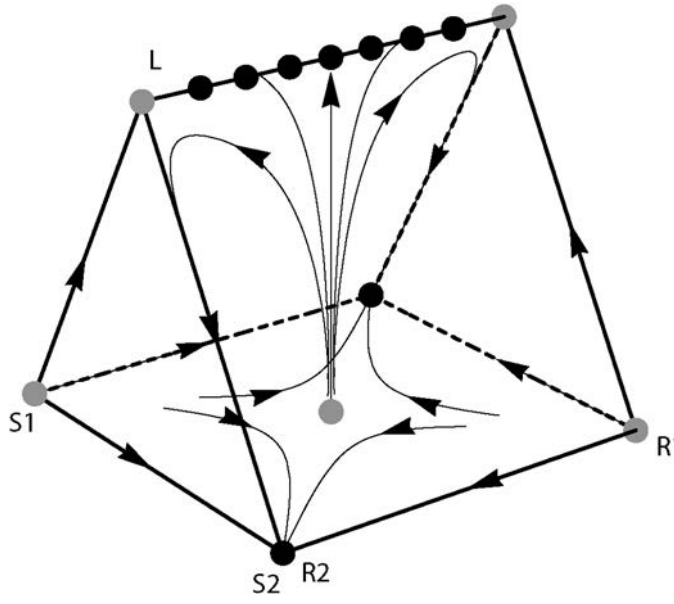


Figure 1. Replicator dynamic in a two-population common-interest sender-receiver model including learners in the sender population: $c = .05$. Pooling strategies have been removed for visual representation. Interior points represent population distributions for the sender population (in the L, S1, S2 space) and the receiver population (in the R1, R2 space). Arrows represent the direction of evolutionary change. Gray points are unstable rest points; black points are (Lyapunov) stable. Phase portraits were created using Mathematica 9.

more than .95 of one population in 22% of simulations at 10^3 generations. This figure dropped to 11.2% at 10^4 generations. Including costs or increasing the mutation rate greatly reduced the prospects for plasticity. A large frequency of plastic types in one population, when observed, was almost always accompanied by a significant mixture of nonplastic separating types in the other population. Although plasticity may not be prominent in the limit of the evolutionary processes, it can have a significant presence in the relatively short term.

4. Plasticity with Opposed-Interest Signaling. Now suppose we have the same setup as before, but where the receiver's payoffs are inverted. The players' interests are now directly opposed: the sender wants the receiver to do a_1 in s_1 and a_2 in s_2, but the receiver wants to do the opposite. Opposed-interest interactions are common in predator-prey interactions, and the evolutionary dynamics in such interactions frequently show cycles similar to the

https://doi.org/10.1086/677685 Published online by Cambridge University Press

well-known Lotka-Voltera models. In some cases of zero common-interest signaling games, the cycles are not regular but chaotic (Wagner 2012).

With directly opposed interests and including only S1, S2, R1, and R2, the population will cycle around the point where all strategies are represented equally. These cycles go through phases in which the senders systematically mislead the receivers and phases in which the receivers take advantage of the information contained in the senders' signals. If we include the other sender and receiver strategies, similar persistent cycles are observed, but they include some proportion of pooling strategies as well. Skyrms (2010) refers to the states in which the sender is manipulating the receiver as cases of "deception."

Skyrms defines a deception in informational terms as a signal that raises the probability of a state that is not the true state to the benefit of the sender and the detriment of the receiver. This definition is not uncontested (Godfrey-Smith 2012). For our purposes here, however, I take Skyrms's definition for granted and use the term "deception" in his sense. Deception is only possible in signaling games without perfect common interest; games with diverging interests are known to provide settings more conducive to the evolution of plasticity (Smead 2014).

Restricting the strategy space to separating strategies and including L in the sender population results in the evolution of plasticity (see fig. 2). This result is not as fragile as the evolution of plasticity in the common-interest case. Although the same Lyapunov-stable equilibrium structure is observed, these points are now the only attracting equilibria in the system. If drift leads to the predominance of one receiving strategy, evolution will favor one of the two nonplastic sender strategies. However, once such a nonplastic strategy invades, selection will cause the receiver population to shift and thus restore L's advantage. This restricted signaling game provides an interaction environment that strongly favors learning.

Including pooling strategies, however, shows that this result is not robust. Simulations both with and without mutation show that if plasticity carries a cost ($c = .05$), L never persists in large frequencies when all strategies are included. In every case, the populations are composed of pooling strategies almost entirely. The reason for this is that pooling strategies cannot be exploited by L, cost nothing to learn, and outperform separating strategies in the presence of learners. Costly plasticity does poorly in the short run as well. After 1,000 generations with no mutation or low mutation rates ($\mu = .001$), no simulations had significant numbers of plastic strategies in either population. The decline of costly plasticity happens more rapidly in the case of opposing interests than in the case of common interest, and separating strategies are driven out as well. In opposed-interest signaling games, costly plasticity and deception work against one another and neither evolves.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

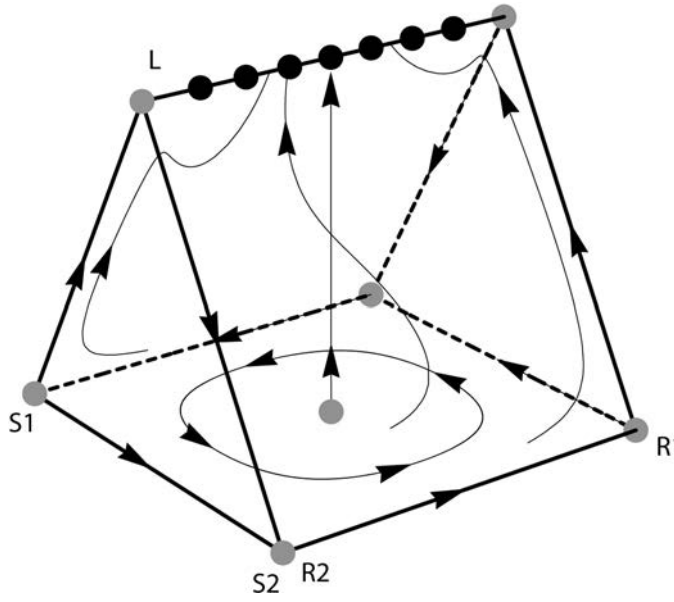


Figure 2. Replicator dynamic in a opposed-interest sender-receiver game including learners in the sender population: $c = .05$. Pooling strategies are omitted.

Costless plasticity fares much better. For $\mu = .001$, both populations converge to a state with ≈.599 L, .2 of each pooling type, and a small proportion of separating strategies maintained by mutation: $\mu = .01$ produces similar results. The constant small presence of separating strategies is what creates the advantage for L over pooling strategies. The possible, although rare, advantage of deception—and the advantage of avoiding being deceived—allows learners to persist at high frequencies. Despite this dynamic, these populations do not have a particularly high degree of deceptive behavior. When two learners interact, they often learn to adopt pooling strategies, or the receiver gets the upper hand and avoids defection. Consequently, the overall level of deceptive behavior in these populations is comparable to that of populations with no learners. So, although deception is where plasticity gains its advantage, the presence of L does not lead to a substantial change in the frequency of deception.

5. Plasticity with Partial Common Interest Signaling. Neither common nor opposed-interest signaling is particularly conducive to the evolution of costly plasticity. However, many instances of signaling games found in nature probably lie somewhere between these two extremes. Indeed, models of costly signaling—a central focus in the study of animal communication—

https://doi.org/10.1086/677685 Published online by Cambridge University Press

often involve partial common interest (Zahavi 1975; Maynard Smith and Harper 2003; Searcy and Nowicki 2005). In games with partial common interest, there are some states in which successful communication is in the interests of both parties and other cases in which there is an incentive to deceive.[3]

Consider a signaling game with the payoffs in table 1. In this case, when the state of the world is s_1 interests are aligned, and both players want the receiver to do a_2. When the state is s_2, the sender would prefer the receiver to do a_2, while the receiver would prefer to do a_1. If the signals are costless and the states are equiprobable, evolution will always lead to a pooling equilibrium where the sender population mixes between separating and pooling strategies but pools often enough that the receivers always ignore the signal and choose a_1. However, when signal costs of the right kind are introduced, the evolutionary dynamics change significantly.

Suppose that it is cheap to send m_1 in s_1 and m_2 in s_2 and that it costs .1 to send m_2 in s_1 or m_1 in s_2. This may occur because there is a natural salience to the signals that it takes some effort to disguise. Or, this may be because certain separating behavior generates some kind of payoff outside of the game taking place. Whatever the cause, a consequence of this cost structure is that, when facing a receiver who is ignoring the signal, it is better to adopt the separating S1 strategy. With these costs, there is a unique Nash equilibrium of the game: the sender plays S1 with probability 2/3 and S3 with probability 1/3, the receiver plays R2 with probability 1/10 and R3 with probability 9/10. This point is also the only interior rest point in the replicator dynamic—populations will fall into orbits around this equilibrium (see fig. 3). With $\mu = .01$, simulations converge to a point near the equilibrium with frequencies (.628138, .006637, .358536, .006690) of (S1, S2, S3, S4) and (.006579, .101911, .885286, .006224) of (R1, R2, R3, R4), respectively. With $\mu = .001$, simulations fall into to a regular orbit about the equilibrium with average frequencies near the equilibrium point.

3. A notable example is the Sir Philip Sidney game, which has been used to study begging behavior of chicks. In this game, the sender (the chick) is either healthy or needy and, knowing the state, can beg for food or remain silent. The receiver (the parent), upon hearing the signal, can either donate food or not. It is assumed that the chick always wants food, the parent only wants to donate when the chick is needy, and begging carries some cost. Dynamical analysis of this game is complex and has revealed that there are often multiple equilibria, some of which involve orbits of deception and partial communication that should be conducive to plasticity (Huttegger and Zollman 2010; Zollman, Bergstrom, and Huttegger 2013). However, because of the complexity of analysis, I focus on a simpler partial common interest game that has a unique equilibrium; the game is similar to that of Cho and Kreps (1987). Whether the Sir Philip Sidney game and related games are equally conducive to the evolution of plasticity is an open question for future research.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

TABLE 1. PAYOFFS FOR A PARTIAL COMMON INTEREST SENDER-RECEIVER GAME

	a_1	a_2
s_1	0, 0	5, 1
s_2	1, 5	2, 2

Deception is a regular occurrence in these populations. When S3 interacts with R2, S3 raises the probability that the state is s_1 when it is not, and R2 reacts by choosing an action that benefits the sender at the cost of the receiver. To determine whether this more complicated interaction is conducive to the evolution of plasticity, we can introduce L into the model as before. This generates the type game represented in table 2. There are three Nash equilibria in this type game:

1. $X \approx (.1105, .8895, 0, 0, 0)$ and $Y \approx (.0533, 0, 0, 0, .9467)$,
2. $X \approx (0, .6353, 0, 0, .3647)$ and $Y \approx (.5299, 0, 0, 0, .4701)$,
3. $X \approx (.6095, 0, 0, 0, .3905)$ and $Y \approx (0, .6481, 0, 0, .3519)$.

Simulations reveal that only the second and third are evolutionary significant. For $\mu = .01$, populations converge to one of two points that are

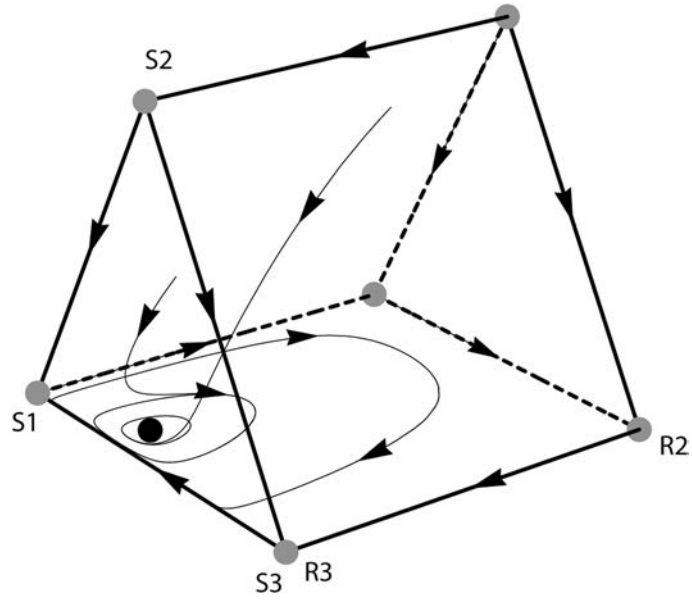


Figure 3. Replicator dynamic in a partial common interest sender-receiver game. Strategies S4, R1, and R4 omitted for visual clarity.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

TABLE 2. Type-Game Payoffs for a Partial Common Interest Sender-Receiver Game with Learners

	R1	R2	R3	R4	*L*
S1	1, 1	3, 3	.5, 2.5	3.5, 1.5	2.167, 2.779
S2	2.9, 3	.9, 1	.4, 2.5	3.4, 1.5	2.06, 2.779
S3	.45, 2.5	3.45, 1.5	.45, 2.5	3.45, 1.5	.686, 2.421
S4	3.45, 1.5	.45, 2.5	.45, 2.5	3.45, 1.5	.686, 2.421
L	3.238, 2.048	3.265, 2.088	.464, 2.500	3.452, 1.500	1.679, 2.433

Note.—Payoff with *L* is approximate (to nearest .001) and was estimated using 10^5 independent simulations.

near the second and third equilibria. For $\mu = .001$, populations fall into an orbit around one of those equilibria. In all cases, although learners are not predominant, they are sustained in both populations at relatively high frequencies. This holds with both high and low mutation rates as well as when learning carries an additional exogenous cost.[4]

The presence of learners also greatly increases the frequency of deceptive signaling. Without learners, the rate of successful deception at the equilibrium is relatively low—.033 of signals—due to the high proportion of the pooling strategy R3. Note that when *L* is included, the equilibria do not include any pooling types. Plastic types learn when to use deceptive signals as a sender and when to trust signals as a receiver. Even with mutation, the advantage of *L* keeps pooling types at small frequencies (less than .06 of any pooling type) and creates populations that are much more responsive to signals. This responsiveness is accompanied by higher degrees of deception, often by learners against separating receivers. Type *L* also can learn that separating senders should be trusted. In this game, plasticity gains an advantage by learning when to bluff and when to take a signal seriously.

6. Discussion and Conclusion. We have considered the evolution of plasticity in three cases of signaling interactions. The best case for the evolution of plasticity involves partial common interest where both honest communication and deception are incentivized. Evolution in this setting creates a varied and complex population where plasticity thrives. This result coheres with the environmental complexity thesis on the evolution of cognition. It is also in line with the idea that social settings involving both cooperation and conflict are important for the evolution of cognition more generally (Byrne and Bates 2007; Sterelny 2012).

4. At 10^6 generations with $c = 0$, the average proportion of *L* in both populations was .355 and .351 for $\mu = .001$ and $\mu = .01$, respectively. With $c = .05$, the average proportion of *L* was .374 and .340.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

Where should we expect to find signaling that is conducive to the evolution of plasticity? Paradigm cases of common-interest signaling might be communication among close relatives or communication among an organism's cells. Interests may diverge if signalers are not closely related or are members of different species. Thus, it may be reasonable to suspect that the important interactions for the evolution of plasticity are those among unrelated conspecifics or those that are interspecific. This also seems to be the case with common-interest signaling: two-population settings are more conducive to the evolution of plasticity than one-population settings. It is natural to interpret the two populations as representing two species, which suggests that interspecific interactions may be particularly important for the evolution of plasticity.

However, it is not necessary to interpret two-population models as representing interspecific interaction, nor is it true that interactions with kin and within-organism interactions are always common interest. Two-population models represent situations in which interactions are directed between (but not within) groups and in which the relevant standard of success is within (but not between) groups. This can occur within a single species (e.g., conditional interaction between different demes, interaction between sexes, or interaction between generational groups). It may also occur within an individual organism (e.g., an organism's cells are interacting with pathogens or symbionts). These cases are also ones that may involve opposed or partially aligned interests. Of course, real interactions usually involve both within- and between-population aspects, often simultaneously. No one would claim that two-population models capture most of the complicated nature of real interactions. Nonetheless, they can capture an important aspect of evolutionary interaction, which can occur within or between organisms, conspecific or heterospecific.

Along these lines, deception seems to occur in all areas of interaction. HIV is an example of a host-pathogen interaction that has been called "deceptive" (Krishnamoorthy et al. 2009). Stomatopods' aggressive display while molting, which renders them unable to fight, is an example of conspecific deception (Steger and Caldwell 1983). "Distracting displays by piping plovers are a case of inter-specific deception" (Ristau 1991). Interestingly, stomatopods and piping plovers are sometimes noted for their apparent cognitive abilities. HIV also displays structural plasticity (Korkut and Hendrickson 2012).

The models explored here show that signaling games of partial common interest can generate complex and varied interaction settings. In such a case, plasticity is able to invade populations and persist in relatively high frequencies. The reason for this is that plastic individuals are adept at both deceiving and avoiding being deceived. These models can be applied to a wide range of possible biological interactions ranging from within-organism signaling to interspecific signaling. It is important to note that plasticity may

https://doi.org/10.1086/677685 Published online by Cambridge University Press

be less likely to occur at the very small scale due to the complexity required for plastic responses. This aside, the key to the evolution of plasticity in signaling games seems to be the character of the signaling interaction rather than where the interaction takes place.

REFERENCES

Beggs, Alan W. 2005. "On the Convergence of Reinforcement Learning." *Journal of Economic Theory* 122:1–36.

Byrne, Richard W., and Lucy A. Bates. 2007. "Sociality, Evolution and Cognition." *Current Biology* 17:R714–R723.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics* 102:179–221.

Dubois, Frédérique, Julie Morand-Ferron, and Luc-Alain Giraldeau. 2010. "Learning in a Game Context: Strategy Choice by Some Keeps Learning from Evolving in Others." *Proceedings of the Royal Society* 277:3609–16.

Godfrey-Smith, Peter. 1996. *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.

———. 2012. Review of *Signals: Evolution, Learning, and Information*, by Brian Skyrms. *Mind* 120:1288–97.

Harley, Calvin B. 1981. "Learning the Evolutionarily Stable Strategy." *Journal of Theoretical Biology* 89:611–33.

Herrnstein, Richard J. 1970. "On the Law of Effect." *Journal of the Experimental Analysis of Behavior* 15:245–66.

Hofbauer, Josef, and Karl Sigmund. 1998. *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.

Humphrey, Nicholas K. 1976. "The Social Function of Intellect." In *Growing Points in Ethology*, ed. P. P. G. Bateson and R. A. Hinde, chap. 9, 303–17. Cambridge: Cambridge University Press.

Huttegger, Simon M., and Kevin J. S. Zollman. 2010. "Dynamic Stability and Basins of Attraction in the Sir Philip Sidney Game." *Proceedings of the Royal Society* 277:1915–22.

———. 2011. "Signaling Games: The Dynamics of Evolution and Learning." In *Language, Games, and Evolution*, ed. Anton Benz, Christian Ebert, Gerhard Jäger, and Robert van Rooij. New York: Springer.

Katsnelson, Edith, Uzi Motro, Marcus W. Feldman, and Arnon Lotem. 2012. "Evolution of Learned Strategy Choice in a Frequency-Dependent Game." *Proceedings of the Royal Society* 279: 1176–84.

Korkut, Anil, and Wayne A. Hendrickson. 2012. "Structural Plasticity and Conformational Transitions of HIV Envelope Glycoprotein gp120." *PLoS ONE* 7, e52170. doi:10.1371/journal.pone.0052170.

Krishnamoorthy, Lakshmi, Julian W. Bess, Alex B. Preston, Kunio Nagashima, and Lara K. Mahal. 2009. "Hiv-1 and Microvesicles from T Cells Share a Common Glycome, Arguing for a Common Origin." *Nature Chemical Biology* 5:244–50.

Lewis, David. 1969. *Convention*. Cambridge, MA: Harvard University Press.

Maynard Smith, John. 1982. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

Maynard Smith, John, and David Harper. 2003. *Animal Signals*. Oxford: Oxford University Press.

Ristau, Carolyn A. 1991. "Aspects of the Cognitive Ethology of an Injury-Feigning Bird, the Piping Plover." In *Cognitive Ethology: The Minds of Other Animals*, ed. Carolyn A. Ristau, chap. 5, 91–126. Hillsdale, NJ: Erlbaum.

Roth, Alvin E., and Ido Erev. 1995. "Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Intermediate Term." *Games and Economic Behavior* 8:164–212.

https://doi.org/10.1086/677685 Published online by Cambridge University Press

Searcy, William A., and Stephen Nowicki. 2005. *The Evolution of Animal Communication: Reliability and Deception in Signaling Systems*. Princeton, NJ: Princeton University Press.
Skyrms, Brian. 1996. *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
———. 2010. *Signals: Evolution, Learning, and the Flow of Information*. Oxford: Oxford University Press.
Smead, Rory. 2012. "Game Theoretic Equilibria and the Evolution of Learning." *Journal of Experimental and Theoretical Artificial Intelligence* 24:301–13.
———. 2014. "The Role of Social Interaction in the Evolution of Learning." *British Journal for the Philosophy of Science*, published online March 21. doi:10.1093/bjps/axt047.
Steger, Rick, and Roy L. Caldwell. 1983. "Intraspecific Deception by Bluffing: A Defense Strategy of Newly Molted Stomatopods (Arthropoda: Crustacea)." *Science* 221:558–60.
Sterelny, Kim. 2012. *The Evolved Apprentice*. Cambridge, MA: MIT Press.
Taylor, Peter D., and Leo B. Jonker. 1978. "Evolutionary Stable Strategies and Game Dynamics." *Mathematical Biosciences* 40:145–56.
Wagner, Elliott O. 2012. "Deterministic Chaos and the Evolution of Meaning." *British Journal for the Philosophy of Science* 63:547–75.
Zahavi, Amotz. 1975. "Mate Selection: A Selection for a Handicap." *Journal of Theoretical Biology* 53:205–14.
Zollman, Kevin J. S., Carl T. Bergstrom, and Simon M. Huttegger. 2013. "Between Cheap and Costly Signals: The Evolution of Partially Honest Communication." *Proceedings of the Royal Society*. doi:10.1098/rspb.2012.1878.
Zollman, Kevin J. S., and Rory Smead. 2010. "Plasticity and Language: An Example of the Baldwin Effect?" *Philosophical Studies* 147:7–21.

https://doi.org/10.1086/677685 Published online by Cambridge University Press