

The Effects of Finite Populations and Selection on the Emergence of Signaling

Kyle I. Harrington, A. Pinar Ozisik, Jordan B. Pollack

DEMO Lab, Computer Science Department, Brandeis University, Waltham, MA 02454

Abstract

In the research described here we examine the emergence of signaling from non-communicative origins, using the Sir Philip Sidney Game as a framework for our analysis. This game is known to exhibit a number of interesting dynamics. In our study, we quantify the difficulty of reaching multiple types of equilibria from initially non-communicative populations with an infinite population model. We then compare the ability of finite populations with typical tournament selection to approximate the behaviors observed in infinite populations. Our findings suggest that honest signaling equilibria are difficult to reach from non-communicative origins. In the second part of the paper, we show that the finite model fails to model dynamics that permit deceptive signaling under typical evolutionary conditions, where infinite populations exhibit spiraling behavior between honest and deceptive signaling.

Introduction

Communication and expression in man and animals has allowed for the formation of complex social organizations. Although sophisticated forms of communication have emerged, such as human language, the origin of animal communication is rooted in the exchange of simple signals. These signals have coevolved between senders and receivers for the communication of attributes such as need, status, and intention. We study a simple signaling game that allows us to address some shortcomings of previous studies on the emergence of signaling. We quantify the difficulty of evolving honest signaling systems, and the failure of some finite models to permit deceptive dynamics. These observations are a step towards understanding how a rational agent could respond to a signal with Sir Philip Sidney's immortal words: "thy necessity is yet greater than mine."

The coevolution of signaling has attracted attention since the inception of the field of artificial life and even earlier in studies of animal behavior and ethology. Evolutionary computation researchers studying the origins of signaling generally employ evolutionary algorithms (EA) in their models, while game theorists employ population dynamics models and analytical tools. Some EA work has used game-theoretic analysis to constrain parameters (Bullock, 1997);

however, we are not aware of any studies of coevolved signals that relate continuous population dynamics to EA dynamics for signaling games. We investigate this relationship and focus on similar EA configurations to those used in previous studies of the emergence of signaling. In particular, we consider the discrepancy between the dynamics of finite population EA's and continuous population dynamics.

The relationship between continuous and finite population evolutionary dynamics has been a contentious topic (Fogel and Fogel, 1995; Ficici et al., 2005; Ficici, 2006; Ficici and Pollack, 2007; Nowak et al., 2004). An evolutionarily stable strategy (ESS) is defined for continuous population dynamics as a strategy that cannot be invaded by a rare mutant (Maynard Smith and Price, 1973; Maynard Smith, 1982). A common question in the study of evolutionary dynamics is when finite populations can achieve an ESS. In particular, the two discoveries that inspire this study are: Best-of-group tournament selection cannot converge to polymorphic Nash equilibria (Ficici et al., 2005), and even with a good selection method, a finite population may be too small to maintain an ESS (Ficici and Pollack, 2007). In a simple signaling game we investigate both the reachability of interesting equilibria from non-communicative origins, and compare the coevolution of continuous population dynamics to finite populations under tournament selection (the most common selection method used in previous work). We find that multiple dynamics involving signaling behavior are more easily reached than the traditional signaling ESS, and one of these dynamics is poorly represented by a finite population.

Background

Zahavi introduced the idea of costly signals as handicaps which lead to reliable signals (Zahavi, 1975). This handicap principle has been used to explain how signaling attributes which would seem to be energetically expensive or superfluous for survival can be selected for, especially in sexual selection. For example, plumage like the peacock's tail signal virility and strength to a peahen because the male has honestly demonstrated that it can carry the unneeded weight of the brilliant tail. For a good web exposition on honest

Table 1: The Sir Philip Sidney game.

(a) Payoff matrix.			(b) Sender strategies.		(c) Donor strategies.	
	Donate	Keep	ID	Signaler strategies	ID	Donor strategies
Potential donor	$1 - d$	1	<i>SH</i>	signal only if healthy	<i>DQ</i>	donate only if no signal
Signaler			<i>ST</i>	signal only if thirsty	<i>DS</i>	donate only if signal
Thirsty	1	$1 - a$	<i>SN</i>	never signal	<i>DN</i>	never donate
Healthy	1	$1 - b$	<i>SA</i>	always signal	<i>DA</i>	always donate
$m = p(\text{thirsty}) = 0.5$	Signal cost = c					

signaling, see (Bergstrom, 2012). This work was later given a rigorous mathematical treatment in (Grafen, 1990) for signals of a continuous range of quality. Later, two simple discrete signaling games were developed: the Sir Philip Sidney game (Maynard Smith, 1991) and the discrete action-response game (Hurd, 1995). The former can be seen as a generalization of the latter, which is a deliberately minimal signaling game. The discrete action-response game is based upon the handicap principle, thus models costly signaling. As the Sir Philip Sidney game is the subject of our study, we will introduce it in greater detail later.

Bullock analytically evaluated the discrete action-response game for parameters that should lead to the emergence of signaling, then used an EA to evolve a finite population (Bullock, 1997). Agents take turns playing an iterated signaling game, and are selected for reproduction using spatial tournaments. The results demonstrate a number of dynamics ranging from evolutionarily stable strategies (ESS) to cycles. It is found that the emergence of honest signaling from a non-communicative state only occurs from a subset of the analytically determined cooperative parameters.

Noble studied a version of the discrete action-response game where only one of the signaler states results in positive payoff for the sender and receiver (Noble, 1999). The criteria for the honest signaling ESS was shown to be when the payoff for signaling is greater than the cost of signaling, and the payoff of responding is greater than the cost of responding. Noble suggested that a signaling game must permit imperfect information, deception, and manipulation to allow for information transmission. While all of these points are present in the described game, we note that the ambivalence of signalers to transmit a signal in one of the two possible states means there is no incentive for deception. We demonstrate situations where an incentive to signal from both signaler states has significant implications on the coevolutionary dynamics of signaling.

We have previously worked on evolution of communication in a group foraging task, although without referencing the signaling literature (Saunders and Pollack, 1996). Similarly, Reggia et al. investigate conditions that enable the emergence of signaling (Reggia et al., 2001). In this work the authors use a 2D simulated world where agent behavior is governed by a finite-state machine, and signaling ability is

encoded in the genome. Agents have an energetic cost of living, and independent experiments are performed for predator signaling, food signaling, and environments where both types of signaling are possible. Their EA operates on a population of size 200 and multiple forms of tournament selection are compared. It is particularly interesting that smaller tournament sizes and spatially-constrained tournaments lead to more signaling. While the authors describe a set of conditions that enable signaling for their world/agent architecture, in this study we will investigate conditions that enable signaling in a simplified environment.

A review of the evolution of signaling systems is beyond the scope of this paper. For an extensive review of studies on simulating the emergence of communication, including signaling, see (Wagner et al., 2003).

Sir Philip Sidney Game

The Sir Philip Sidney (SPS) game was developed by John Maynard Smith as a model of costly signals (Maynard Smith, 1991). It is an extensive form game between two players. The importance of costly signals is based upon Zahavi’s handicap principle (Zahavi, 1975) which states that reliable signals are costly with respect to the signaler’s ecological context. This cost is explicitly introduced as a fitness penalty in the SPS game.

The SPS game is played for a single round between two players: a signaler and a donor. The signaler may be in one of two states: *thirsty* or *healthy*. The probability of the signaler being *thirsty* is m . A *thirsty* signaler has a fitness of $(1 - a)$, and a *healthy* signaler has a fitness of $(1 - b)$. In all cases $a > b$. The strategy of the signaler specifies whether it signals in either, both, or neither states. It costs the signaler c to transmit a signal. In response to receiving a signal the donor decides whether or not to donate to the signaler. Donation comes at a cost, d , to the donor, but heals the signaler to a fitness of 1. Furthermore, a globally-fixed relatedness term, r , is introduced which accounts for the opponent in the inclusive fitness of each player. Labels for signaler and donor strategies are listed in Tables 1(b) and 1(c), respectively. For example, in a game between *ST* and *DS*, if *thirsty* the signaler will transmit a signal and in response the donor donates. The signaler’s fitness is $(1 - c + r(1 - d))$ and the donor’s fitness is $(1 - d + r(1 - c))$. If the game

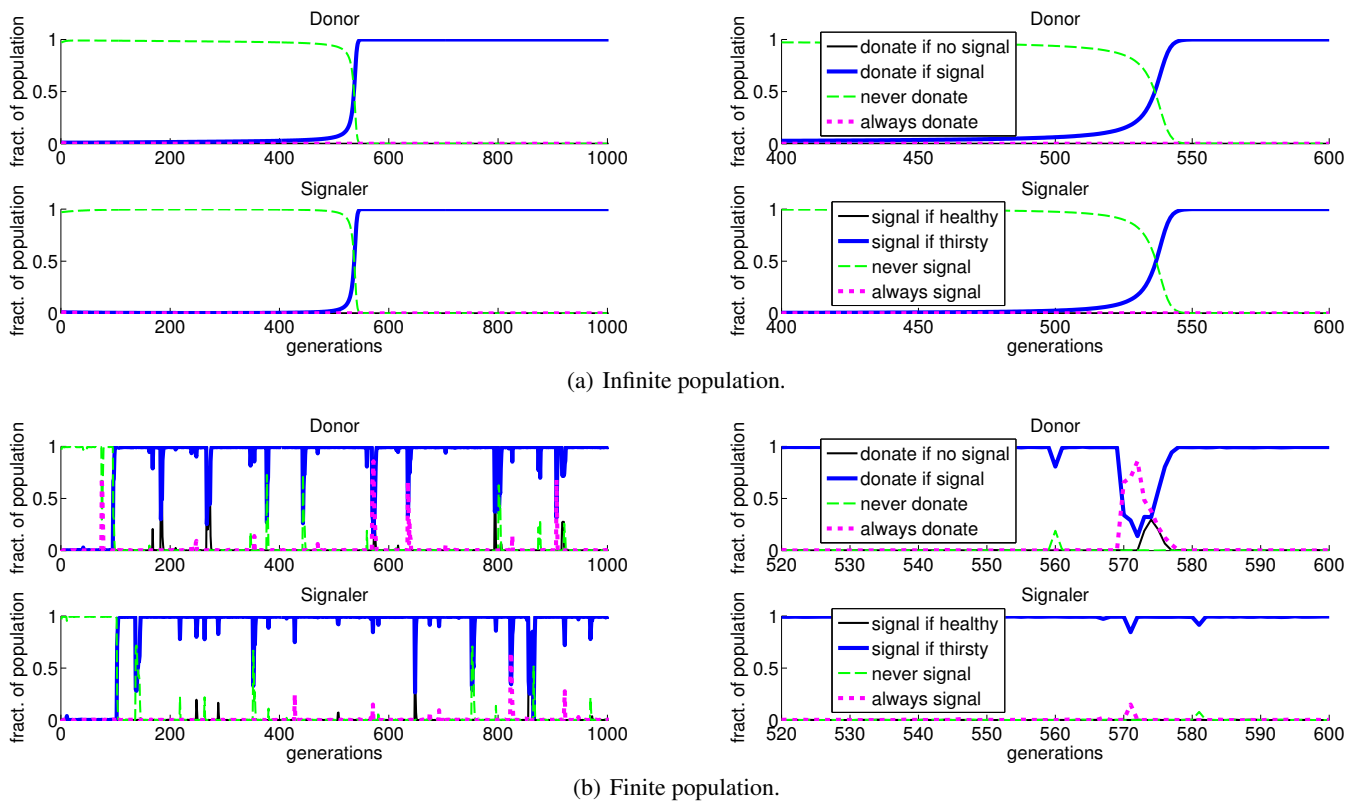


Figure 1: Example of an honest signaling equilibrium. The graphs on the right hand side are zoomed in versions of those on the left. In 1(b) the “always donate” strategy briefly invades the donor population. This is a possibility in the finite model if the SPS parameters are within a range where non-optimal strategies can be mistook for sampling noise.

was played between *ST* and *DQ*, then if *thirsty* the signaler transmits a signal and the donor does not donate. The signaler’s fitness is $(1 - a - c + r)$ and the donor’s fitness is $(1 + r(1 - a - c))$. The payoff matrix is shown in Table 1. Unless otherwise specified we set $m = 0.5$.

The SPS game has been the subject of a number of game theoretic studies, for both the discrete signaling game we study here, and the continuous-version of the SPS game (Johnstone and Grafen, 1992). The interest in this game arises from its facilities for modeling both costly signaling and signaling amongst relatives, where the latter property permits cost-free signaling in a number of conditions.

The key distinction between the discrete action-response (Hurd, 1995) and SPS games is the use of inclusive fitness (Hamilton, 1964), adding the opponent’s score weighted by a “relatedness” term, r . Relatedness accounts for the fact that if a player’s opponent is related to the player, then benefits to the opponent also benefit the player. Inclusive fitness is only utilized when computing the score for a donor and signaler playing a game, as opposed to fitness sharing from genetic algorithms where related individuals in the same population share the fitness of a given niche. In (Ozisk and Harrington, 2012) it was shown that relatedness based upon

tags, unique phenotypic identifiers, destabilizes honest signaling equilibria in finite models.

Non-communicative Equilibria

In this study we are interested in the emergence of signaling from non-communicative initial conditions. While there are multiple combinations of signaler and donor strategies that do not transfer information, we will be particularly interested in the *SN* and *DN* combination of strategies, because the two populations will be initially composed of predominately *SN* and *DN* individuals. Bergstrom and Lachmann (Bergstrom and Lachmann, 1997) have shown the *SN* and *DN* pair to be a Nash equilibrium if

$$d > r(ma + (1 - m)b)$$

Huttegger and Zollman (Huttegger and Zollman, 2010) note that reversing the inequality leads to the *SN* and *DA* pair of strategies being a Nash equilibrium. They refer to these as “pooling equilibria.”

Signaling Equilibria

One of most commonly studied type of equilibria in signaling games with handicap signals is the signaling ESS, sometimes referred to as separating equilibria. In these equilibria

ria the *ST* and *DS* strategies are dominant. Bergstrom and Lachmann (Bergstrom and Lachmann, 1997) show this is a Nash equilibrium when

$$a \geq c + rd \geq b \quad \text{and} \quad a \geq d/k \geq b.$$

An example of this type of signaling equilibrium is shown in Figure 1. We will refer to this type of signaling equilibrium as the honest signaling equilibrium.

Another type of signaling equilibrium is possible where the *SH* and *DQ* strategies are dominant. Huttegger and Zollman (Huttegger and Zollman, 2010) show this is a Nash equilibrium when

$$a \geq rd - c \geq b \quad \text{and} \quad a \geq d/k \geq b.$$

In previous work on evolving communicative agents we have seen this type of strategy pattern emerge (Saunders and Pollack, 1996). We will refer to this type of signaling equilibrium as the inverse honest signaling equilibrium.

Hybrid Equilibria

A dynamic of particular interest in the SPS game is that of hybrid equilibria, whose name is taken from the economics literature. First formally presented for the SPS game in (Huttegger and Zollman, 2010), hybrid equilibria are actually a family of polymorphic mixed Nash equilibria. In practice these hybrid equilibria can be observed in the SPS game as a spiraling phenomenon (Figure 2). The system first approaches a signaling equilibrium, such as *ST* and *DS*, and upon reaching a certain fraction of signalers and responsive donors *SA* signalers begin to take advantage of the donors. The introduction of these deceptive signalers into the population causes the *DN* strategy to increase in the donor population. As the *DN* strategy increases it becomes less favorable to signal. The *SA* strategy signals both when *thirsty* and *healthy*, as opposed to the *ST* strategy which only signals when *thirsty*, which means that the *SA* strategy has a lower fitness than *ST* when playing against the *DN* strategy, thus the *SA* strategy will be more strongly selected against. As *ST* begins to take over the signaler population the *DS* strategy also increases. Huttegger and Zollman (Huttegger and Zollman, 2010) show that the polymorphisms of the hybrid equilibria are mixed Nash equilibria given by $\lambda ST + (1 - \lambda)SA$ and $\mu DS + (1 - \mu)DN$, where

$$\lambda = \frac{r(ma + (1-m)b) - d}{(1-m)(rb - d)} \quad \text{and} \quad \mu = \frac{c}{b - kd}$$

both of which must be well-defined, and thus

$$a > d/k > b \quad \text{and} \quad b - kd > c.$$

must also be true. Furthermore, the condition

$$d > r(ma + (1 - m)b)$$

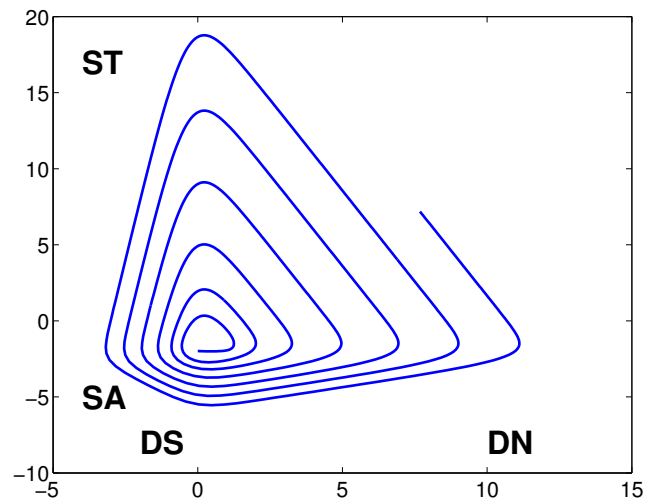


Figure 2: Example of a phase plot of strategies involved in hybrid equilibria. The evolutionary trajectory begins at the center of the spiral and moves outwards over time. X- and Y-coordinates denote the difference between the \log_{10} of the population fraction for the respective strategies.

is also required. An example of a hybrid equilibrium is shown in Figure 3. This evolutionary dynamic is reminiscent of the complex evolutionary dynamics which have been observed in continuous populations of Prisoner’s Dilemma strategies (Lindgren, 1991). However, Lindgren’s system eventually leads to an ESS, while hybrid equilibria spiral *ad infinitum* (Huttegger and Zollman, 2010).

Note that in the case of hybrid equilibria $b > 0$. This serves as an incentive for deceptive signaling, which is not a possibility in the case of Noble’s game (Noble, 1999).

Population Dynamics

We evolve infinite populations with a two-population version of the discrete-time replicator equation (Sigmund and Hofbauer, 1998)

$$x_i(t + 1) = \frac{\pi(Z(t), x_i)x_i(t)}{\sum_j \pi(Z(t), x_j)x_j(t)}$$

$$z_i(t + 1) = \frac{\pi(X(t), z_i)z_i(t)}{\sum_j \pi(X(t), z_j)z_j(t)}$$

where $x_i(t)$ is the fraction of strategy i in the first population X at time t , $\pi(P, s)$ is the payoff of strategy s against population P , and $z_i(t)$ is the fraction of strategy i in the second population Z at time t . The fitness of a particular strategy is dependent upon the strategy distribution of the other population. This assumes complete mixing and that each strategy plays each other strategy.

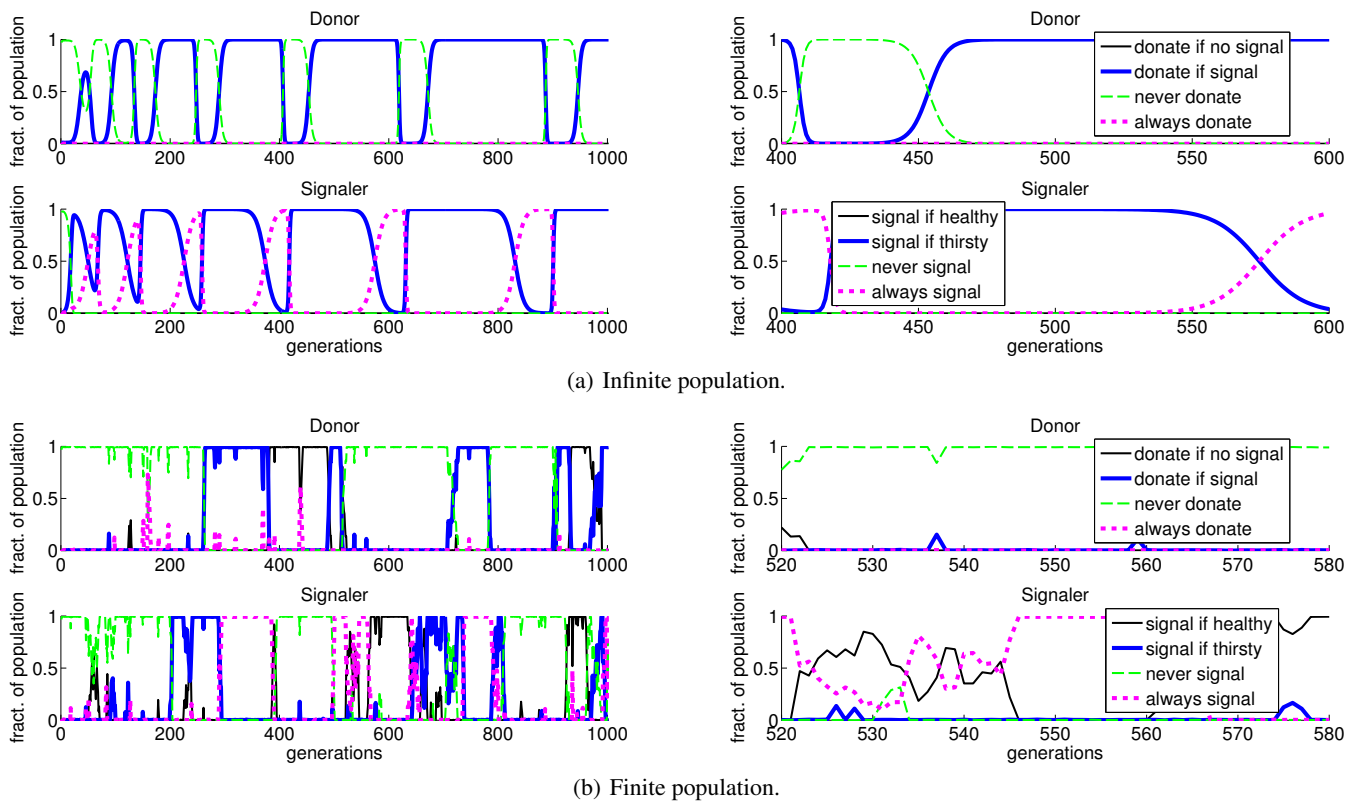


Figure 3: Example of a hybrid equilibrium. The graphs on the right hand side are zoomed in versions of those on the left. Note that the “signal if healthy” strategy invades the signaler population in the finite model. This strategy is essentially non-existent in the continuous model.

Evolutionary Algorithms

When evaluating finite populations we employ a simple genetic algorithm (Mitchell, 1996). In both populations individuals are represented as integers between 1 and 4 representing the strategies listed in Tables 1(c) and 1(b). Strategies are mutated with a probability of 0.01, and no crossover is used. Mutation is performed by replacing an individual with a randomly generated strategy. Each individual plays 50 games against randomly selected individuals from the opposing population, and the average payoff of these games is treated as the individual’s fitness.

A number of selection methods have been employed in evolutionary algorithms. In this study we focus on tournament selection due to its prevalence in the study of the emergence of signaling. In tournament selection, individuals are selected for reproduction by repeatedly choosing the best individuals from small randomly picked subsets. It has been shown that this “best-of-group” version of tournament selection has pathological behavior in terms of maintaining an ESS (Ficici, 2006). This finding helps motivate our hypothesis that this pathology might be present in studies of the emergence of signaling. (Nowak et al., 2004) extends the idea of ESS to finite populations as ESS_N where N is

the population size. We ensure that all individuals have an equal opportunity to compete by constructing tournaments with random permutations of the population.

Results

The results are presented in two sections. We first investigate the difficulty of reaching particular types of equilibria from non-communicative initial population distributions. We then use the parameters from the first investigation in a comparison of infinite and finite population sizes, the latter are investigated with multiple tournament sizes.

Emergence of Signaling

Game theoretic studies of the SPS game generally lead to statements about the existence of particular types of equilibria if certain conditions hold true for a given set of parameters. However, the existence of an equilibrium does not imply that the equilibrium is reachable from arbitrary population distributions. This has significant implications for the emergence of signaling. Under what conditions can an equilibrium be reached from a non-communicative origin?

We approach this question empirically. For each type of

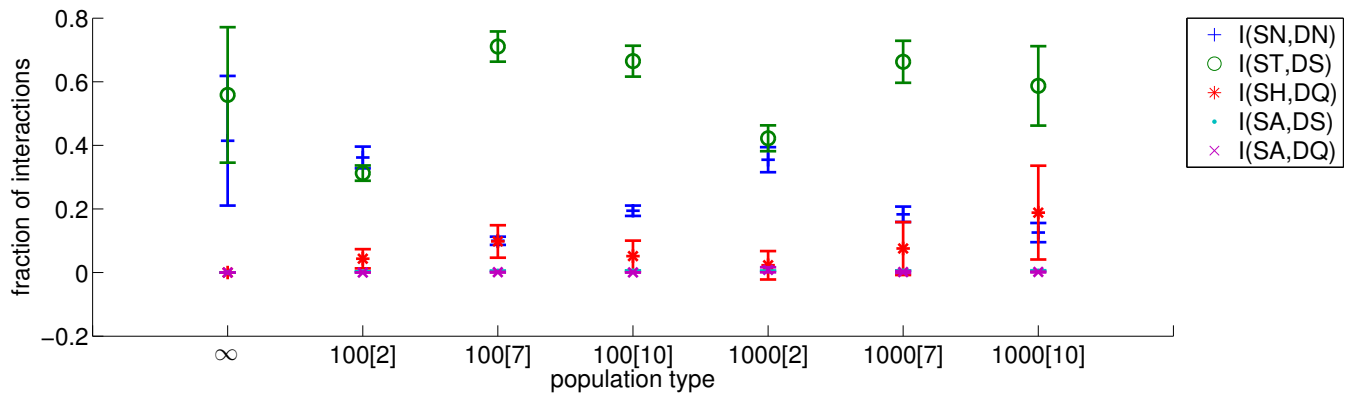


Figure 4: Results for honest signaling equilibria. $I(SX, DX)$, where SX and DX denote signaler and donor strategies, indicates the mean expected number of interactions over time with error bars showing standard deviation. “Infinite” identifies results from the continuous model. The rest of the labels in the form of $x[y]$, denote population and tournament size, respectively.




Figure 5: Results for inverse honest signaling equilibria.

equilibria we generate 1,000,000 random parameters¹ that satisfy the conditions presented in the sections describing equilibria, and test to see whether a continuous model initialized with non-communicative population distributions actually reaches the target equilibrium. The success rate for a given equilibria type quantifies the size of the basin of attraction in parameter space.

Populations are initialized with primarily non-signalers and non-donors (97% of the population) and small fractions of the remaining strategies (1%). We evolve the populations with the discrete-time replicator for 1,000 generations and test to see if the evolutionary trajectory matches that of the corresponding equilibria. For signaling and noncommunicative equilibria, we assume that the system has reached the target if the dominant strategy for signalers and donors matches that of the given equilibrium. For hybrid and pooling equilibria, we compute the mean of the distribution of strategies over time. We look for a match using these means

¹While 1,000,000 may seem like a large number of parameters to test, evaluations of the continuous model are very fast.

for dominant signaler and donor strategies, assuming that strategies with continuously small distributions are eliminated. All parameters that produce the appropriate behavior within 1,000 generations are recorded. In Table 2 we present the success rate for reaching particular equilibria from non-communicative initial conditions.

We can see that honest signaling, followed by hybrid, are the hardest type of equilibria to reach given noncommunicative population distributions. This is followed by inverse honest signaling and pooling II (where donor strategies are a mix of DA and DQ against SN) equilibria. We observe that of the 1,000,000 parameter sets generated for each, less than 10% were able to reach any of these target equilibria. It is not particularly intuitive that inverse honest signaling equilibria are easier to reach than the honest signaling equilibria. However, we note that in order to reach an inverse honest signaling equilibrium the system must pass through a configuration like that of a pooling equilibrium. The pooling equilibrium that it passes through is the SN and DA/DQ profile. Additionally, it can be seen that it is easier to reach an hybrid equilibrium than an honest signaling equilibrium




Figure 6: Results for hybrid equilibria.

Equilibrium type	Success rate
Honest signaling	0.0027
Inverse honest signaling	0.0628
Pooling I	0.9996
Pooling II	0.0969
Hybrid	0.0129

Table 2: Success rate for reaching the appropriate equilibrium from non-communicative initial conditions. Rates are computed based upon 1,000,000 randomly generated parameters that satisfy the conditions of the respective equilibria.

from non-communicative initial conditions.

Infinite and Finite Populations

We are interested in the emergence of signaling, as such all simulations are initialized with populations of primarily non-signalers and non-donors. The populations are evolved for 1,000 iterations for both infinite and finite populations. For each equilibria, 200 parameter sets are randomly chosen from those that reached target in the previous search. Then each evolutionary configuration is evaluated on a given parameter set. Finite populations are repeated 10 times and averaged.

We measure the distance from the true equilibrium with the expected number of interactions given the current population distributions. This is denoted as

$$I(SX, DX) = \frac{|SX| * |DX|}{\sum_{Si, Dj} |Si| * |Dj|}$$

where SX is the signaler strategy of interest, DX is the donor strategy of interest, $Si \in \{ST, SH, SA, SN\}$, and $Dj \in \{DS, DQ, DA, DN\}$.

We take the mean for each expected interaction over time for both the continuous and finite models. For finite populations we look at population sizes of 100 and 1,000, and tournament sizes of 2, 7, and 10. These population sizes span the order of magnitudes that are generally used in studies of the emergence of signaling. Likewise, these tournament sizes span the range commonly used in such studies.

Figure 4 and 5 suggest that the finite model is a good approximation of the continuous model for Nash equilibria. The expected interactions for finite populations of both sizes roughly estimate those calculated in the continuous model (labeled infinite on the x-axis) for tournament sizes greater than 2. Figure 4 suggests that the finite populations approach the behavior of the infinite population as tournament size increases. Tournaments of size 2 perform particularly poorly relative to bigger tournament sizes in the case of signaling equilibria. This is counter to Reggia et al.'s finding where they see that smaller tournament sizes actually lead to higher proportions of signalers in the population (Reggia et al., 2001). This leads us to suggest that in their case the complex environment and agent architecture may have a bias towards signaling behavior.

Figure 6 we see that the finite model fails to capture the complex dynamics of hybrid equilibria. This is because hybrid equilibria are actually collections of polymorphic mixed Nash equilibria. It has previously shown that tournament selection cannot converge to polymorphic Nash equilibria in both one- (Ficici et al., 2005) and two-population coevolution (Ficici, 2006). This leads us to question the significance of the dynamics observed in previous studies of the emergence of signaling. If it is not possible for a simple evolutionary model with tournament selection to maintain a polymorphic Nash equilibrium, then what are the complex dynamics that have previously been observed (Bullock, 1997)? We suggest that these types of dynamics may be a direct result of the spatial selection mechanism based upon previous findings that spatial games can produce behaviors ranging from chaotic dynamics to asymptotically predictable population dynamics (Nowak and May, 1992; Roca et al., 2009).

Conclusion

We have presented a coevolutionary study of the effects of evolutionary mechanics on the emergence of signaling. In doing so we quantify Bullock's previous finding that the existence of a signaling equilibrium does not imply that it can be reached from an initially non-communicative state (Bullock, 1997). It is also shown that it is significantly easier for signaling to evolve from non-communication to an inverse signaling equilibrium than to the signaling equilibrium traditionally studied in the SPS game. Recall that the difference between these two signaling equilibria is when the signal is sent, while the donor adopts the response corresponding to honest signaling. This observation aligns with the signal of the peacock's tail to the peahen, which is a demonstration of virility not aridity.

Finally, we have shown that finite population models with tournament selection can fail to capture the dynamics of hybrid equilibria, one of the most attractive dynamics of the SPS game. These equilibria (which are actually families of polymorphic Nash equilibria) follow a spiraling trajectory that switch between honest and deceptive signaling. The inability of tournament selection to maintain polymorphic Nash equilibria is already known (Ficici et al., 2005). The enhanced reachability of hybrid equilibria relative to traditional signaling ESS's suggests that the generalizability of evolutionary models which fail to capture this phenomenon are limited.

Acknowledgments

We thank the DEMO lab, past and present; in particular, Sevan Ficici. Computing support was provided by the Brandeis HPC.

References

- Bergstrom, C. (2012). The Theory of Honest Signaling. http://octavia.zoology.washington.edu/handicap/honest_biology_01.html.
- Bergstrom, C. and Lachmann, M. (1997). Signalling among relatives. I. Is costly signalling too costly? *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 352(1353):609–617.
- Bullock, S. (1997). An exploration of signalling behaviour by both analytic and simulation means for both discrete and continuous models. In *Proceedings of the Fourth European Conference on Artificial Life*, pages 454–463.
- Ficici, S. (2006). A game-theoretic investigation of selection methods in two-population coevolution. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 321–328.
- Ficici, S., Melnik, O., and Pollack, J. (2005). A game-theoretic and dynamical-systems analysis of selection methods in coevolution. *Evolutionary Computation, IEEE Transactions on*, 9(6):580–602.
- Ficici, S. and Pollack, J. (2007). Evolutionary dynamics of finite populations in games with polymorphic fitness equilibria. *Journal of Theoretical Biology*, 247(3):426–441.
- Fogel, D. and Fogel, G. (1995). Evolutionary stable strategies are not always stable under evolutionary dynamics. *Evolutionary Programming IV*, pages 565–577.
- Grafen, A. (1990). Biological signals as handicaps*. *Journal of Theoretical Biology*, 144(4):517–546.
- Hamilton, W. (1964). The genetical evolution of social behaviour. II. *Journal of Theoretical Biology*, 7(1):17–52.
- Hurd, P. (1995). Communication in discrete action-response games. *Journal of Theoretical Biology*, 174(2):217–222.
- Huttegger, S. and Zollman, K. (2010). Dynamic stability and basins of attraction in the Sir Philip Sidney game. *Proceedings of the Royal Society B: Biological Sciences*, 277(1689):1915–1922.
- Johnstone, R. and Grafen, A. (1992). The continuous Sir Philip Sidney game: a simple model of biological signalling. *Journal of Theoretical Biology*, 156(2):215–234.
- Lindgren, K. (1991). Evolutionary phenomena in simple dynamics. *Artificial Life II*, 10:295–312.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*.
- Maynard Smith, J. (1991). Honest signalling: The Philip Sidney game. *Animal Behaviour*, 42:1034–1035.
- Maynard Smith, J. and Price, G. R. (1973). The Logic of Animal Conflict. *Nature*, 246(5427):15–18.
- Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press.
- Noble, J. (1999). Cooperation, conflict and the evolution of communication. *Adaptive Behavior*, 7(3-4):349–369.
- Nowak, M. and May, R. (1992). Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829.
- Nowak, M., Sasaki, A., Taylor, C., and Fudenberg, D. (2004). Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428(6983):646–650.
- Ozisk, A. P. and Harrington, K. I. (2012). The Effects of Tags on the Evolution of Honest Signaling. In *ECOMASS Workshop, Workshop Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2012)*, in press.
- Reggia, J., Schulz, R., Wilkinson, G., and Uriagereka, J. (2001). Conditions enabling the evolution of inter-agent signaling in an artificial world. *Artificial Life*, 7(1):3–32.
- Roca, C., Cuesta, J., and Sánchez, A. (2009). Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics. *Physics of Life Reviews*, 6(4):208–249.
- Saunders, G. and Pollack, J. (1996). The Evolution of Communication Schemes Over Continuous Channels. In *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 580–589.
- Sigmund, K. and Hofbauer, J. (1998). Evolutionary games and population dynamics. *Cambridge Univ. Press*.
- Wagner, K., Reggia, J., Uriagereka, J., and Wilkinson, G. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69.
- Zahavi, A. (1975). Mate selection—a selection for a handicap. *Journal of Theoretical Biology*, 53:205–214.