Ambiguity in Cooperative Signaling

Carlos Santana*

In game-theoretic signaling models, evolution tends to favor perfectly precise signaling systems, but in the natural world communication is almost always imprecise. I argue that standard explanations for this discrepancy are only partially sufficient, and I show that communication is often ambiguous because signal senders take advantage of context sensitivity. As evidence, I make two additions to the signaling model: a cost for more complex signaling strategies and the ability to combine information in signals with independent information. Analysis and simulation of the altered model shows that it leads to the predicted outcome of evolution favoring ambiguous signaling.

1. Introduction. Biological signaling is a messy affair. Rarely do signals correspond one-to-one with the states they represent or effects they cause. Nature more often than not employs a single signal for multiple jobs. This is true of cellular signals like hormones, nearly all of which perform multiple functions. The human peptide hormone cholecystokinin, for example, produces either satiety, anxiety, or both, depending on contextual factors (Greenough et al. 1998). It is true of animal vocalizations, such as meerkat alarm calls, which are reactions to a varied class of stimuli including "predators, large non-predatory animals, falling trees, and social disturbances within the group" (Seyfarth and Cheney 2003, 35). And it is especially true of natural language. Content words are usually polysemous or semantically ambiguous, and frequent words like 'run' and 'feel' are especially flexible in their meanings.

The frequency of these one-to-many mappings between biological signals and their meanings can be puzzling, given that at first glance they seem to be suboptimal compared to one-to-one signaling systems. It certainly seems as if it would better serve the meerkats' interests if they had separate signals for different types of threats or that it would be better for us *Homo*

Received June 2013; revised January 2014.

*To contact the author, please write to: Department of Philosophy, University of Pennsylvania, 433 Cohen Hall, Philadelphia, PA 19104; e-mail: csantana@sas.upenn.edu.

Philosophy of Science, 81 (July 2014) pp. 398–422. 0031-8248/2014/8103-0005$10.00
Copyright 2014 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

sapiens if a hormone released to stop us from eating too much did not sometimes cause incidental unwarranted anxiety. Yet although one-to-many signals are less informative and more error prone than their one-to-one counterparts in the abstract, I will show how under the right conditions—conditions under which most biological signaling occurs—they are not necessarily suboptimal and may even be favored by natural selection.

I do so within the signaling games framework introduced by Lewis in *Convention* (1969). Lewis uses game theory to explain how an arbitrary signal can obtain meaning through the need for conventions of coordination. Others have fruitfully developed Lewis's ideas in evolutionary models of natural signaling systems. Lewis and his successors, however, generally share an assumption that an optimal signaling system provides a one-to-one mapping from signals to states.[1] For example, Lewis remarks of a system of communication in which a single signal represents multiple states: "It is not a signaling system" (134). Likewise, Skyrms considers an evolutionary dynamic that leads to a smaller set of signals being used to signify a larger set of states "bad news" (2010, 68). Optimal signaling systems on this account must have a one-to-one correspondence between signals and the states of the world they represent. They resemble the theoretical logically perfect language of Frege and Russell, more than they resemble any natural signaling system, including human language.

The Lewisian model thus seems to fail to account for the ubiquity of imprecise (one-to-many) signals, especially since other theories often assign value to imprecision. For instance, Sowa (drawing on Pierce) argues that since natural languages have finite vocabulary, infinite extensibility, and limited utterance length, "vagueness and ambiguity are not defects in language, but essential properties that enable it to express anything and everything that people need to say" (2007, 17). And van Benthem (2000) hypothesizes that the need to balance expressive power and resource complexity explains the imprecision of language, but he wonders how to give this intuition formal game-theoretic support.

In what follows, I outline some plausible conservative adjustments to Lewisian signaling games that produce behavior in which ambiguous (one-to-many) signals are features, not bugs. Specifically, I show that in a model in which signal users are context sensitive, and precision bears a slight cost, ambiguous signaling systems can be advantageous. In doing so, I aim to give formal support to the plausible idea that imprecision can be economical and help enable communication. In the process, by showing how we can make Lewisian signaling models more closely resemble real-world communication, we increase the credibility of Lewis's account of the origins of communication.

1. What counts as a distinct state depends on the interests of the signal users (see sec. 3).

https://doi.org/10.1086/676652 Published online by Cambridge University Press

2. Signaling Games. Signaling games originated in Lewis (1969). In their basic form, signaling games are played by a *sender* and a *receiver*. The sender has privileged information about a state of the world; call this the sender's *type* $t \in \mathbf{T}$, where $\mathbf{T}$ is the set of possible types $\{t_1 \ldots t_n\}$. The sender plays first, selecting one signal $s \in \mathbf{S}$, where $\mathbf{S}$ is the set of possible signals $\{s_1 \ldots s_n\}$. The receiver notes the signal and then chooses a response $r \in \mathbf{R}$, where $\mathbf{R}$ is the set of possible responses $\{r_1 \ldots r_n\}$. Sender and receiver share a payoff function, which is defined by $\mathbf{T} \times \mathbf{R} \rightarrow \mathbb{R}$, meaning that each type-response pair is assigned a numerical utility to capture the agents' preferences. In Lewis's original exposition, the payoff function is restricted to $\mathbf{T} \times \mathbf{R} \rightarrow \{0, 1\}$, since the response is either correct given the type or not. Additionally, he stipulates that there is only one correct response to each type, so the payoff matrix is a permutation matrix of dimensions $n \times n$. A sender's strategy in this game consists of a one-to-one function from $\mathbf{T}$ to $\mathbf{S}$, and a receiver's strategy consists of a one-to-one function from $\mathbf{S}$ to $\mathbf{R}$.[2] Lewis defines a *signaling system* as a pair of strategies such that every play of the game yields a payoff of 1. That is, in a signaling system the receiver's response to the signal is always the correct response to the type which elicited that signal. This can only occur when signals perfectly represent types, and signaling systems are Nash equilibria (NE) of a coordination problem.[3]

To illustrate, consider the following game: $\mathbf{T} = \{\text{hot, cold}\}$, $\mathbf{S} = \{\text{'warm', 'chilly'}\}$, and $\mathbf{R} = \{\text{jacket, shorts}\}$. Figure 1 gives the payoff matrix. Ignoring those strategies in which agents simply ignore the signal, the sender has two possible pure strategies: $\{\langle \text{hot, 'warm'} \rangle, \langle \text{cold, 'chilly'} \rangle\}$ and $\{\langle \text{hot, 'chilly'} \rangle, \langle \text{cold, 'warm'} \rangle\}$. Likewise the receiver has two possible strategies: $\{\langle \text{'warm', shorts} \rangle, \langle \text{'chilly', jacket} \rangle\}$ and $\{\langle \text{'warm', jacket} \rangle, \langle \text{'chilly', shorts} \rangle\}$.[4] Of the four possible pairings of sender's and receiver's strategies, two correctly match the receiver's response to the type. These two pairings are the signaling systems, and since neither party can benefit by unilaterally switching strategies if the players are already receiving the payoff, the two signaling systems are NE. Note that the assignment of signals to types is purely arbitrary. Neither signaling system is superior, so whether 'warm' means hot or cold depends only on which equilibrium state the players reach first. 'Warm' and 'chilly' were meaningless signs before the game, but because of the sender's and receiver's interest in coordinating, they gain meaning through cooperative use.

2. This ensures that strategies must be *pure* strategies: they involve no nondegenerate probabilistic behavior.

3. An NE is a game state in which no player would benefit from unilaterally changing strategies.

4. Send the signal 'warm' when it is hot and 'chilly' when it is cold.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

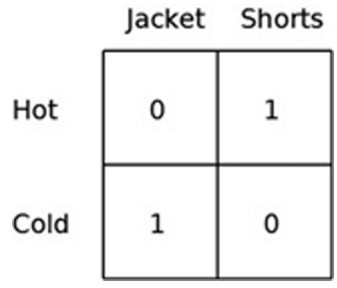


Figure 1. Payoff matrix for 2 × 2 game.

While Lewis deals only with games of pure coordination in which sender and receiver share a payoff function, economists have explored the development of signals in situations of partial conflict. In an influential paper, Crawford and Sobel (1982) analyze a signaling game in which **T**, **S**, and **R** are continuous rather than discreet sets, and the sender and receiver have independent payoff functions. Among their results is the conclusion that perfect signaling typically does not take place in situations with divergent interests. Kreps and Sobel (1994) explore games with a larger repertoire of actions and show that agents can design mechanisms to elicit more informative and honest signals. These developments of signaling games, however, assume that agents are cognitively sophisticated and have existing media of communication. Given our interest in the Lewisian project of explaining the origins of conventional meaning, we cannot make those assumptions.

More relevant are developments of the model from an evolutionary perspective. Evolutionary game theory (Maynard Smith 1973) makes no assumption that players are cognitively sophisticated, only that they use a heritable strategy. Practitioners of more successful strategies reproduce at a higher rate, and the long-term behavior of a population of mixed strategies can be explored using either simulations or analysis. An early analysis of signaling games using evolutionary game theory found that the only evolutionarily stable strategies (ESS) of perfectly cooperative signaling games were Lewis's signaling systems (Wärneryd 1993).[5] Wärneryd's analysis suggests that a cooperative population should be able to develop signaling systems out of a world with no communication, a promising extension of Lewis's argument to populations of subrational agents. Subsequent innovations have generalized this result to a variety of evolutionary signaling games and populations of learners (Huttegger 2007; Skyrms 2010; Hut-

5. A strategy is an ESS if, when the entire population is playing an ESS, no other strategy can invade. An ESS must therefore be an NE, but not all NE are ESS.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

tegger and Zollman 2011). Since we find signaling systems throughout the living world, even in organisms as simple as bacteria, evolutionary signaling games provide a useful means to explore the origins of the common phenomenon of communication.

3. Imprecision. Evolutionary signaling games have been successful in modeling the development of signaling systems, but the explanatory goal is the evolution of natural communication, not Lewis's perfect one-to-one systems. A signaling system as Lewis defines it has no imprecision, but in actuality signals successfully transmit information even given some departures from the ideal. A sender's strategy, for example, could involve using more than one signal to signify a particular type on separate occasions. In this situation of *synonymy*, the signaling system still carries full information. But information is lost in the inverse case, when a sender's strategy assigns a single signal to multiple types. In *pooling*, the sender sends the same signal for every type, and nothing is communicated. But we are interested in the intermediate case, where individual signals signify a restricted set of types. This situation, *partial pooling*, corresponds to the phenomenon of communicative *ambiguity*. Figure 2 depicts each of the three types of imprecise sender's strategies, with 'T' standing for types and 'S' for signals.

I have claimed that sender's strategies without a one-to-one type to signal function are imprecise, but without a principled method for counting and differentiating types this is an empty claim. As Lewis correctly observed (1969, 135), we can reclassify the set of types to transform any perfect signaling system into an imperfect one, or vice versa. Consider the ambiguous strategy depicted in figure 2, which is imprecise because S2 signifies T2 and T3. We could redescribe the game by introducing a new type, T2* = (T2 or T3), in which case the sender's strategy is a perfectly precise one-to-one function: $\{\langle T1, S1\rangle, \langle T2^*, S2\rangle\}$. Whether a signaling system is imprecise depends on the granularity of our model description. As long as we can arbitrarily reclassify types, imprecision is a property of the model only and not its target system.

In response we need a justified constraint on type classification. Given that we explain signal use in terms of benefit to signal users, the correct constraint is that relevant properties affect the sender's and receiver's payoffs.[6] To put it another way, a property is sufficient to delimit a distinct type if the best response of the receiver depends on the presence of that property. Grouping states of the world according to a receiver's best re-

6. We must resist the temptation to think that types should match on to some sort of natural or metaphysical kind. For a useful discussion of why the way signals partition the world comes apart from metaphysics, see Barrett (2007).

https://doi.org/10.1086/676652 Published online by Cambridge University Press

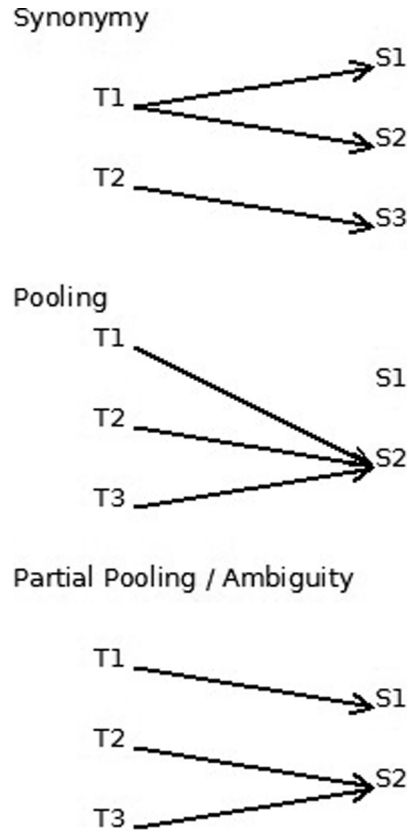


Figure 2. Imprecise sender's strategies.

sponse gives us a principled method for classifying types, so we can discuss imperfect signaling systems without our models being vulnerable to redescription that eliminates imprecision.

Additionally, it bounds the set of communicative phenomena we are interested in. Linguists often restrict the term 'ambiguity' to surface forms that have multiple underlying syntactic or semantic structures (Zwicky and Sadock 1975). If the underlying structure itself can represent a range of states of affairs, this is not 'ambiguity' in their sense but *indeterminacy*. This conception of ambiguity should not be confused with the current phenomenon of interest. Signaling games encompass a broad range of biological communication, most of which is performed without complex underlying psycholinguistic structures. So the line I have drawn between ambiguous and unambiguous signaling systems necessarily departs from the usage of

https://doi.org/10.1086/676652 Published online by Cambridge University Press

'ambiguity' in linguistics and includes most cases of what linguists call ambiguity as well as some cases of indeterminacy. We are kept from having to accept all cases of indeterminacy as ambiguous by the conclusion that only indeterminacy that has action-guiding consequences counts as ambiguity. In these terms, our central question is, Why do real-world signalers not send a unique signal for each best response? The next section presents some proposed answers to this question and explains why they are insufficient.

4. A Few Attempts to Explain Ambiguity. 'Ambiguity' here means partial pooling: a pairing of sender's and receiver's strategies in which at least one signal signifies more than one type.[7] Ambiguous strategy pairings are by Lewis's definition not signaling systems. Furthermore, unambiguous signaling is favored in basic rational signaling games (Lewis 1969), in basic evolutionary games simulated under both the replicator dynamics (Skyrms 2010) and the Moran process (Huttegger and Zollman 2011), on grid structures (Zollman 2005), and from multiple varieties of reinforcement learning (Skyrms 2009; Huttegger and Zollman 2011). If the evolution of signaling systems really works in accord with these models, then evolution often discourages ambiguity. So why is ambiguity so ubiquitous? And what about our intuition that ambiguity can be useful?

One response denies that ambiguity is useful and explains the emergence of ambiguity as the failure of a system to reach perfect signaling for incidental reasons. For example, if a game is stipulated to have fewer signals than types, senders are forced to be ambiguous (Skyrms 2010, 15).[8] Donaldson, Lachmann, and Bergstrom (2007) claim that this explains the prevalence of partial pooling in animal communication. But although restricted lexicon size might contribute to ambiguity, it does not seem to fully explain it. Consider the signaling behavior of putty-nosed monkeys. Arnold and Zuberbühler (2006, 651) report that male putty-nosed monkeys use a series of 'hack' calls to alert others to the presence of eagles but also use the same call "when baboons were fighting close by, when trees fell near the monkeys, when branches broke as they landed on them, in response to a rape alarm and when a duck flew close by." Males of the same species have a second 'pyow' call that sees a similar range of use, as well as calls made up of a mixed sequence of 'hack' and 'pyow' (Arnold and Zuberbühler 2008). If this ambiguity is due solely to the restricted lexicon size, then we need an explanation for why the lexicon contains only these signals.

7. Specifically, an ambiguous signal signals more than one type disjunctively. A signal that represents more than one type conjuctively (i.e., 'A and B') is unambiguous.

8. If DO NOTHING is an available action, senders can also just ignore some states and signal the rest perfectly.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

Given that monkeys are reasonably intelligent, with some species capable of learning dozens of commands, it is unlikely that the putty-nosed monkeys are limited to three signals by their cognitive capacity. They may not be smart enough to learn a human-sized vocabulary, but certainly they could have had a slightly larger lexicon of five or six signals. Nor is it plausible that physiological constraints on production limit the putty-nosed monkeys to only two sound sequences. Indeed, Arnold and Zuberbühler (2006) report that the monkeys can also emit 'booms' and 'chirps' but that these are rarely used by mature males and were not produced in response to any experimental stimuli. The monkeys' reliance on predominantly two signals is therefore unlikely to result from psychological or physiological constraints. When it comes to *Homo sapiens* it is even more clear that ambiguity cannot be explained wholly by a limited repository of signals. Human speakers are generally more ambiguous than we are required to be by the constraint of our lexicon size. Human miscommunication, although frequent, can usually be avoided by conveying additional information or using less ambiguous language. We cannot fully explain ambiguity merely by saying there are not enough signals to go around.

Various variants of the basic Lewis game do lead to ambiguous signaling. If the distribution of types is uneven, the receiver might do well enough by always using the best response for the most common types and thus never learn to pay attention to signals. For this reason, in games with more than two signals and states, uneven type distribution can lead to partial pooling (Huttegger 2007). This will explain some cases of pooling or partial pooling, but it does not explain cases in which the distribution of types is fairly even, nor does it explain why, even in cases of uneven distribution, ambiguity is not just a possibility but nearly an inevitability.

Similarly, signaling games with more than two types, signals, and responses sometimes lead to ambiguous signaling systems under some dynamics (Barrett 2006). Pawlowitsch (2008) discovered that this occurs because partial pooling can be a neutrally stable strategy (NSS). Whereas ESS are asymptotically stable,[9] under simple replicator dynamics, NSS are only Lyapunov stable.[10] But Lyapunov stability is enough to guarantee that the population can be trapped in a state of partial pooling, despite its suboptimality. The occasional evolution of ambiguity, according to this explanation, is just a mathematical inevitability.

9. Nearby states in the state space will converge toward the state where the entire population plays the ESS.

10. If the state is near enough to the state where the entire population plays the NSS, it will stay near that state forever. This is a weaker type of stability than asymptotic stability.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

Some ambiguity might be explained by the fact that evolution sometimes leads to suboptimal outcomes, but we should not overestimate how much of the prevalence of ambiguity this explains. We are motivated by a suspicion that ambiguity makes a positive contribution to communication, so an explanation that writes off ambiguous signaling systems as suboptimal is unsatisfying. Moreover, Pawlowitsch shows that NSS are not stochastically stable, meaning that small perturbations can cause the population to jump out of a partial pooling trap. In biological evolution, mutation provides just such a source of perturbation, and it turns out that under dynamics with perturbations representing mutation perfect signaling is favored more strongly (Huttegger 2007; Skyrms 2010; Huttegger and Zollman 2011). For learning dynamics, forgetting plays a similar role in helping the agents escape partial pooling strategies (Barrett and Zollman 2009). Having a larger set of potential signals can also increase the probability that a system converges to perfect signaling (Barrett 2006). In short, once we include realistic perturbations in signaling models we cannot so easily chalk up ambiguity to simple mathematical bad luck.

In one variant of signaling games, however, ambiguity does inevitably emerge and play an important role in enabling communication. Up to this point we have discussed signaling games as games of pure common interest. But much of communication in the living world occurs in situations of partial conflict, and partial conflict gives communicators reason to be only partially informative. Consider the case (Aragones and Neeman 2000) of a political candidate speaking to voters about her political platform. Assume she has the option of either (1) specifying precisely what she would do once in office or (2) characterizing her position in broad, ambiguous terms that merely situate her within a certain range on the political spectrum. She should almost certainly choose 2, since it allows her to appeal to a broader range of voters, while simultaneously avoiding the costly commitments that action 1 would tie her to. Although the voters would benefit from precision, the politician benefits from only communicating partial information. The mismatch of interests between sender (politician) and receiver (voters) gives the sender reason to use ambiguous signals. Crawford and Sobel (1982) prove this in the general case, demonstrating that the level of pooling in signaling games increases as the degree of common interest diminishes. In zero-sum games in which agents share no common interest, only total pooling or babbling equilibria emerge.[11] In pure coordination games in which agents have interests entirely in common, perfect signaling systems are the only strict NE. In the intermediate case—partial

11. In total pooling, the sender sends the same signal regardless of type. In babbling equilibria, the sender sends a random signal regardless of type. Signals carry no information either way.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

conflict—perfect signaling is not stable, so if any communication is to occur at all it must be ambiguous.

Partial conflict is thus able to explain much of the ambiguity in the real world, but I do not think it accounts for all of it. After all, almost all real-world signals are ambiguous, but much communication occurs in situations in which conflict is minimal. If I ask you to 'Get me the book on the table' when there is more than one book on the table, it is not because I have any interest in deceiving you as to which book I want. And if I report that 'John chased Bill until he ran out of breath', leaving the anaphor for 'he' ambiguous, it is unlikely that I profit from your uncertainty about whether it was John or Bill who tired first. So while partial conflict is an effective explanation for some ambiguity, it fails to explain a large number of ambiguous utterances.

To recap: although ambiguity is everywhere in human and animal communication, models based on Lewis signaling games predict the frequent emergence of one-to-one mappings from signals to types. While partial pooling does sometimes occur in evolutionarily models, evolutionary signaling games generally favor less ambiguous signaling systems, especially when the models include stochastic elements such as mutation. Limitations on the set of available signals and situations of partial conflict do produce ambiguity in signaling games, but these explain only a subset of cases of ambiguity, especially when it comes to communication among highly intelligent and cooperative organisms.[12] Furthermore, many theorists have argued that ambiguity serves a communicative function in arbitrating a trade-off between expressive range and precision. Lewis signaling games thus fail to capture both the prevalence of ambiguity and its utility. Let us see whether we can solve that problem.

5. Cost, Context, and Ambiguity. To make signaling games able to accommodate ambiguity in cooperative contexts, we need to alter the model. The adjustments I am going to make are motivated by two facts about the model's target system. First, in determining a response, receivers generally have more information about the state of the world than merely what is contained in the signal. Although the sender has some privileged information, receivers still have perceptual data and prior expectations that may bear on the probability of types. The receiver's additional information may allow the sender to get away with ambiguity. Adding receiver's information

12. Signaling games in which the distribution of types is uneven can also result in ambiguity. For example, if in a 2×2 game one type occurs much more frequently than the other, the receiver can do well enough by always using the best response for that type. It thus never learns to pay attention to signals. This fact, however, is a better explanation for total pooling in these cases rather than partial pooling in general.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

to the model should not be controversial,[13] because context nearly always affects interpretation of both human and animal communication. Considering, for example, the wide range of circumstances—eagle, baboons, rape—under which the putty-nosed monkeys emit 'pyow', they must be using contextual information to determine the correct response. Including receiver's information merely accommodates this type of context sensitivity.

Second, signals and strategies vary in cost. In the basic Lewis signaling game, utilities are a function of types and responses alone, but it could be that either the particular signal used affects utility or strategies themselves have varying costs.[14] If in some situations ambiguity is effective enough to allow successful coordination but less costly than precision, cost would explain both why ambiguity exists and why it is useful. A skeptical reader might insist that the variance in signal cost would be trivial and that precision would not always be more costly than ambiguity anyhow. I grant that the variance in utterance cost would be minimal, but when all else is equal even minimal differences matter. In the model I will outline, cost differences will be minimal but still have an effect.

I can also grant that ambiguity will not always be less costly than precision, but I think that more often than not it is. To support this claim, I appeal to both common sense and various findings in psycholinguistics. In signaling, the cost of a signal boils down to the players' preferences and the physiological and cognitive load incurred in producing or interpreting the signal. For human language, at least, it seems plausible that more marked, complicated utterances are less preferred and require greater cognitive time and effort.[15] Research in psycholinguistics has confirmed that production is more difficult as sentences are more complex (Ferreira 1991) and that stuttering is more frequent at the beginning of longer and more complex sentences due to the difficulty of motor planning (Tornick and Bloodstein 1976; Jayaram 1984). Since longer, more complex utterances tend to be more precise, these facts suggest that there is a greater cognitive and physiological cost to precision. Additionally, children acquiring a first language tend to learn and produce semantically general terms more readily than more specific words (Theakston et al. 2004). It is debated whether this preference for general terms is due to cognitive factors or merely an effect of expo-

13. Nor am I the first to do so. Although their aims are distinct from mine, Zollman (2011) and Franke (2012) both discuss signaling models in which the receiver has private information.

14. I take this idea from Rubinstein (1998).

15. In fact, this is a commonly accepted assumption in the game-theoretic literature on human communication. See, e.g., Parikh (2000) and Jäger (2013).

https://doi.org/10.1086/676652 Published online by Cambridge University Press

sure frequency. If it is due to cognitive factors, then ambiguous terms would be less cognitively costly. Alternatively, if this effect is due to a higher frequency of exposure, this would suggest a general preference for using more ambiguous words—which also supports my claim that ambiguity is less costly. As for adult senders' preferences, Herbert Clark and collaborators (Clark and Wilkes-Gibbs 1986; Wilkes-Gibbs and Clark 1992) have repeatedly documented a preference among speakers to employ the minimal level of precision required to establish reference, and they have proposed a "principle of least common effort" to capture this preference. Similarly, studies using visual world eye tracking have demonstrated that in establishing reference, speakers prefer to use ambiguous terminology as long as pragmatic factors will allow the hearer to successfully disambiguate (Brown-Schmidt, Campana, and Tanenhaus 2002). Finally, in a signaling game experiment with human subjects, Agranov and Schotter (2010) found that sender's strategies involving too large a lexicon tended to confuse participants and decrease the efficiency of communication, suggesting that precise strategy profiles can be costly as well.

I could go on, but hopefully these few examples will convince the skeptic that I have empirical support in assigning a cost to precision. The preference for ambiguity in these examples, I should note, appears to be nonmonotonic. While more ambiguous utterances tend to be preferred over the most precise options, communicators do not always prefer the most ambiguous option, such as a short grunt. It is telling, however, that we are perfectly happy to communicate even with a short grunt if that grunt can get the job done. This suggests a possible preference for ambiguity going all the way down, but one that is overridden by the need for effective communication. The analysis to follow explores this balance between the cost of precision and the benefit of clear communication.

We have established motivation for adjusting the signaling game model in two ways: adding contextual information and assigning a cost to signals. For simplicity's sake, we will treat more precise sender's strategies as costlier. The modified game looks as follows:

A set of types $\mathbf{T} = \{1, 2, 3, 4\}$.
A set of signals $\mathbf{S} = \{a, b, c, d\}$.
A set of receiver's information $\Sigma = \{A, B\}$.
A set of responses $\mathbf{R} = \{I, II, III, IV\}$.
Sender's strategies $f(\mathbf{T}) \rightarrow \mathbf{S}$.
We restrict the space of sender's strategies to the following two:
standard strategies, which are one-to-one from **T** to **S**.
ambiguous strategies, which map two elements of **T** each to two elements of **S** and leave two elements of **S** unused.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

For present purposes we ignore other possible strategies because we are specifically interested in the success of ambiguous and standard strategies when competing against each other.

Receiver's strategies $f(\mathbf{S}, \Sigma) \rightarrow \mathbf{R}$.

Agents consisting of a strategy pair: a sender's strategy paired with a receiver's strategy.

A payoff function taking as inputs the type, the response, and the kind of sender's strategy.

Specifically, if a player's sender's strategy is ambiguous, its payoff is 1 if the Arabic value of the type is equivalent to the Roman value of the response, and 0 otherwise. For example, if the type is 4, response IV receives a payoff of 1, and all other responses receive a payoff of 0. If the player's sender's strategy is standard, the payoff is $1 - \epsilon$ if the values are equivalent, and 0 otherwise. Note that the cost of precision in this game is arbitrarily small.

The game is played as follows:

1. Nature moves first and randomly assigns one player to be sender and one player to be receiver. This symmetrizes the game.
2. Nature randomly chooses a type and reveals it to the sender.
3. If the type is in $\{1, 2\}$, nature reveals A to the receiver; if the type is in $\{3, 4\}$, it reveals B. (If the receiver's strategy is sensitive to this information, it will partition **T**, and thus nature's signal from Σ represents the receiver's background and contextual knowledge.)
4. The sender sends a signal according to its strategy.
5. The receiver chooses a response according to its strategy.
6. Payoffs are awarded, and the game ends.

By introducing a cost to standard strategies, it may seem that I have merely presented a restriction on lexicon size in another guise. But the cost I have attached to strategies that use four signals rather than two is not a hard cap: unambiguous strategies are still available. Additionally, the size of the cost (ϵ) is outweighed by the benefit from successful coordination (1), so it is far from a trivial result that the cost will matter. Hard caps on lexicon size, by contrast, trivially ensure that ambiguity will result, by not providing enough signals to enable precision.

It may also appear that I have introduced a slight amount of conflict into the model by allowing the sender and the receiver to receive a payoff that differs by ϵ. But true conflict involves not only a difference in payoff function but also a difference in preference ordering. In my modified signaling game, players prefer the same outcomes in all cases regardless of their

https://doi.org/10.1086/676652 Published online by Cambridge University Press

strategy profile, so it is still a nonconflictual coordination game. With these concerns addressed, we can turn to analyzing the game.

What we discover is that any population of standard signalers can eventually be invaded by ambiguous signalers, but populations of ambiguous signalers can only be invaded by other ambiguous strategies. A full proof is found in the appendix, but I give a summary here. Intuitively it might seem obvious that the evolutionary endpoint of this game would be ambiguous signaling since it has a lower cost, but one of the insights of game theory is that a result need not be optimal to be an equilibrium. For example, in the simple coordination game defined by the payoff matrix in table 1, both {A, A} and {B, B} are strict NE, even though the payoff for {A, A} is less by half. This is because neither player would benefit from deviating from a scenario in which each is already choosing A.

We need to show that a similar situation does not hold in the modified signaling game. In other words, we need to show that in a situation in which the entire population is using standard signaling, one player could benefit from switching strategies.

Consider a finite population of signal users under a dynamic fitting the following criteria: (1) individuals reproduce in proportion to their relative fitness, which is determined on the basis of their payoffs from playing the signaling game, and (2) individuals have some probability to switch to a random strategy each generation.[16] This strategy switching simulates mutation, the (re)introduction of strategies into the population. It also allows for drift, the redistribution of strategy frequencies independent of fitness.

Imagine that the entire population is using an identical standard signaling strategy, one that always accurately transmits information to itself. Consider, for example, the strategy whose signal for type 1 is "a" and whose response to "a" is I, with similar arrangements for types 2–4. Each member of this population will receive the respectable payoff of $1 - \epsilon$, which is better than almost any other strategy would expect if it joined this population. In most cases, an ambiguous signaler who tried to join the population would die off because it could not communicate—even if it understood the population's signals, the standard signalers would probably not be able to understand its own, so it would receive a low expected payoff.

The strategies that signalers use for receiving in this game, however, have to include eight contingency plans: one for each pairing of a sender's signal and nature's signal. Only four of these contingencies matter to a uniform population like the one we have imagined: in our example, the signal pairings of ⟨a, A⟩, ⟨b, A⟩, ⟨c, B⟩, and ⟨d, B⟩. Other possible pairings like

16. This can be generalized to the infinite case, in which mutation would mean a small proportion of the population adopts a random strategy.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

TABLE 1. Simple Coordination Game

	A	B
A	1	0
B	0	2

⟨a, B⟩ will never come up. This means that a few other standard strategies are able to join the previously uniform population, and these strategies would differ only in how they would respond to the unused pairings.

Thus, the population would not remain uniform but drift randomly among the standard strategies that could all communicate with each other. Some configurations of the four unused contingency plans, however, would allow a strategy to successfully interpret a particular ambiguous signaler. For instance, if an ambiguous signaler signaled "a" for both 1 and 3, and a member of our example population had a previously unused contingency plan to respond to ⟨a, B⟩ with response III, it would understand the ambiguous signaler. The right ambiguous signaler could then join the population and both understand and be understood. Furthermore, it would receive a payoff of 1, higher than its standard conspecifics. It would then drive out the standard signalers. So an ambiguous signaler could invade a population of standard signalers (but not vice versa). In other words, in this modified signaling game, 'languages' with maximally precise signals are vulnerable to invasion by ambiguous signalers. This analysis demonstrates that in a four type, four response signaling game with contextual information and a cost to precision, natural selection favors ambiguous signaling strategies over perfect signaling.

Although this suggests that there are pathways for a population to change from standard to ambiguous signaling but not vice versa, that fact alone does not guarantee that ambiguous signaling is the only (or even a primary) evolutionary endpoint (Huttegger and Zollman 2013). The actual evolutionary behavior of a population will depend on the particulars of the dynamics, and one way to explore the dynamic properties of the game is by using individual-based modeling.

The purpose of my first simulation was to see how rapidly a small number of ambiguous signalers could invade a population of standard signalers given plausible parameters. In the initial condition, the population consists of 999 equivalent standard signalers and one ambiguous signaler with the same receiver's strategy. Fitness is set to the payoff of the game and is implemented as the probability of reproduction; that is, if an individual plays and receives a payoff of 1, that individual reproduces with 100% probability. The dynamics are a birth-death process (Nowak 2006): each individual plays against a random member of the population each generation to de-

https://doi.org/10.1086/676652 Published online by Cambridge University Press

termine whether it reproduces, and if it does it creates a clone of itself that replaces a random member of the population. With ϵ set at .03, I ran 5,000 simulations of this dynamic for 750 generations. In all but six simulations ambiguous signalers had completely driven out the original population of standard signalers. In the remaining six, standard signalers comprised less than 10% of the population. The rate at which ambiguous signalers replace standard signalers would increase at higher values for ϵ and decrease at lower values, and time to total invasion depends on population size. Nevertheless, this simulation makes it clear that, given the right conditions, ambiguous signaling could replace standard signaling in a relatively small number of generations.

I designed a second set of simulations to highlight the mechanisms leading to the adoption of ambiguous strategies. For one simulation I initialized a population of 300 individuals each of which selects a strategy randomly from among the possible strategies. The dynamics of the game are a discrete counterpart to the replicator dynamics (Taylor and Jonker 1978): at each time step the new proportion of the population playing a strategy S is set to the old proportion playing that strategy multiplied by the ratio of the fitness of that strategy to the average fitness. As in the previous simulation, ϵ is set to .03. No mutation occurs.

Without mutation, the population tends to adopt whichever is the most prevalent behavior, with slight advantage to ambiguous signaling. Over 500 trials of 100 generations each, the proportion of signalers using an ambiguous strategy was, on average, .51, with most trials ending up with the population entirely at one extreme as shown in figure 3. This suggests that the evolution of ambiguity is not merely a trivial effect of adding a cost to standard signaling strategies. A population will tend to move toward uniformly using the most popular signaling strategy regardless of a slight inefficiency.

Compare that result with a second simulation. Population size and number of generations are the same, but reproduction includes a mutation probability of .01. Mutations are point adjustments to the offspring's strategy—a change from response III to response II at one point of the receiver's strategy, for example: 500 trials of 100 generations each produced the distribution in figure 4. With mutation, ambiguous signaling is favored much more strongly. It also takes the population longer to reach local attractors, so intermediate proportions persist in higher numbers. Running the simulation for more generations produces fewer populations in an intermediate state and a greater probability that a nearly homogenous population will practice ambiguous signaling. These results suggest two things. First, comparison with the nonmutation simulation suggests that populations which would have reached a state of standard signaling without mutation are pushed by mutation into ambiguous signaling (as we would expect from the analytical proof). Second, in a dynamic with fairly realistic parameters, a small population with no

https://doi.org/10.1086/676652 Published online by Cambridge University Press

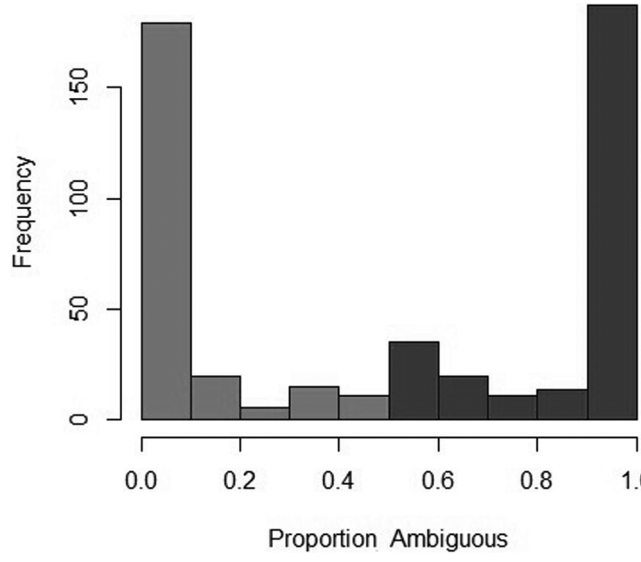


Figure 3. Replicators without mutation.

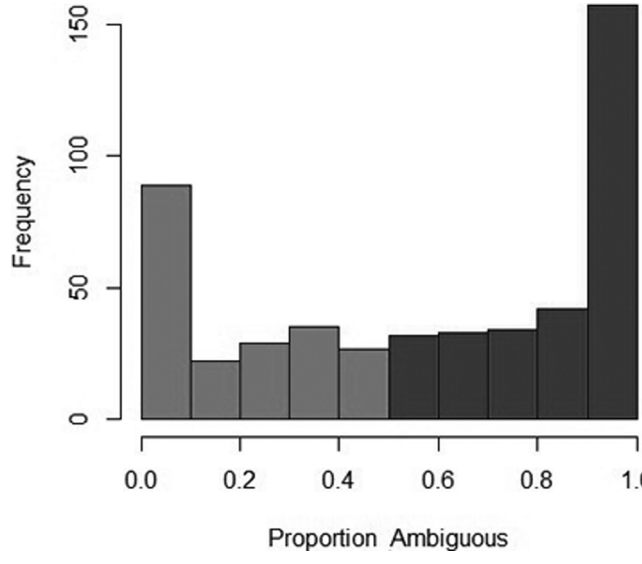


Figure 4. Simulation with mutation.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

initial coordination can move to consist of mostly ambiguous signalers in relatively few generations.

6. Conclusion. I have shown that two simple, plausible adjustments to the basic Lewisian signaling game are sufficient to enable ambiguity to evolve. In systems in which more complex signaling strategies are more costly and in which receivers receive some incomplete information from the context, partial pooling is generally favored over perfect precision.

With regard to our understanding of communication in general, the model I have explored adds support to the theory that ambiguity serves a useful communicative function even in cooperative contexts. Ambiguity allows senders to off-load some of the work of information transfer to the context, allowing for successful coordination with less effort. In other words, ambiguity can contribute to the efficiency of communication. By implication, we should consider treating ambiguity as part of the explanatory goal of modeling communication.

Some ideas for how to further pursue this: in this article I have focused on the relationship between ambiguity and the concerns of the signaler, so one direction for future work is to explore the receiver's contextual information. Assigning a cost to accessing contextual information[17] and varying its informativeness[18] are both promising directions for further modeling variations. In this work I have left the model general in scope, but it could be refined to fit more closely with specific target systems. I hope, by laying this ground, that in the future instead of viewing models that predict ambiguity as "bad news" or an unhappy outcome, we can align our modeling more closely with real-world phenomena and theoretical intuition by aiming to create models in which ambiguity is a feature, not a bug.

Appendix

THEOREM. If individuals reproduce in proportion to their relative fitness, which is determined on the basis of their payoffs from playing the modified signaling game, and individuals have some probability to switch to a random strategy each generation, then all populations either can (a) be invaded by ambiguous signaling strategies or (b) evolve to population compositions that can in turn be invaded by ambiguous signaling strategies. This is not true for nonsignaling and standard signaling strategies.

Proof. Define a *signaling strategy* as a strategy that, in any given play of the game against itself, is guaranteed a strictly positive payoff. All strate-

17. Suggested by Ryan Muldoon.
18. Suggested by an anonymous reviewer.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

gies are classified as *standard* or *ambiguous* according to the sender's strategy employed. The function $u(x, y)$ gives the expected payoff (EV) for strategy x when playing a game against strategy y.

Strategies will be represented as tables, with the upper panel representing a player's sender's strategy and the lower panel representing its receiver's strategy. Asterisks are wildcards, so what follows is the class of all possible strategies (with the proviso that the sender's strategy must be admissible):

[*name*]	(1)	(2)	(3)	(4)
f(**T**)	*	*	*	*
	(a)	(b)	(c)	(d)
f(**S**, A)	*	*	*	*
f(**S**, B)	*	*	*	*

QED

LEMMA 1. No strategy is an ESS.

Proof. Because this is a symmetrized version of a game with information asymmetry, only pure strategies can be ESS (Selten 1980). For any pure strategy, at most four of the eight cells in the receiver's strategy table determine the strategy's payoff against itself. The four cells can be populated with any member of **R**.

ss1	(1)	(2)	(3)	(4)
f(**T**)	a	b	c	d
	(a)	(b)	(c)	(d)
f(**S**, A)	I	II	III	IV
f(**S**, B)	I	II	III	IV

For example, $u(\text{ss1}, \text{ss1}) = 1 - \epsilon$ (ss1 is a standard signaling strategy). But consider the class of strategies ss1*:

ss1*	(1)	(2)	(3)	(4)
f(**T**)	a	b	c	d
	(a)	(b)	(c)	(d)
f(**S**, A)	I	II	*	*
f(**S**, B)	*	*	III	IV

https://doi.org/10.1086/676652 Published online by Cambridge University Press

Here, u(any member of ss1*, any member of ss1*) $= 1 - \epsilon$. Since ss1 is a member of ss1*, ss1 will perform only as well against members of ss1* as it does against itself, and they will perform just as well against themselves as they do against ss1. Thus, ss1 will not be able to drive them out through selection.

Since, like ss1, all strategies will only have four cells in the receiver's strategy table that determine the strategy's payoff against itself, every strategy will have at least 255 other strategies that could not be driven out by selection were they to invade. So no strategy is an ESS. QED

DEFINITION Let *equivalence class* (EClass) refer to a set of strategies, such as ss1*, in which every member of the set does exactly as well as every other member of the set when paired with any member of the set.

LEMMA 2. Populations consisting of members of an EClass of nonsignal ing strategies can be invaded by signaling strategies.

Proof. Since a population consisting solely of members of an EClass will drift among all members of the class, it is sufficient to show that one member of the class is invasible to show that the population is.

Consider first the case of an EClass in which the average payoff is 0. Any signaling strategy can invade because it will receive a nonzero payoff when paired with itself. EClasses in which the average payoff is >0 must consist of only one type of strategy (standard or ambiguous), since when standard and ambiguous strategies are paired, they never receive identical nonzero payoffs.

Consider an arbitrary standard nonsignaling strategy, ns1: u(ns1, ns1) $= p$, where, because ns1 is a nonsignaling strategy, $p < 1 - \epsilon$. Let ss1 be a standard signaling strategy with the same sender's strategy as ns1: u(ss1, ns1) $= (1 - \epsilon + p)/2$, which is greater than p. Furthermore, u(ss1, ss1) $= 1 - \epsilon$, which is also greater than p. Because ss1's payoffs are higher, it can invade ns1.

Consider an arbitrary ambiguous nonsignaling strategy, an1. Either an1's signal combined with nature's signal will fully partition **T** (i.e., an1 sends one signal for $\{1, 3\}$ and one signal for $\{2, 4\}$) or an1's signal will provide no information not found in nature's signal (like nature, it sends one signal for $\{1, 2\}$ and one for $\{3, 4\}$). In the case that **T** is fully partitioned, an1 will be invaded by an ambiguous signaling strategy for reasons analogous to those in the previous paragraph. Consider, then, the case in which an1 provides no information not provided by nature: u(an1, an1) $\leq 1/2$. Take, for example, EClass an1*:

https://doi.org/10.1086/676652 Published online by Cambridge University Press

an1*	(1)	(2)	(3)	(4)
f(**T**)	a	a	b	b
	(a)	(b)	(c)	(d)
f(**S**, A)	I	*	*	*
f(**S**, B)	*	III	*	*

There is a member of an1's equivalence class (call this an2) whose receiver's function is identical to the receiver's function of an ambiguous signaling strategy (call this as1). At best, $u(\text{an2}, \text{an2}) = 1/2$. But $u(\text{as1}, \text{as1}) = 1$, and as1 would have an EV of 1 when playing sender against an2. When playing receiver against an2, both players would have an EV in $[0, 1/2]$. Consequently, $u(\text{as1}, \text{an2}) \geq 1/2$. In other words, as1 would be at least as good at responding to an2 as an2 is to itself but better at playing against itself than an2 is. Thus, a population of an2 can be invaded by as1. Since arbitrary EClasses of standard and both types of ambiguous nonsignaling strategies can be invaded by signaling strategies, all populations consisting of EClasses of nonsignaling strategies are invasible. QED

LEMMA 3. Populations consisting of members of an EClass of standard signaling strategies can be invaded by ambiguous signaling strategies.

Proof. Let ss1 be any standard signaling strategy, with ss1* as its EClass: ss1* will have two wildcards in the upper row of the receiver's function and two in the lower row. For instance:

ss1*	(1)	(2)	(3)	(4)
f(**T**)	a	c	b	d
	(a)	(b)	(c)	(d)
f(**S**, A)	I	*	II	*
f(**S**, B)	*	III	*	IV

Choose one nonwildcard cell from each row of the receiver function of ss1*, making sure to pick one member of {I, II} and one member of {III, IV}. For each cell you picked, fill in the wildcard of the vertically adjacent cell with the complement of the other cell you picked. For example, if you picked a cell with a I and a cell with a III, put a IV vertically adjacent to the I and a II vertically adjacent to the III. Fill in the remain wildcards however you would like. For example:

https://doi.org/10.1086/676652 Published online by Cambridge University Press

ss2	(1)	(2)	(3)	(4)
f(**T**)	a	c	b	d
	(a)	(b)	(c)	(d)
f(**S**, A)	I	II	II	?
f(**S**, B)	IV	III	?	IV

Call this new strategy, a member of ss1's equivalence class, ss2. The receiver's strategy in ss2 is identical to the receiver's strategy of an ambiguous signaling strategy. This must be the case because the method you used to construct ss2 ensured that all four members of **R** are uniquely identified by a function from Σ and only two members of **S**—the two heading the columns of the cells you chose. Take those two members of **S** and place them on the f(**T**) row of a strategy table under the Arabic counterparts of the Roman numerals found in their column. Keep the receiver's strategy the same as in ss2. This strategy (call it as1) is an ambiguous signaling strategy that shares a receiver's function with ss2. Continuing with the same example, we have

as1	(1)	(2)	(3)	(4)
f(**T**)	a	b	b	a
	(a)	(b)	(c)	(d)
f(**S**, A)	I	II	II	?
f(**S**, B)	IV	III	?	IV

Because both are signaling strategies and both share a receiver's function, ss2 and as1 coordinate perfectly with themselves and with each other. In all these interactions, however, ss2 has an EV of $1 - \epsilon$, but as1 has an EV of 1.[19] Thus, ss2 is vulnerable to invasion by as1. Since ss1* is a generic EClass of standard signaling strategies, all populations consisting of members of an EClass of standard signalers can be invaded by ambiguous signalers. QED

LEMMA 4. Populations consisting of equivalence classes of ambiguous signaling strategies can only be invaded by other ambiguous signaling strategies.

Proof. Ambiguous signaling strategies have an EV of 1 when playing against themselves. They therefore have an EV of 1 when playing against

19. This is the point where all else is equal, so the cost of precision matters even if vanishingly slight.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

other members of their EClass. Since 1 is the highest possible EV, ambiguous signaling strategies are best responses to every member of their equivalence class. To invade, a mutant would need an EV of 1 against some member of the EClass and an EV of 1 against itself. Only ambiguous signaling strategies ever attain an EV of 1 against any strategy, so only ambiguous signaling strategies could invade a population made up of members of an EClass of ambiguous signaling strategies. QED

LEMMA 5. Equilibrium population states consisting of multiple EClasses of signaling strategies are susceptible to transition via invasion to equilibrium states in which the population consists of members of only one EClass.

Proof. Consider a population consisting of any number of distinct EClasses of signaling strategies. Call each EClass a *subpopulation*. Allow a single mutation to produce an additional member of any one subpopulation. Before the mutation, average fitness must be $< 1 - \epsilon$, or else the population would already consist of only one EClass. After the mutation, members of the mutant's subpopulation would have $\text{EV} \geq 1 - \epsilon$ against the mutant. Their fitness would therefore increase postmutation. Members of other subpopulations with $\text{EV} < 1 - \epsilon$ against the mutant would have lower fitness relative to the mutant's subpopulation, so they would be driven out of the population. Members of other subpopulations could, however, have $\text{EV} \geq 1 - \epsilon$ against the mutant. If this means that not all nonmutant subpopulations are driven out, repeat the process with a new mutant from any surviving subpopulation until only one remains. QED

The modified signaling game has no noninvasible equilibria, but we can characterize a set of transition rules between equilibrium population states. Lemmas 1 and 5 indicate that we are concerned with transitions between equilibrium populations consisting of EClasses. Lemma 2 specifies a rule allowing transition from nonsignaling to signaling equilibria. Lemma 3 specifies a rule allowing transition from standard signaling equilibria to ambiguous signaling equilibria. Lemma 4 shows that transitions from ambiguous to ambiguous are allowed, but transitions from ambiguous to nonambiguous are not. So although no particular ambiguous signaling system is stable, once a population has fully adopted ambiguous strategies, ambiguity itself is.

REFERENCES

Agranov, Marina, and Andrew Schotter. 2010. "An Experimental Study of Ambiguity and Vagueness in the Announcement Game." Working Paper, New York University.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

Aragones, Enriqueta, and Zvika Neeman. 2000. "Strategic Ambiguity in Electoral Competition." *Journal of Theoretical Politics* 12 (2): 183–204.
Arnold, Kate, and Klaus Zuberbühler. 2006. "The Alarm-Calling System of Adult Male Putty-Nosed Monkeys, *Cercopithecus Nictitans Martini*." *Animal Behaviour* 72 (3): 643–53.
———. 2008. "Meaningful Call Combinations in a Non-human Primate." *Current Biology* 18 (5): R202–R203.
Barrett, Jeffrey. 2006. "Numerical Simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar." Paper 54, Institute for Mathematical Behavioral Sciences, University of California, Irvine.
———. 2007. "Dynamic Partitioning and the Conventionality of Kinds." *Philosophy of Science* 74 (4): 527–46.
Barrett, Jeffrey, and Kevin Zollman. 2009. "The Role of Forgetting in the Evolution and Learning of Language." *Journal of Experimental and Theoretical Artificial Intelligence* 21 (4): 293–309.
Brown-Schmidt, Sarah, Ellen Campana, and Michael Tanenhaus. 2002. "Reference Resolution in the Wild: On-Line Circumscription of Referential Domains in a Natural Interactive Problem-Solving Task." In *Proceedings of the 24th Annual Meeting of the Cognitive Science Society*, ed. Wayne D. Gray and Christian Schunn. Austin, TX: Cognitive Science Society.
Clark, Herbert, and Deanna Wilkes-Gibbs. 1986. "Referring as a Collaborative Process." *Cognition* 22 (1): 1–39.
Crawford, Vincent, and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.
Donaldson, Matina, Michael Lachmann, and Carl Bergstrom. 2007. "The Evolution of Functionally Referential Meaning in a Structured World." *Journal of Theoretical Biology* 246 (2): 225–33.
Ferreira, Fernanda. 1991. "Effects of Length and Syntactic Complexity on Initiation Times for Prepared Utterances." *Journal of Memory and Language* 30:210–33.
Franke, Michael. 2012. "On Assertoric and Directive Signals and the Evolution of Dynamic Meaning." *International Review of Pragmatics* 4 (2): 232–60.
Greenough, A., G. Cole, J. Lewis, A. Lockton, and J. Blundell. 1998. "Untangling the Effects of Hunger, Anxiety, and Nausea on Energy Intake during Intravenous Cholecystokinin Octapeptide (CCK-8) Infusion." *Physiology and Behavior* 65 (2): 303–10.
Huttegger, Simon. 2007. "Evolution and the Explanation of Meaning." *Philosophy of Science* 74 (1): 1–27.
Huttegger, Simon, and Kevin Zollman. 2011. "Signaling Games: Dynamics of Evolution and Learning." In *Language, Games, and Evolution*, ed. Anton Benz, Christian Ebert, Gerhard Jäger, and Robert van Rooj, 160–76. Berlin: Springer.
———. 2013. "Methodology in Biological Game Theory." *British Journal for the Philosophy of Science* 64 (3): 637–58.
Jäger, Gerhard. 2013. "Rationalizable Signaling." *Erkenntnis*. doi:10.1007/s10670-013-9462-3.
Jayaram, M. 1984. "Distribution of Stuttering in Sentences." *Journal of Speech and Hearing Research* 27:338–41.
Kreps, David, and Joel Sobel. 1994. "Signalling." In *Handbook of Game Theory with Economic Applications*, ed. Robert J. Aumann and Sergiu Hart, 849–67. New York: Elsevier.
Lewis, David. 1969. *Convention*. Oxford: Blackwell.
Maynard Smith, John. 1973. "The Logic of Animal Conflict." *Nature* 246 (5427): 15–18.
Nowak, Martin. 2006. *Evolutionary Dynamics*. Cambridge, MA: Harvard University Press.
Parikh, Prashant. 2000. "Communication, Meaning, and Interpretation." *Linguistics and Philosophy* 23 (2): 185–212.
Pawlowitsch, Christina. 2008. "Why Evolution Does Not Always Lead to an Optimal Signaling System." *Games and Economic Behavior* 63 (1): 203–26.
Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.
Selten, Reinhard. 1980. "A Note on Evolutionarily Stable Strategies in Asymmetric Animal Conflicts." *Journal of Theoretical Biology* 84 (1): 93–101.
Seyfarth, R. M., and D. L. Cheney. 2003. "Meaning and Emotion in Animal Vocalizations." *Annals of the New York Academy of Sciences* 1000 (1): 32–55.

https://doi.org/10.1086/676652 Published online by Cambridge University Press

Skyrms, Brian. 2009. "Evolution of Signalling Systems with Multiple Senders and Receivers." *Philosophical Transactions of the Royal Society* B 364 (1518): 771–79.
———. 2010. *Signals: Evolution, Learning and Information*. Oxford: Oxford University Press.
Sowa, John. 2007. "Language Games, a Foundation for Semantics and Ontology." In *Game Theory and Linguistic Meaning*, ed. Ahti-Viekko Pietarinen, 17–37. Oxford: Elsevier.
Taylor, Peter D., and Leo B. Jonker. 1978. "Evolutionarily Stable Strategies and Game Dynamics." *Mathematical Biosciences* 40 (1): 145–56.
Theakston, Anna, Elena Lieven, Julian Pine, and Caroline Rowland. 2004. "Semantic Generality, Input Frequency and the Acquisition of Syntax." *Journal of Child Language* 31 (1): 61–99.
Tornick, Gaye, and Oliver Bloodstein. 1976. "Stuttering and Sentence Length." *Journal of Speech and Hearing Research* 19 (4): 651–54.
van Benthem, Johan. 2000. "Economics and Language." In *Economics and Language*, 93–107. Cambridge: Cambridge University Press.
Wärneryd, Karl. 1993. "Cheap Talk, Coordination, and Evolutionary Stability." *Games and Economic Behavior* 5:532–46.
Wilkes-Gibbs, Deanna, and Herbert Clark. 1992. "Coordinating Beliefs in Conversation." *Journal of Memory and Language* 31 (2): 183–94.
Zollman, Kevin. 2005. "Talking to Neighbors: The Evolution of Regional Meaning." *Philosophy of Science* 72:69–85.
———. 2011. "Separating Directives and Assertions Using Simple Signaling Games." *Journal of Philosophy* 108 (3): 158–69.
Zwicky, Arnold, and Jerrold Sadock. 1975. "Ambiguity Tests and How to Fail Them." *Syntax and Semantics* 4 (1): 1–36.

https://doi.org/10.1086/676652 Published online by Cambridge University Press