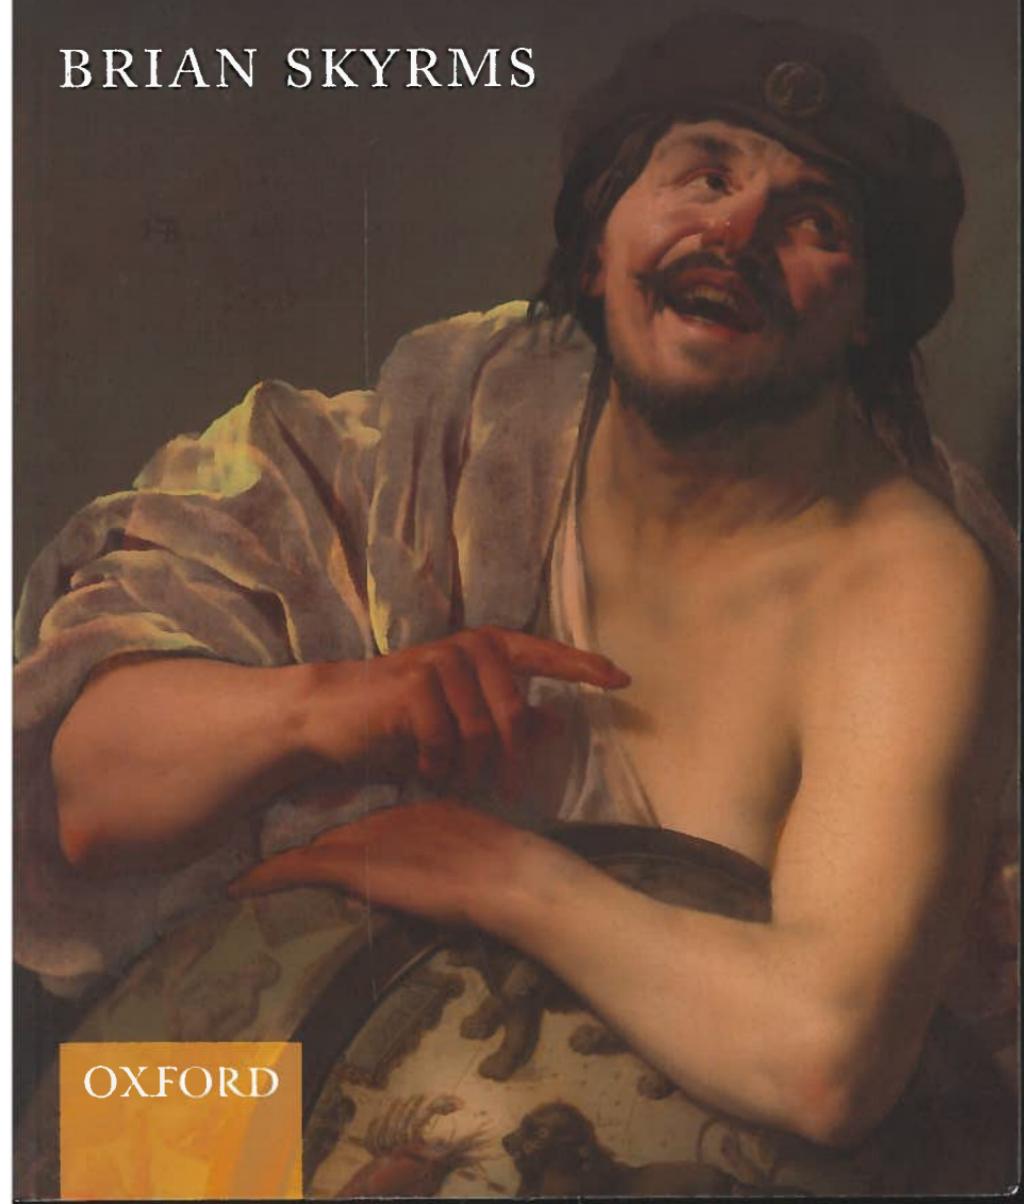


SIGNALS

Evolution, Learning, & Information

BRIAN SKYRMS



OXFORD

addresses cooperation and collective action—issues that usually reside in social and political philosophy. It does not quite fit into any of these categories, and gives each a somewhat novel slant. That's good, because the theory of signaling is full of fascinating unexplored questions.

1

Signals

“Two savages, who had never been taught to speak, but had been brought up remote from the societies of men, would naturally begin to form that language by which they would endeavor to make their mutual wants intelligible to each other . . .”

Adam Smith, *Considerations Concerning the First Formation of Languages*

What is the origin of signaling systems? Adam Smith suggests that there is nothing mysterious about it. Two perfectly ordinary people who did not have a signaling system would naturally invent one. In the first century BC, Vitruvius says much the same thing:

In that time of men when utterance of a sound was purely individual, from daily habits they fixed on articulate words just as they happened to come; then, from indicating by name things in common use, the result was in this chance way they began to talk, and thus originated conversation with one another.

Vitruvius is echoing the view of the great atomist Democritus, who lived four centuries earlier. Democritus held that signals were conventional and that they arose by chance.¹ Can it be true? If so, *how* can it be true?

¹ Another echo is to be found in Diodorus of Sicily:

“The sounds they made had no sense and were confused; but gradually they articulated their expressions, and by establishing symbols among themselves for every sort of object they came to express themselves on all matters in a way intelligible to one another. Such groups

The leading alternative view was that some signals, at least originally, had their meaning “by nature”—that is, that there was an innate signaling system.² At the time this may have seemed like an acceptable explanation, but after Darwin, we must say that it is no explanation at all. Bare postulation of an evolutionary miracle is no more explanatory than postulation of a miraculous invention. Either way, some work needs to be done.

Whatever one thinks of human signals, it must be acknowledged that information is transmitted by signaling systems at all levels of biological organization. Monkeys,³ birds,⁴ bees, and even bacteria⁵ have signaling systems. Multicellular organisms are only possible because internal signals coordinate the actions of their constituents. We will survey some of the signaling systems in nature in Chapter 2. Some of these signaling systems are innate in the strongest sense. Some are not.

We now have not one but two questions: *How can interacting individuals spontaneously learn to signal? How can species spontaneously evolve signaling systems?*

I would like to indicate how we can bring contemporary theoretical tools to bear on these questions.

came into existence throughout the inhabited world, and not all men had the same language, since each group organized their expressions as chance had it.”

Translation from Barnes 2001: 221.

See also Verlinski 2005 and Barnes 2001: 223. Proclus says:

“Both Pythagoras and Epicurus were of Cratylus’ opinion. Democritus and Aristotle were of Hermogenes” (5,2526).

and:

“Democritus who said that names are conventional formulated this principle in four dialectical proofs... Therefore names are arbitrary, not natural.” (6,207,1)

Translation from Duvick 2007.

² I am, of necessity, drastically oversimplifying the ancient debate here. See van den Berg 2008.

³ Cheney and Seyfarth 1990.

⁴ See Charrier and Sturdy 2005 for an avian signaling system with syntactical rules, and Marler 1999 for shadings of “innateness” in sparrow songs.

⁵ See the review article of Taga and Bassler 2003.

Sender-receiver

In 1969 David Lewis framed the problem in a clean and simple way by introducing sender-receiver games.⁶ There are two players, the sender and the receiver. Nature chooses a state at random and the sender observes the state chosen. The sender then sends a signal to the receiver, who cannot observe the state directly but does observe the signal. The receiver then chooses an act, the outcome of which affects them both, with the payoff depending on the state. Both have pure common interest—they get the same payoff—and there is exactly one “correct” act for each state. In the correct act-state combination they both get positive payoff; otherwise payoff is zero. The simplest case is one where there are the same number of states, acts, and signals. This is where we will begin.

Signals are not endowed with any intrinsic meaning. If they are to acquire meaning, the players must somehow find their way to information transmission. Lewis confines his analysis to equilibria of the game, although more generally we would want to investigate information transmission out of equilibrium as well. When transmission is perfect, so that the act always matches the state and the payoff is optimal, Lewis calls the equilibrium a *signaling system*. It is a virtue of Lewis’s formulation that we do not have to endow the sender and receiver with a pre-existing mental language in order to define a signaling system.

That is not to say that mental language is precluded. The state that the sender observes might be “What I want to communicate” and the receiver’s act might be concluding “Oh, she intended to communicate that.” Accounts framed in terms of mental language,⁷ or ideas or intentions can fit perfectly well within sender-receiver games. But the framework also accommodates signaling where no plausible account of mental life is available.

⁶ Russell 1921 is a precursor to Lewis. In an important paper, Crawford and Sobel 1982 analyze a model that generalizes signaling games in a different direction from that pursued here.

⁷ Such as Hurford 1989 and Komarova, Niyogi, and Nowak 2001.

If we start with a pair of sender and receiver strategies, and switch the messages around the same way in both, we get the same payoffs. In particular, permutation of messages takes one signaling-system equilibrium into another. This fundamental symmetry is what makes Lewis signaling games a model in which the meaning of signals is *purely conventional*.⁸ It also raises in stark form a question that bothered some philosophers from ancient times onward. There seems to be no *sufficient reason* why one signaling system rather than another should evolve. Of course, there may be many signaling systems in nature which got an initial boost from some sort of natural salience. But it is worth considering, with Lewis, the worst case scenario in which natural salience is absent and signaling systems are purely conventional.

Information in signals

Signals carry information.⁹ The natural way to measure the information in a signal is to measure the extent that the use of that particular signal changes probabilities.¹⁰ Accordingly, there are two kinds of information in the signals in Lewis sender-receiver games: information about what state the sender has observed and information about what act the receiver will take. The first kind of information measures effectiveness of the sender's use of signals to discriminate states; the second kind measures the effectiveness of the signal in changing the receiver's probabilities of action.¹¹

⁸ Some signaling interactions may not have this strong symmetry and then signals may not be perfectly conventional. There may be some natural salience for a particular signaling system. Here we are addressing the worst case for the spontaneous emergence of signaling.

⁹ I follow Dretske 1981 in taking the transmission of information as one of the fundamental issues of epistemology.

¹⁰ This can be measured in a principled way using the discrimination information of Kullback and Leibler 1951; Kullback 1959. We will look at this more closely in Chapter 3.

¹¹ Corresponding to these two types of information, we can talk about two types of content of a signal. See Russell 1921; Millikan 1984; Harms 2004.

Both kinds of information are maximal in a signaling-system equilibrium. But this does not uniquely characterize a signaling system. Both kinds of information can also be maximal in a state in which the players *miscoordinate*, and the receiver always does an act that is wrong for the state. Then there is plenty of information of both kinds, but it seems natural to say that information has not been successfully *transmitted* (or perhaps that *misinformation* is transmitted.)

Transmission of information clearly consists of more than the quantity of information in the signal. To deal with this example, you might think that we have to build in mentalistic concept of information—specifying what the sender *intended* the signal to mean and what the receiver took it to mean. Within the framework of Lewis signaling games this is not necessary. Sender and receiver have pure common interest. Perfect information about the state is transmitted perfectly if the receiver acts just as he would if he had direct knowledge of the state. As Democritus said, “The word is the shadow of the act.”¹²

A general treatment of information in signaling requires a lot more than this simple observation. In Chapter 3, I will develop a unified framework for both *informational quantity* and *informational content* of signals. The notion of informational content will be new, and will allow a resolution of some philosophical puzzles.

Evolution

As a simple explicit model of evolution, we start with the *replicator dynamics*.¹³ This has interpretations both for genetic evolution and for cultural evolution. The population is large, and either differential reproduction or differential imitation lead the population proportion of strategy A, $p(A)$, to change as:

¹² Barnes 1982: 468.

¹³ For a canonical reference, see Hofbauer and Sigmund 1998.

$$dp(A)/dt = p(A)[U(A) - U]$$

where $U(A)$ is the average payoff to strategy A and U is the average payoff in the population.

Evolutionary dynamics could operate on one population of senders and another of receivers as in some cases of interspecies communication, or it could operate on a single population, where individuals sometimes find themselves in the role of sender and sometimes in the role of receiver.

Consider the two-population model for the simplest Lewis signaling game—two states, two signals, two acts. Nature chooses a state by flipping a fair coin. And for further simplification, suppose the population has senders who only send different signals for different states and receivers who only perform different acts when they get different signals. There are then only two sender's strategies:

S_1 : State 1 => Signal 1

State 2 => Signal 2

S_2 : State 1 => Signal 2

State 2 => Signal 1

and only two receiver's strategies:

R_1 : Signal 1 => Act 1

Signal 2 => Act 2

R_2 : Signal 1 => Act 2

Signal 2 => Act 1

The pairs $\langle S_1, R_1 \rangle$ and $\langle S_2, R_2 \rangle$ are the signaling system equilibria. (We will consider varying the numbers of states, signals and acts, and the probabilities of the states, and the payoffs in subsequent chapters.)

The population dynamics lives on a square, with $p(S_2)$, the proportion of senders playing strategy S_2 , on the y axis and $p(R_2)$, the proportion of receivers playing strategy R_2 , on the x axis. It looks like this:

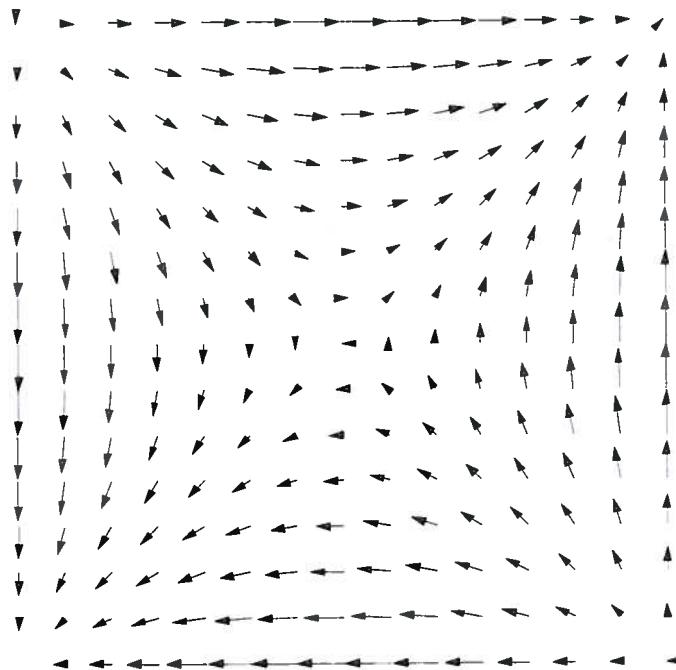


Figure 1.1: Replicator dynamics, two populations.

There are 5 dynamic equilibria—the four corners and one in the center of the square—but three of them are dynamically unstable. The two signaling systems are the only stable equilibria, and evolution carries almost every state of the combination of populations to either one signaling system or another.

Consider a one-population model where the agent's contingency plans, *if sender... and if receiver...* correspond to the four corners of the model we just considered. The dynamics lives on a tetrahedron. It looks like this:

The vertices are dynamic equilibria, and in addition there is a line of equilibria running through the center of the tetrahedron. But again, all the equilibria are unstable except for the signaling systems. All states to one side of a plane cutting through the tetrahedron are carried to one signaling system; all to the other side to the other signaling system.

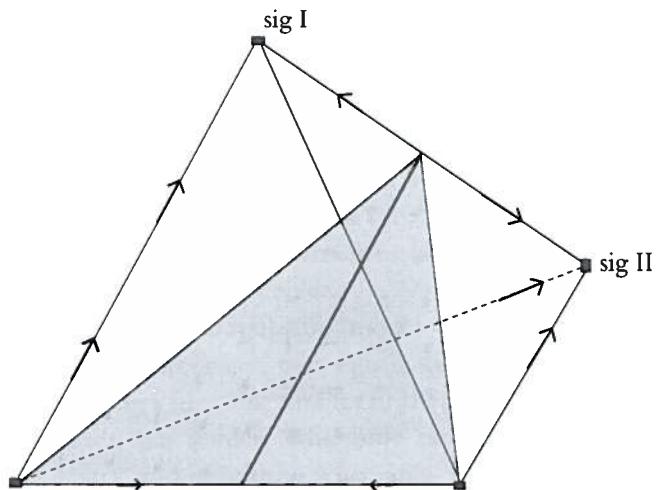


Figure 1.2: Replicator dynamics, one population.

Almost every possible state of the population is carried to a signaling system. More complex cases are discussed in Chapters 4 and 5.

We see in these simple cases how a perfectly symmetric model can be expected to yield an asymmetric outcome. In our two examples, the principle of sufficient reason is defeated by *symmetry breaking* in the evolutionary dynamics. The population moves to a signaling system as if—one might say—guided by an unseen hand.

Learning strategies

As a simple explicit model of unsophisticated learning, we start with reinforcement according to Richard Herrnstein's *matching law*—the probability of choosing an action is proportional to its accumulated rewards.¹⁴ We start with some initial weights, perhaps equal,

¹⁴ First proposed in Herrnstein 1970 as a quantification of Thorndike's law of effect, later used by Roth and Erev 1995 to model experimental human data on learning in games, by Othmer and Stevens 1997 to model chemotaxis in social bacteria, and by Skyrms and Pemantle 2000 to model social network formation.

assigned to each action. An act is chosen with probability proportional to its weight. The payoff gained is added to the weight for the act that was chosen, and the process repeats. As the weights build up, the process slows down in accordance with what psychologists call the law of practice.

Consider repeated interactions between two individuals, one sender and one receiver, who learn strategies by this kind of reinforcement. This set-up resembles the two-population evolutionary model, except that the process is not deterministic, but chancy. For a nice tractable example consider the two-state, two-signal, two-act signaling game of the last section. Computer simulations show agents always learning to signal, and learning is reasonably fast.

Learning actions

We helped the emergence of signaling in the foregoing model by letting reinforcement work on *complete strategies* in the signaling game—on functions from input to output. Essentially, the modeler has done some of the work for the learners. I take this as contrary to the spirit of Democritus, according to which the learners should not have to conceive of the problem strategically. Let us reconceptualize the problem by having reinforcement work on *single actions* and see if we still get the same result.

To implement this for the simplest Lewis signaling game, the sender has separate reinforcements for each state. You can think of it as an urn for state 1, with red balls for signal 1 and black balls for signal 2; and another such urn for state 2. The receiver also has two urns, one for each signal received, and each containing balls for the two acts. Nature flips a fair coin to choose the state. The sender observes the state and draws a ball from the corresponding urn to choose a signal. The receiver observes the signal and draws a ball from the corresponding urn to choose an act. The act is either successful, in being the act that pays off in that state, or not. Reinforcement for a successful act is like adding a ball of the

color drawn to the sender and receiver urn just sampled. The individuals are being reinforced for “what to do on this sort of occasion.” We can then ask what happens when these occasions fit together to form a signaling game.

This model appears to be more challenging than the one in the previous section. There are now four interacting reinforcement processes instead of two. Equilibria where the sender ignores the state and the receiver ignores the signal are no longer ruled out by appeal to the agents’ intelligence and good intentions. Nevertheless, there is now an analytic proof¹⁵ that reinforcement learning converges to a signaling system with probability one. The robustness of this result over a range of learning rules is discussed in Chapters 6 and 7.

States, acts, and signals

In the simplest Lewis signaling games, the number of states, acts, and signals are assumed to be the same. Why should this be so? What if there is a mismatch? There may be extra signals, or too few signals, or not enough acts. All these possibilities raise questions that are interesting both philosophically and mathematically.

Suppose there are too many signals. Do synonyms persist, or do some signals fall out of use until only the number required to identify the states remain in use? Suppose there are too few signals. Then there is, of necessity, an information bottleneck. Does efficient signaling evolve; do the players learn to do as well as possible? Suppose there are lots of states, but not many acts. How do the acts affect how the signaling system partitions the states?

If we have two states, two acts and three signals, we could imagine that the third signal gets in the way of efficient signaling, or that one signal falls out of use and one ends up with essentially a two-signal system, or that one signal comes to stand for one state and the other two persist as synonyms for the other state. Simulations of the

learning process of the last section always produce efficient signaling, often with the persistence of synonyms. Learning is about as fast as in the case where there are only two signals.

If we have three states, three acts and only two signals, there is an information bottleneck. The best that the players could do is to get it right $\frac{2}{3}$ of the time. This could be managed in various ways. The sender might use signals deterministically to partition the states—for example, send signal 1 in state 1 and signal 2 otherwise. An optimal receiver’s strategy in reply would be to do act 1 when receiving signal 1, and to randomize between acts 2 and 3 with any probability. This identifies a whole line of equilibria, corresponding to the randomizing probability. Alternatively, the receiver could be deterministic—for example, doing act 1 for signal 1 and act 2 for signal 2. If so, an optimal sender’s strategy to pair with this would always do sending signal 1 in state 1 and signal 2 in state 2, but randomizing in state 3. This identifies another line of efficient equilibria.¹⁶ There are, of course, also lots of inefficient equilibria. Simulations always deliver efficient equilibria. They are always of the first kind, not the second. That is to say the signaling system always partitions the states. Learning is still fast.

If we have three states, but only two signals and two acts, we can have act 1 right for state 1, and act 2 right for state 3, and then vary the payoffs for state 2:

Payoffs	State 1	State 2	State 3
Act 1	1	$1-e$	0
Act 2	0	e	1

If $e < .5$ it is best to have one signal (which elicits act 1) sent in both state 1 and state 2; and the other signal (which elicits act 2) sent in state 3. If $e > .5$ an efficient equilibrium lumps states 2 and 3 together. The optimal payoff possible depends on e : $\frac{1}{3}$ for $e = .5$ and 1 for $e = 0$ or $e = 1$. For the whole range of values, optimal

¹⁵ Argiento, Pemantle, Skyrms, and Volkov 2009.

¹⁶ Notice that these two lines share a point. If we consider all the lines of efficient equilibria, we have a cycle.

signaling emerges. These generalized signaling games are discussed in Chapter 8. The signaling game itself may not be fixed. The game structure itself may evolve. A model of signaling with invention of new signals is introduced in Chapter 9. The combination of simple signals to form complex signals is discussed in Chapter 11.

Signaling networks

Signaling is not restricted to the simple 1-sender, 1-receiver case discussed so far. Alarm calls usually involve one sender and many receivers, perhaps with some of the receivers being eavesdroppers from other species. Quorum signaling in bacteria has many individuals playing the role of both sender and receiver. The brain continually receives and dispatches multiple signals, as do many of its constituents. Most natural signaling occurs in networks. A signaling network can be thought of as a directed graph, with an edge directed from node A to node B signifying that A sends signals to B. All our examples so far have been instantiations of the simplest possible case; one sender sends signals to one receiver.



There are other simple topologies that are of interest. One that I discussed elsewhere¹⁷ involved multiple senders and one receiver. I imagined two senders who observed different partitions of the possible states.



In the context of alarm calls, if one sender observes a snake or leopard is present, and another observes that there is no snake, a receiving monkey might be well advised to take the action appropriate to evade a leopard. Multiple senders who transmit different information leave the receiver with a problem of logical inference.

¹⁷ Skyrms 2000, 2004.

It is not simply the problem of drawing a correct inference, but rather the problem of drawing the correct inference relevant to his decision problem. For instance, suppose sender 1 observes the truth value of p and then sends signal A or signal B, and sender 2 observes the truth value of q and sends C or D. Maximum specificity is required where the receiver has four acts, one right for each combination of truth values. But a different decision problem might require the receiver to compute the truth value (p exclusive or q) and to do one act if true and another if false.

Senders may observe different aspects of nature by chance, but they might also be able to choose what they observe. Nature may present receivers with different decision problems. Thus, a receiver might be in a situation where he would like to ask a sender to make the right observation. This calls for a dialogue, where information flows in both directions.



Nature flips a coin and presents player 2 with one or another decision problem. Player 2 sends one of two signals to player 1. Player 1 selects one of two partitions of the state of nature to observe. Nature flips a coin and presents player 1 with the true state. Player 1 sends one of two signals to player 2. Player 2 chooses one of two acts. Here a question and answer signaling system can guarantee that player 2 always does the right thing.

A sender may distribute information to several receivers.



One instance is the case of eavesdropping, where a third individual listens in to a two-person sender-receiver game, with the act of the third person having payoff consequences for himself, but not for the other two.¹⁸ In a somewhat more demanding setup, the sender sends separate signals to multiple receivers who then have to

¹⁸ There are also more complicated forms of eavesdropping, where the third party's actions have consequences for the signalers and there is conflict of interest. For a fascinating instance, where plants eavesdrop on bacteria, see Bauer and Mathesius 2004.

perform complementary acts for everyone to get paid. For instance, each receiver must choose one of two acts, and the sender observes one of four states of nature and sends one of two signals to each receiver. Each combination of acts pays off in exactly one state.

Signalers may form chains, where information is passed along.



In one scenario, the first individual observes the state and signals the state, and the second observes the signal and signals the third, which must perform the right act to ensure a common payoff. There is no requirement that the second individual sends the same signal that she receives. She might function as a translator from one signaling system to another.

When we extend the basic Lewis game to each of these networks, computer simulations show reinforcement learning converging to signaling systems—although a full mathematical analysis of these cases remains to be done. It is remarkable that such an unsophisticated form of learning can arrive at optimal solutions to these various problems. Simple signaling networks are discussed as a locus of information processing in Chapter 10 and as a component of teamwork in Chapter 13.

These networks are the simplest examples of large classes of phenomena of general interest. They also can be thought of as modules, which appear as constituents of more complex and interesting networks that process and transmit information. It is possible for modules to be learned in simple signaling interactions, and then assembled into complex networks by either reinforcement or some more sophisticated form of learning. The analogous process operates in evolution. The dynamics of formation of a simple signaling network is discussed in Chapter 14.

Conclusion

How do these results generalize? This is not so much a single question as an invitation to explore an emerging field. Even the simplest extensions of the models I have shown here are full of surprising and interesting phenomena. We have seen the importance of focusing on adaptive dynamics. The dynamics can be varied. On the evolutionary side, we can add mutation to differential reproduction. In addition, we might move from the large population, deterministic model of the replicator dynamics to a small population stochastic model. The mathematical structure of one natural stochastic model of differential reproduction is remarkably similar to our model of reinforcement learning.¹⁹ On the learning side, we should also consider more sophisticated types of learning. From considering evolution in a fixed signaling game we might move to evolution of the game structure itself. We should explore both signaling on various kinds of networks, but also the dynamics of formation of signaling networks. The rest of this book is an introduction to these topics.

We started with a fundamental question. Suppose we start without pre-existing meaning. Is it possible that, under favorable conditions, unsophisticated learning dynamics can spontaneously generate meaningful signaling? The answer is affirmative. The parallel question for evolution turns out to be not so different, and is answered in the same way. The adaptive dynamics achieves meaning by breaking symmetry. Democritus was right. It remains to explore all the ways in which he was right.

¹⁹ Schreiber 2001.