

Talking to Neighbors: The Evolution of Regional Meaning*

Kevin J. S. Zollman^{†‡}

In seeking to explain the evolution of social cooperation, many scholars are using increasingly complex game-theoretic models. These complexities often model readily observable features of human and animal populations. In the case of previous games analyzed in the literature, these modifications have had radical effects on the stability and efficiency properties of the models. We will analyze the effect of adding spatial structure to two communication games: the Lewis Sender-Receiver game and a modified Stag Hunt game. For the Stag Hunt, we find that the results depart strikingly from previous models. In all cases, the departures increase the explanatory value of the models for social phenomena.

1. Introduction. In seeking to explain the emergence and maintenance of cooperative social norms, philosophers, biologists, and economists regularly turn to game theory. Often the explanations and predictions provided by traditional equilibrium analysis are not adequate to account for the practices we observe. As a result, scholars are increasingly using dynamic models. It is common to find the replicator dynamics, a simple model of large population change, used in the philosophical and biological literature. Economists often use other similar models called myopic best response or probabilistic best response. Unfortunately, these dynamic models have been unable to provide an adequate explanation for many observed social behaviors. Even when these dynamics provide excellent explanations, we might worry that these explanations are peculiar to the particular dynamics and thus are not general explanations.

*Received November 2003; revised March 2004.

[†]To contact the author, please write to: Department of Logic and Philosophy of Science, 3151 Social Science Plaza A, University of California, Irvine, CA 92697-5100; e-mail: kzollman@uci.edu.

[‡]My gratitude is owed to Brian Skyrms, Jason McKenzie Alexander, Sam Hillier, the anonymous referees, and the participants in the Dynamics Seminar at UCI for their assistance. Generous financial support was provided by the School of Social Sciences and the Institute for Mathematical Behavioral Sciences at UCI.

Philosophy of Science, 72 (January 2005) pp. 69–85. 0031-8248/2005/7201-0004\$10.00
Copyright 2005 by the Philosophy of Science Association. All rights reserved.

In the search for an adequate explanation for cooperative social behavior, scholars are finding new ways to add analogues of animal social life to their models. If the addition increases the propensity for cooperative behavior in the model, then we have a potential explanation for the existence of such behavior in animal populations. For instance, some scholars are accommodating communication or spatial arrangement in their models. Although these are not the only relevant features of animal social life, both are readily present and are easily fit into the more general dynamic models. If there is a general lesson to be drawn from the existing literature on these two modifications it is that cooperation is rarely harmed and sometimes helped by increasing the reality of the model. Of course, this general lesson radically oversimplifies matters. Often in adding these two complexities, new pitfalls for the evolution of cooperation are introduced. Occasionally the pitfalls are substantial enough to make cooperation less likely than in simpler models.

Although both communication and spatial arrangement have been studied in great detail individually, both modifications have not been studied in combination. Studying the combined models provides us with two opportunities to improve our explanations of social cooperation. First, we can discover if increasing the complexity of the model adds explanatory value. It might be the case that the more complex models behave in ways analogous to the simpler models. If we discover this, we have some reason to think that the simpler models provide the best possible explanation for social cooperation without the addition of new, as of yet unstudied, features in animal populations. On the other hand, new features might emerge and provide new analogues for other types of social behavior that could not be explained by simpler models. Second, we can assess our general lesson that the probability of cooperation increases as our model becomes more realistic. It would be a mistake to uncritically extend the general lesson, expecting that the combination would simply result in even more cooperation. It is conceivable that the modifications might interfere with each other, making cooperation less likely than in the simpler models. This would of course be a disaster for those interested in modeling animal cooperation, since both features are present simultaneously in cooperating populations.

As a result, combining two previously studied modifications is of substantial importance to those interested in providing evolutionary explanations for social behavior. In order to begin this study, we will examine two games: the Sender-Receiver game and Stag Hunt. Both games will be embedded in a spatial network where each player is restricted to interacting only with her neighbors. In the case of the Stag Hunt, players will also be allowed to send signals to their partners prior to choosing an action in the base game. In both cases we discover that the richer model

provides equally good or better explanations for the evolution of social behavior than prior models. The results for the Stag Hunt are striking because they depart radically from prior analysis, demonstrating that this avenue of study has not been exhausted by analyzing the individual modifications independently.

In both of these games we will look at two types of cooperation. In each game, some strategies are cooperative strategies, i.e., all players would do better if they all played one of those strategies collectively. Achieving cooperation is not as easy as it might seem; in these games there is either more than one cooperative strategy or achieving the cooperative payoff requires a risk on the part of the individual players. Investigating this type of cooperation will be our primary focus. Since both games involve signals sent between two players, we will also analyze the status of meaning of those signals. In asking if a signal has meaning, we are simply asking whether, upon reception of the signal, a player can be said to have any more information about either an unobserved state of the world or about his opponent than he had prior to receiving that signal.¹ Formally, this is a question about the probabilistic dependence between the signals and other unobserved information. In the case of the Stag Hunt we are interested in determining the effect that signals have on the population and if they have any unique effect on the outcomes we observe.

2. The Sender-Receiver Game. In *Convention* (1969), David Lewis suggests that meaning can be completely explained by describing the communicators as using a strategy in a repeated cooperation game. Consider for example a simple situation of communication. A person, ‘the sender’, has some private information about the state of the world. Her partner, ‘the receiver’, must choose an action. If the receiver acts in accordance with the state of the world, then both players receive a positive payoff, otherwise they get nothing. The sender has at her disposal a set of signals which she can employ in coordinating the receiver’s actions with the state of the world. To simplify the case, suppose there are two acts, signals, and states of the world, and each state of the world is equally likely to occur. In this case there are two equilibria which achieve the highest possible payoff:² one where signal 1 is only sent in state one and signal 2 is sent in state two and a second where the signals are inverted. There

1. For a discussion of the philosophical import of analyzing this feature of meaning see Grim et al. 2001.

2. Here we are referring to pure strategy Nash Equilibrium (Nash 1950). A set of strategy choices is a Nash Equilibrium if and only if no player can do better by unilaterally changing her strategy.

are also babbling equilibria. Babbling equilibria are composed of players who send the same signal regardless of the state of the world or who take the same act given either signal. Since the signals do not contribute to correct action in these equilibria, the players do not receive the highest payoff. However, since these states are Nash Equilibria, no player can improve by unilaterally changing her strategy.

Unfortunately, the assumptions employed by standard equilibrium analysis require such extensive cognitive capabilities that one might wonder if humans are capable of such complexity (much less animals with lower cognitive abilities). While rational choice theory might be able to account for the emergence of signaling systems in populations of humans that already possess some language, it cannot explain the emergence of the first human signaling systems, nor can it explain the emergence of signaling systems among other living things.

Yet signaling systems *do* emerge in these populations. Lewis suggests that features such as natural salience might account for the emergence of signaling systems. Certainly this might be the case in the emergence of some signals (e.g., hieroglyphic representations), but it cannot account for all signaling systems. Crawford and Sobel (1982) suggest that signaling equilibria are focal and thus chosen by players observing them. Unfortunately, this explanation only pushes the evolutionary question back onto the psychological functions that result in particular equilibria being focal. In addition, in many signaling games it is hard to see how one signaling system is focal. Furthermore, neither Lewis' nor Crawford and Sobel's suggestions can account for the emergence of signaling in creatures with low cognitive capacities.

In an attempt to determine whether signaling systems could emerge among humans, absent focal points, or natural salience, Blume et al. (1998) undertook an experiment to determine if signaling systems would spontaneously emerge. They discovered that with signals devoid of prior meaning, small groups of human players were able to converge to one signaling system within a very short time (under 20 plays). So, how can we explain this emergence?

Using the tools of evolutionary game theory, Skyrms (1996) attempts to provide this explanation. Here, Skyrms makes no assumptions about the rationality or knowledge of the players. Each player is described entirely by a strategy choice which involves a choice of a unique signal in each state and a unique act given a signal. After playing as both the sender and receiver against a random opponent, the players then reproduce proportional to their payoff relative to other players in the population. This rule governing the population's transition is commonly called the replicator dynamics. Here we discover, in accordance with our observations about the natural world, that signaling systems evolve from a vast majority

of population starting points. In fact, signaling systems are the *only* evolutionarily stable strategies—they cannot be invaded by nonsignalers and they will invade any other population. Despite the loss of rational actors and complete knowledge, we still seem entitled to use Lewis' signaling systems as an explanation of the emergence and continuance of simple languages in a vast array of contexts.

One might object that Skyrms has merely abandoned one set of implausible assumptions in favor of another. The replicator dynamics makes a series of assumptions regarding the type of population and interactions unlikely to occur in the natural world. In order to account for this worry, we might switch and use a different model where players are restricted to playing only with a few people. If by making this switch we discover radically different results, we might question the applicability of Skyrms' results. There are many possible kinds of spatial models; here we will use one of them which allows for the best analysis of signaling games.

Our new model involves a group of 10,000 players arranged in a square. Each player plays the game against each of eight neighbors (known as the Moore-8 neighborhood).³ In order to avoid problems that might arise for players on the edges, the square is mapped onto a torus so every player has eight neighbors. In each round every player acts both as a sender and receiver with each neighbor. After playing with every neighbor, each player chooses a new strategy based on the performance of his neighbors. If a neighbor has done better than him, he switches to that neighbor's strategy.⁴ If more than one strategy has done better than him, but equally well as each other, he chooses one of these strategies at random. To simplify the strategy choice, we will define a strategy as both a sender strategy and a receiver strategy.

If we randomly assign each player to one of the sixteen available strategies (four sender strategies combined with four receiver strategies), do we still observe the emergence of signaling systems in this new model? Indeed we do, but in this case we almost always observe the presence of *both* possible signaling systems. An example of this is illustrated in Figure 1,⁵ where white represents one signaling system and black indicates another.

3. A model similar this has been used to study a Prisoner's Dilemma like game by Nowak and May (1993) and Nowak, Bonhoeffer, and May (1994) and a repeated Prisoner's Dilemma by Grim, Mar, and St. Denis (1998). This spatial model was applied to bargaining games by Alexander and Skyrms (1999).

4. We can interpret this change either as a player imitating the actions taken by another who does better or that unsuccessful players die off and are replaced by more successful neighbors.

5. Both figures were generated with Jason McKenzie Alexander's *Evolutionary Modeling Lab* software.



Figure 1. Spatial Sender-Receiver.

In Skyrms' model using the replicator dynamics, it was possible for both signaling systems to be present in an equilibrium. If each of the two signaling systems composed exactly half of the population then the population would remain in that state. But in Skyrms' model this state is highly unstable; if either signaling system gets the smallest gain over the other, it will quickly take over the population. In the spatial model, this state is *neutrally stable*, which is to say that although some possible mutations will not be eliminated, they will not cause large scale changes in the structure of the population.⁶ The largest change caused by a mutation would be the elimination of one of the smaller pockets that is surrounded by another signaling system.

In a random sample of 10,000 possible starting points *all* populations evolved to a state containing only signaling systems. However, it is important to note that there does exist a small proportion of starting points which will result in populations that are not engaged in a signaling system.

6. Even if we relax our dynamics slightly and allow players to switch if a neighbor does equally as well as he, this state remains stable since every border player sees an interior player of his type but no (completely) interior player of the other type. As a result, the border player will constantly imitate his interior neighbor.

These populations are either entirely composed of babblers or each signaler in the population is completely surrounded by babblers. Babblers come in two varieties. One employs half of a signaling system—when they are the sender they send a different signal in each state of the world, but always take the same action (or vice versa). The other variety consistently babbles—they send the same signal in both states and take the same action regardless of the signal they receive. All of these states can be invaded by signalers. If the population is composed of babblers of the first type, then a single mutation to the appropriate signaling system will quickly lead the population to a state composed entirely of that signaling system. On the other hand, if the population is composed of the second type, a successful invasion requires that two signalers be neighbors. This is not as unlikely as it might seem; since a single mutation will not be eliminated in a population composed of babblers, it will do no worse than its neighbors. Our mutant need only wait for one of her neighbors to switch. Populations composed of both types of babblers may also be invaded by either one or two mutations.

As suggested earlier, one might worry that these results represent an artifact of the dynamics and not a more general phenomenon. Happily, there are two similar results in the literature using either different games or different dynamics. Berninghaus and Schwalbe (1996) analyze a class of games known as cooperative games which are very similar to ours. In their model the players interact with four neighbors (N, S, E, W) and determine their strategy based on myopic best-reply. Myopic best-reply is also a deterministic dynamic for strategy choice; however with best-reply, each player chooses the strategy that would provide the highest payoff against the strategies used by her opponents on the previous round. Despite the prevalence of myopic best-reply in the literature, we choose not to use it. As an explanation of simple animal behavior, myopic best-reply is overly complex. Bacteria are not capable of determining even the myopic best-reply. On the other hand, as an explanation of human behavior, this dynamic is overly simple. It does not allow for evaluation of the future nor does it allow one to choose different strategies against different opponents. As alluded to above, the dispute is purely academic since Berninghaus and Schwalbe also observe that two conventions can coexist in spatial populations.

Grim et al. (2001) analyze a slightly different game using our imitation dynamics. The most notable difference in their game is that the sender does not receive any payoff for successful communication. Grim et al. also find similar results in their game. Two signaling systems quickly take over and peacefully coexist with each other. They also consider another dynamic where the players ‘breed’ with their most successful neighbor

TABLE 1.

	Stag	Hare
Stag	(4, 4)	(0, 3)
Hare	(3, 0)	(3, 3)

rather than merely imitating them. In this dynamic two signaling systems emerge but one eventually completely takes over the population.

In these models the signals' meaning must be analyzed differently than it was in the replicator dynamic model. If the population is taken as a whole, neither signal provides much information about the state of the world. Receiving signal one from a random member of the population leaves one with only the smallest increase in information about the state of the world than one had prior to the signal.⁷ Leaving the analysis of meaning at that would be to miss something important about the population. Our dynamics almost always result in a population where there are no babblers; many players perfectly coordinate with their neighbors. Notice that if you are aware of the region from which the signaler was chosen, you now have perfect information about the state of the world. So each signal can be said to have acquired regional meaning.

It appears that replacing the replicator dynamics with a more realistic dynamic has not harmed our explanation for the evolution of meaning in simple sender-receiver games. Our certainty of the evolution of meaning has not been swayed. In fact, as we increased the realism of our model, we also gained an explanation of another observed result: different signaling systems evolve in the world and the signaling system a player uses is dependent on his location in the population.

3. The Stag Hunt. In the case of the Lewis Sender-Receiver game, one is not risking much by playing a signaling system. Although one runs the risk that one's neighbors will all choose the opposite signaling system, for the most part playing a signaling system is a relatively safe endeavor. This game certainly does not correspond to all situations in which communication might arise. In fact, we often need to determine if we should take a risk in order to secure a better result based on what we expect others to do. One of these situations is well modeled in the Stag Hunt (illustrated in Table 1). Two players go out to hunt and each individual chooses independently to hunt either a stag or a hare. If they cooperate they can kill a stag, which provides the highest payoff. If a player chooses to hunt a hare, she may kill it without the aid of her partner, however

7. In the example pictured in Figure 1 the signal does give some information, since one signaling system is more prevalent than another. However, the difference is small and so too would be the information provided by the signal.

TABLE 2.

	Stag	Hare
Stag	(15, 15)	(0, 8)
Hare	(8, 0)	(7, 7)

the value of the hare is less than that of the stag. This game has two pure strategy Nash equilibria: one where both players hunt stag and one where both players hunt hare. The stag hunting equilibrium is payoff dominant because in this equilibrium both players receive higher payoffs than they would in any other equilibrium. The hare hunting equilibrium is less risky since each player is guaranteed to receive a payoff regardless of the actions of her partner.

If we analyze this game using the replicator dynamics, we discover that stag hunting will take over in less than one quarter of the starting population states. If more than three quarters of the population hunts stag, stag hunting will take over; hare hunting will take over otherwise. The two stable equilibrium states of the population are “All hunt stag” and “All hunt hare.” In order to switch from the hare hunting equilibrium to the other, over three quarters of the population must simultaneously switch from one strategy to the other. As the ratio between the payoff for stag hunting and the payoff for hare hunting becomes smaller (e.g., making Hare pay 3.5 instead of 3), the basin of attraction for stag hunting shrinks. Also, as this ratio becomes smaller it becomes easier to switch from a stag hunting equilibrium to a hare hunting equilibrium and harder to switch from hare to stag. It seems the standard replicator dynamics does not do well in explaining the evolution of cooperation in a game like the stag hunt.

This game can be modified in any number of ways. In keeping with the Sender-Receiver game, we will focus on communication and spatial structure. First, we can allow each player to send a signal to the other. Through the evolutionary process these signals might acquire meaning that would assist in the evolution of cooperation. Second, we can change the interaction structure as we did above—placing the individuals on a surface where they are constrained to interact only with those around them. Each of these modifications has been studied in some detail individually. Here we will briefly summarize the results in turn.

3.1. Stag Hunt with Communication. Aumann (1990) raises a concern about the prospect of costless preplay communication in a game like the stag hunt. If we change the simple stag hunt above to a slightly different payoff structure, pictured in Table 2, we might lose communication as a

facilitator of cooperation.⁸ Aumann argues that since both stag hunters and hare hunters would prefer their partners to hunt stag, the signals would have no meaning. Receiving a cooperative signal from one's partner would give a player no information regarding her partner's intentions since her partner might plan to hunt hare but wish to receive the higher payoff. This conjecture has received substantial treatment in the literature and it is not our aim to address its plausibility for rational choice theories here. However, it would be interesting to know if such a modification would have a negative effect on the evolution of cooperation.

Skyrms (2002) analyzed this stag hunt with signals devoid of prior meaning using the replicator dynamics.⁹ Each player had two signals which she could send to her partner, she could then choose an action which could be dependent on the signal sent by her opponent. In a random sample of 10,000 starting points 75% stabilized with all players hunting stag. The remaining populations were divided between two population states: 13% ended with all players hunting hare and 11% ended in a new polymorphic equilibrium which will be discussed in more detail below.¹⁰ So it appears that contrary to Aumann's conjecture, signals can provide a substantial catalyst for cooperative behavior.

3.2. Spatial Stag Hunt. Ellison (1993) analyzed the Stag Hunt on a circle, where each player plays with his two neighbors. Players update their strategies based on the best-reply dynamics used by Berninghaus and Schwalb. Most populations will quickly converge to all hunting hare. It would require almost miraculous mutations for the population to escape this state. As previously noted, one might worry about the best-reply dynamics as providing adequate explanations for the behavior of any creature. In addition, one might worry about the circle—since it allows for only one type of border between players, it might prohibit certain types of coexistence that would be possible in differently arranged populations.

In order to address the second worry, Lee and Valentinyi (2000) studied the best-reply dynamics on a two dimensional lattice. As in Berninghaus

8. In fact, this payoff structure is even more amenable to cooperative play than the one suggested by Aumann. Aumann's Stag Hunt involves higher payoffs for hare hunting (while maintaining the ordinal structure of the payoffs). We will call any Stag Hunt where a hare hunter would prefer his partner hunt stag an Aumann Stag Hunt.

9. Skyrms presents his results for the Stag Hunt pictured in Table 2. Since this game is the least hostile to cooperative play among possible Aumann Stag Hunts, the results were most striking. He also analyzed several other payoff structures and achieved similar results.

10. The missing 1% is due to rounding error; for the exact numbers see Skyrms 2002.

and Schwalbe (1996), players were constrained to interact only with their nearest four neighbors. Lee and Valentinyi analyzed a model where there were no mutations. They analyzed the limit behavior of the population as the size of the lattice grew to infinity. Even without mutations, hare hunting would take over the population with probability one.

Concern over the former of the two worries led Skyrms (2004) to use the same spatial model used in the Sender-Receiver game to analyze the Stag Hunt. In simulation he finds that stag hunting takes over in 99% of the starting populations he sampled. In order to determine the culprit in Ellison's model, Skyrms also considers populations placed on a circle where the players update based on imitate-the-best dynamics. Here we find that neither hare hunting nor stag hunting is contagious. If there is any group of three or more stag hunters surrounded by hare hunters they will remain, since the edge players see an interior player who successfully hunts stag with each of his neighbors. The border hare hunters will not switch since they do better than the border stag hunters (the only stag hunters they see). If we allow the players to see beyond their interaction neighborhood (so a player might imitate either his neighbors or his neighbors' neighbors), then the border hare hunter will see a successful stag hunter and stag hunting will take over.

It is hard to draw a general lesson from all of these different interactions. Stag hunting does well in two-dimensional populations with imitate-the-best dynamics and one-dimensional populations with large imitation neighborhoods. Stag hunting can survive while not invading in one-dimensional populations using imitate-the-best dynamics. And stag hunting is worse off with best-reply dynamics in spatial populations than it was with any dynamic in a random encounter model.

So how do we resolve these divergent results? In spatial games it seems that although some dynamics help stag hunting, others harm it. If we add communication instead of spatial structure, stag hunting seems to be helped somewhat although not as much as in some spatial models. In order to resolve this question, we can look at one model where both spatial structure and communication are present. We will use the same communication structure used by Skyrms (2002) and the same dynamics we used above to analyze the Sender-Receiver game.

3.3. Spatial Stag Hunt with Communication. Following Skyrms (2002) we will represent each strategy with a 3-tuple: ⟨Signal to send, Act to take if signal 1, Act to take if signal 2⟩. We have eight strategies, four types for each signal. For the sake of simplicity, we will name each of the four strategies. *Stag Hunters* will hunt stag with players who send either signal, *Nationalists* will hunt stag only with those who send the same signal they sent their partner, *Individualists* will hunt stag only with a player who sent

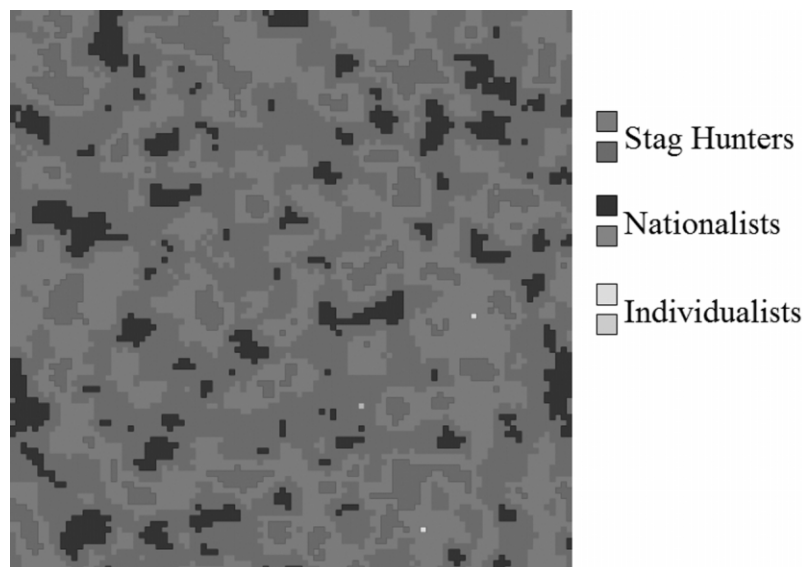


Figure 2. Spatial Communication Stag Hunt.

the opposite signal, and *Hare Hunters* always hunt hare. As with the Sender-Receiver game, we will place the players in a 100 by 100 lattice and use the imitate-the-best dynamics.

Again starting the population at a random starting point and allowing it to evolve until it reaches a stable state, we find that almost all populations evolve to a state like the one pictured in Figure 2. Here we have *six* of the eight strategies peacefully coexisting. Despite the presence of so many strategies, everyone is hunting stag! Nationalists are surrounded by Stag Hunters who send the same signal. Individualists are surrounded by Stag Hunters who send the opposite signal. And the only prolonged borders between two strategies who send different signals are between the two types of Stag Hunters.

Unlike the Sender-Receiver game, the stability of this state depends on inertia built into the model. Remember, a player will only switch his strategy if a neighbor has done *better* than him. Since in Figure 2 all players are hunting stag with all of their neighbors, all players are receiving the same payoff. Should we relax the inertia assumption and allow players to switch to an equally good strategy with some probability, the state will regularly change. The Nationalists will disperse themselves throughout the island of Stag Hunters in which they are contained. However, the Nationalists will always be surrounded by Stag Hunters who send the

same signal. One should not expect this to result in any one strategy invading, since each border player is equally likely to switch to a neighbor's strategy. However, the state will not be stable in the same way.

Although many starting points evolved to states where the population did not completely hunt stag, stag hunting was very prevalent (average proportion of the players hunting stag was greater than .999 and no population hunted less than 98.9% stag in 10,000 trials). In a significant number of the cases the population did not reach a stable state in 10,000 generations (approximately 13%). But even in those populations, most players were hunting stag in the final generation. These results were not substantially altered by increasing the return for hare hunting (without surpassing the payoff for stag-stag),¹¹ by modifying the mutation rate, or by changing the structure to an Aumann's Stag Hunt.¹² Aumann Stag Hunts took on average three times longer to reach stability (among those populations that eventually stabilized), but they were mostly hunting stag when they did.

This is not the only good news for social cooperation. As was the case in Skyrms' 2002 article, any population of Hare Hunters with an unused signal can be invaded by two neighboring Nationalists using the extra signal as a "secret handshake."¹³ A Nationalist who uses the other signal will hunt hare against all of her Hare Hunter neighbors and so will do no worse. If one of her neighbors should switch to her strategy, they will now cooperate only with each other and as a result achieve a higher payoff than any of their neighbors. After a sufficient number of generations, the entire population will convert to their strategy. This new population is neutrally stable.

Unlike the simple communication game, the secret handshake can be used even in populations with both signals. Consider the following population: the only two strategies are the Hare Hunters and one agent has eight neighbors who all send the same signal. This agent can now use the

11. The results reported here were for a borderline Stag Hunt where stag-stag paid 2 and hare hunting paid 1. This situation is the most conducive to cooperation since a player only needs half of her neighbors to hunt stag to make stag hunting as profitable as hare hunting. Although the average number of stag hunters decreased as the payoffs grew closer together, it did not drop below 96% until hare hunting paid 7/8 that of stag hunting. It is unsurprising that this ratio would be important since each player has eight neighbors.

12. In the standard Stag Hunt 99.86% of the population hunted stag at the end of the run (stability or 10,000 generations). In the Aumann Stag Hunt, with payoffs 4, 3, and 2, 99.51% hunted stag at the end. As with the standard Stag Hunt as the return for hare hunting grew, the proportion of hare hunters grew, but only by a relatively small amount.

13. The notion of a secret handshake is due to Robson (1990).

other signal as a secret handshake. She need only wait for one of her neighbors to mutate to a cooperative strategy, and in one generation they will convert the three-by-three square. Now if they encounter Hare Hunters who use the other signal, the signal is no longer a secret handshake. This is of no consequence because there are enough of them to overcome the harm of the Hare Hunters. Since most starting populations composed entirely of Hare Hunters will have one player entirely surrounded by one strategy, we should expect most Hare Hunting populations will be invaded by cooperative strategies.¹⁴

This invasion is only possible if the payoff for hare hunting is less than $5/8$ s the payoff for stag hunting. Otherwise the three-by-three square of stag hunters simply remains at that size. However, if we enlarge the imitation neighborhood, so that each player will imitate-the-best neighbor out of the nearest twenty-four or more players, *all* payoff structures can be invaded (so long as the game is still a Stag Hunt).¹⁵ This invasion is possible since a border Hare Hunter could see an interior Stag Hunter who perfectly coordinates with her neighbors. Since every Stag Hunt game involves a higher payoff for Stag Hunting, this interior Stag Hunter will receive a higher payoff than any Hare Hunter.

Skyrms discovered that signaling introduced a new polymorphic equilibrium into his replicator dynamic model. In this equilibrium the population is evenly divided between the two types of Individualists. Here each half of the population will hunt stag with the other half, but not with themselves! This seems the most counterintuitive result of this model since it is rare to find groups that refuse to cooperate with their own kind but are willing to cooperate with others. In the replicator dynamic model, this state is completely stable, any mutation will be eliminated by the dynamics. In our spatial model, however, this population can now be invaded. The state itself takes on some strange properties when placed in a spatial setting. The population is never stable—most players switch to

14. We can look at this game on a circle and we find similar results to the ones we observed with the regular spatial Stag Hunt on a circle. Groups of three or more Hare Hunters who all send the same signal can be invaded by Nationalists using the other signal as a secret handshake. Unless the imitation neighborhood is large, then the invasion only stops with each of those groups.

15. This requirement is even more strict if we are considering an Aumann stag hunt. There the payoff for hunting stag must be greater than the payoff for hunting hare against a Hare Hunter plus $3/5$ ths of the payoff for hunting hare against a Stag Hunter. But this condition can also be eliminated by enlarging the imitation neighborhood.

playing the opposite strategy in each generation.¹⁶ However, if a player mutates to a Stag Hunter and his neighbors will cooperate with him, then he will invade the population. In simulation, with an average of one mutation per generation, every population starting with only Individualists resulted in a state with all players hunting stag. Convergence to this state was relatively fast, occurring on average in 107 generations.

One might wonder how much work the signals are doing in our new simulation. After all, in the spatial Stag Hunt using the imitation dynamics, stag hunting took over in 99% of the trials.¹⁷ Perhaps the signals are merely ‘staying out of the way’ while stag hunting takes over. Even if that were the case, it would have been an important result that the signals did stay out of the way. One might have expected that the spatial signal game would look more like the replicator dynamic game than the spatial Stag Hunt game. But in fact, signals have contributed something. In the spatial Stag Hunt without communication, a population of all hare hunters requires at least six simultaneous mutations in order to be invaded by stag hunters.¹⁸ In the spatial Stag Hunt with communication, only two non-simultaneous mutations are required to invade all hare hunting populations.

In addition, one might look at the information conveyed by the signals in the most common end population (e.g., the one pictured in Figure 2). In that population, the signals have more global meaning than they had initially. Each signal corresponds to one of three types, not four. And the probability that a signal 1 sender is a Stag Hunter or a Nationalist is very high. In addition, signals provide strong evidence of a player’s propensity to cooperate with another who sends the same signal.¹⁹ Finally, the signals have acquired a similar type of regional meaning we observed in the spatial

16. The switching is caused by corner players who have more Individualists who send the opposite signal surrounding them. Since this corner player does well (he successfully cooperates with most of his neighbors), his neighbors imitate him on the next round. But now, the benefit is lost since he is surrounded by others that send the same signal.

17. This is comparing our game to the Spatial Stag Hunt using imitation dynamics. It would be nice to compare our results to other dynamics like myopic best response. Unfortunately, best response dynamics are not well defined for communication games like these. For any strategy there is a minimum of two best responses. So, the dynamics becomes much more complex. Either the entire population chooses a best response at random and is constantly changing, or we make a particular strategy the default and run the risk of tainting the results.

18. This number depends on the specifics of the payoff matrix, but must be at least six if the game is a Stag Hunt.

19. This is not perfect information since a few Individualists have survived.

Sender-Receiver game. In different regions, the signal conveys perfect information about a player's type.²⁰

4. Conclusion. What of our two concerns that motivated us to begin this investigation? Our first concern was to discover if the combination of two previously studied modifications (i.e., signaling and spatial arrangement) would result in any new results or if it would just be more of the same. We found that things were much different in the combined models than they were in the simpler models. In the case of the Sender-Receiver game, we found that we had a new explanation for the emergence of different signaling systems in the same population; this could not be explained by the simpler model. In the Stag Hunt, we observed a radically different population than could even be conceived in either of the simpler models present in the literature.

Our second endeavor was to test the general lesson—that cooperation is helped by increasing the complexity of our model. The answer to this question is more complicated. At the outset we defined two types of cooperation. The first type of cooperation is achievement of the pareto optimum payoff (i.e., the payoff which benefited both players the most). Here we found the general lesson was confirmed. In the case of the Sender-Receiver game, most players continued to achieve the highest payoff by coordinating with all of their neighbors. Cooperation did well in the Stag Hunt; a noncooperative equilibrium (composed of individualists) was eliminated from the simple communication game and stable populations in the communication game could now be invaded. Our second type of cooperation is the achievement of meaning. Here we found that population-wide meaning was harmed by the addition of spatial structure. However, we were able to save a weaker version of meaning by relativizing meaning to regions in our model. Overall, the general lesson seems to be confirmed: increasing the reality of the model does assist in the evolution of cooperative behavior in the Sender-Receiver and Stag Hunt games. However, it seems that the simpler models have not said all there is to say about cooperation and spatial structure as explanations of social cooperation.

REFERENCES

- Alexander, J. McKenzie, and Brian Skyrms (1999), "Bargaining with Neighbors: Is Justice Contagious?", *Journal of Philosophy* 96 (11): 588–598.
 Aumann, Robert J. (1990), "Nash Equilibria Are Not Self-Enforcing", in Jean Jaskold Gabzewicz, Jean François Richard, and Laurence A. Wolsey (eds.), *Economic Decision*

20. If we relax the inertia assumption as suggested earlier, some of the regional meaning will be lost.

- Making: Games, Econometrics and Optimization*. Amsterdam: North Holland Press, 201–206.
- Berninghaus, Siegfried K., and Ulrich Schwalbe (1996), “Conventions, Local Interaction, and Automata Networks”, *Journal of Evolutionary Economics* 6: 297–312.
- Blume, Andreas, Douglas V. DeJong, Yong-Gwan Kim, and Geoffrey B. Sprinkle (1998), “Experimental Evidence on the Evolution of Meaning of Messages in Sender-Receiver Games”, *American Economic Review* 88 (5): 1323–1340.
- Crawford, Vincent P., and Joel Sobel (1982), “Strategic Information Transmission”, *Econometrica* 50 (6): 1431–1451.
- Ellison, Glen (1993), “Learning, Local Interaction, and Coordination”, *Econometrica* 61 (5): 1047–1071.
- Grim, Patrick, Trina Kokalis, Ali Alai-Tafti, Nicholas Kilb, and Paul St. Denis (2001), *Making Meaning Happen*, Research Report #01–02, Group for Logic and Formal Semantics, SUNY, Stony Brook.
- Grim, Patrick, Gary Mar, and Paul St. Denis (1998), *The Philosophical Computer*. Cambridge, MA: MIT Press.
- Lee, In Ho, and Ákos Valentinyi (2000), “Noisy Contagion without Mutation”, *Review of Economic Studies* 67: 47–56.
- Lewis, David (1969), *Convention: A Philosophical Study*. Oxford: Blackwell.
- Nash, John F. (1950), “Equilibrium Points in n-Person Games”, *Proceedings of the National Academy of Sciences of the United States of America* 36 (1): 48–49.
- Nowak, Martin A., Sebastian Bonhoeffer, and Robert M. May (1994), “More Spatial Games”, *International Journal of Bifurcation and Chaos* 4 (1): 33–56.
- Nowak, Martin A., and Robert M. May (1993), “The Spatial Dilemmas of Evolution”, *International Journal of Bifurcation and Chaos* 3 (1): 35–78.
- Robson, Arthur J. (1990), “Efficiency in Evolutionary Games: Darwin, Nash, and the Secret Handshake”, *Journal of Theoretical Biology* 144: 379–396.
- Skyrms, Brian (1996), *Evolution of the Social Contract*. New York: Cambridge University Press.
- (2002), “Signals, Evolution and the Explanatory Power of Transient Information”, *Philosophy of Science* 69: 407–428.
- (2004), *The Stag Hunt and the Evolution of Social Structure*. New York: Cambridge University Press.