MBS TECHNICAL REPORT 18-03

Hierarchical Models for the Evolution of Compositional Language

Jeffrey A. Barrett, Brian Skyrms, Calvin Cochran

May 30, 2018

Abstract

We present three hierarchical models for the evolution of compositional language. Each has the basic structure of a two-sender/one receiver Lewis signaling game augmented with executive agents who can learn to influence the behavior of the basic senders and receiver. With each game, we move from stronger to weaker modeling assumptions. The first game shows how the basic senders and receiver might evolve a compositional language when the two senders have pre-established representational roles. The second shows how the two senders might coevolve representational roles as they evolve a reliable compositional language. Both of these games impose an efficiency demand on the agents. The third game shows how costly signaling alone might lead role-free agents to evolve a compositional language.

1 Introduction

Humans and some animals use languages that allow for the functional composition of basic terms to form more complex expressions. The meanings of the more complex expressions are influenced, sometimes determined, by the meanings of their parts. Animals where there is compelling evidence for such compositional languages include putty nosed monkeys [2, 3], Campbell's Monkeys [20], suricates [18], prairie dogs [12, 29], and some species of birds [27].[1] Here we are concerned

---

[1] The sort of composition exhibited can be subtle. Suricates, for example, compose acoustically two aspects of their call to indicate predictor class and urgency.

with how a basic compositional language might evolve. The aim is to model the evolution of a very simple compositional language using the basic structure of a Lewis signaling game.[2]

The simplest sort of Lewis-signaling game involves two players: a sender and a receiver. The sender is equipped with N signals which she can send to the receiver, and the receiver is equipped with N acts to choose from. Each play of the game, nature chooses one of N states with unbiased probabilities and reveals it to the sender. The sender then sends one of her N signals to the receiver, who cannot see the state of nature directly. The receiver then chooses an act conditional on the signal. Each state of nature corresponds to exactly one act, and the players both win if and only if the act chosen by the receiver corresponds to the current state of nature. When successful, the agents evolve a signaling system where the sender associates exactly one signal with each state of nature, the receiver associates each of these signals with the corresponding act, and both players always succeed regardless of the state of nature.

Signaling games have been studied extensively under different learning and evolutionary dynamics, and there are a number of analytic and simulation results [1, 14, 24, 5, 9, 15, 16]. And many variations of the basic game have also been studied. These include games where nature is biased, where the agents have too few or too many signals, and where there are multiple senders or receivers [4, 6, 25].

The games we are concerned with here have two basic senders and one receiver, but they are distinguished from other such games by their hierarchical structure. In these games the basic agents play a sender or receiver role as in a standard Lewis signaling game while the executive sender and executive receiver track and control aspects of the behavior of the basic agents. Together the agents must evolve a simple compositional language in order to be successful. In the *special composition game*, the basic agents have preassigned representational roles which they have to learn to perform while the executive agents learn how to control basic agents with those particular assigned roles. In the *role-free composition game*, the basic agents have no preassigned representational roles. Here the executive agents have to learn what roles the basic agents are playing and how they might be used for successful action even as the basic agents evolve representational roles. The dispositions of all of the agents coevolve as the compositional

---

[2] See [11] and [30] for alternative models of compositional signaling.

language evolves. The *general composition game* is just like the role-free game except that no explicit efficiency demands are placed on the basic agents. Here the agents evolve a compositional language as a result of costly signaling alone.

Costly signaling is ubiquitous in nature. In bacteria, and other microorganisms, each signal sent may involve producing a molecule which diffuses into the vicinity. Here each signal has a metabolic cost. For higher organisms, if giving an alarm call exposes an individual to increased immediate danger, then pausing to give two signals might increase danger. Again, this gives each signal a cost. More generally, costly signaling in nature is a well-studied phenomena with an extensive literature.[3]

The games we are concerned with here are significantly more subtle than Lewis' original signaling game. They are *generalized signaling games* with a hierarchical structure. As with simpler signaling games, generalized signaling games might self-assemble from the ritualization of individual actions of the agents [7, 8]. The behavior of the agents in the complex game is forged in the context of an evolutionary process. In the present models, this is a learning dynamics. We will refer to the parts of the model as agents, but they may be understood as functional components of a social group or of an individual agent.

2 The special composition game

The special composition game is a variant of the traditional Lewis signaling game. It is a cooperative game with two basic senders and one basic receiver where the agents must evolve a particular sort of signaling system in order to be uniformly successful. In addition to the basic agents, the game has an executive sender and an executive receiver. These hierarchical agents can learn to influence the behavior of the basic agents.

The state of nature features two *properties* and a *context*. For concreteness we will take the two properties to be *color*, which is either *black* or *white*, and *animal*, which is either *dog* or *cat*. The state of nature on a particular play of the game will feature either *black dog*, *black cat*, *white dog*, or *white cat*. The context indicates which of the properties the receiver will need to know in order to

---

[3] See, for example, [28], [23], [31], and [22].

perform a successful action on the current play. The context is either *color*, *animal*, or *both*.

The game is played by two senders and an executive sender and one receiver and an executive receiver. While we will drop this condition for the role-free and general composition games, on the special composition game each basic sender is assigned a particular property and only has access to that aspect of nature. So one basic sender sees *color* and the other sees *animal*. The executive sender sees *context*. Initially, the executive sender randomly determines whether the color sender, the animal sender, or both will send a signal. Over time the executive sender may learn what type of signal the current context demands.

The basic receiver sees the signals sent by the two basic senders and which sender sent each. The executive receiver also sees who sent the signals. The executive receiver determines whether the basic receiver will interpret the signal as a color, an animal, or both. The receiver performs an action based on his interpretation. The receiver's possible actions are *black*, *white*, *dog*, *cat*, *black dog*, *white dog*, *black cat*, or *white cat*. See figure 1.

The agents are successful on a particular play of the game if and only if (1) the receiver performs the correct action given the current context and (2) the senders only sent the signals required for success given the context. To be successful, then, the receiver must match the state of nature and the senders must be as efficient as possible, not sending any irrelevant signals. We will suppose that the agents learn by simple reinforcement.[4]

The game is implemented as follows. On each play, nature chooses a value for each of the two properties and the context randomly and with uniform probabilities. The executive sender is equipped with an urn for each of the three possible contexts *color*, *animal*, and *both*. Each urn begins with one ball of each type *color-sender* (c), *animal-sender* (a), and *both* (b). Upon witnessing the context, the executive sender randomly draws a ball from the corresponding urn. The drawn ball determines whether the color sender, animal sender, or both will send a signal.

The color sender is equipped with a white urn and a black urn, each initially containing a 0 ball and a 1 ball. If the executive sender draws a ball that requires the color sender to send a signal, the color sender

---

[4] See [13, 10, 21] for discussions of reinforcement learning generally and as a model for human learning.

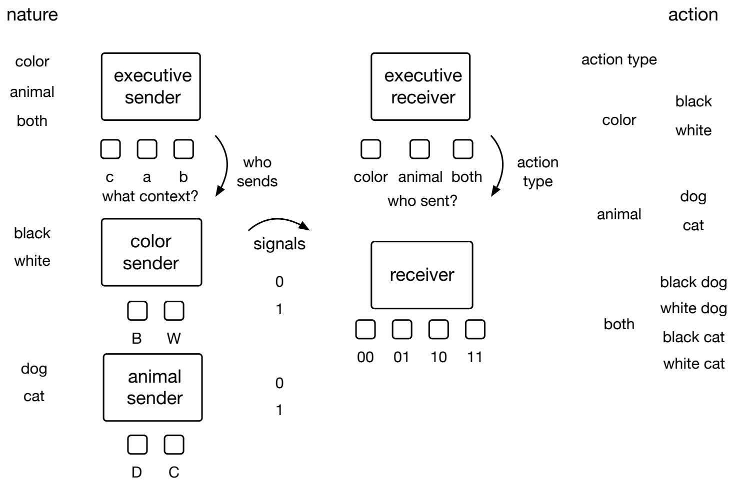

randomly draws a ball from the urn corresponding to the color she sees and sends the corresponding signal. Similarly, the animal sender is equipped with a dog urn and a cat urn, each initially containing a 0 ball and a 1 ball. If required to do so by the executive sender, she draws a ball from the urn corresponding to the animal she sees and sends the corresponding signal.

The receiver has four urns, one for every ordered pair of signals she might receive from the color sender and animal sender respectively: 00, 01, 10, and 11. Each urn begins with one ball for each of the possible color-animal pairs *black dog*, *white dog*, *black cat*, or *white cat*. If both senders send a signal, then the receiver draws a ball at random from the corresponding urn. If only one sender sends a signal, then the receiver randomly chooses, with unbiased probabilities, one of the two urns that correspond to the sender's signal then draws a ball at random from that urn.

The executive receiver determines how the receiver will interpret the *type* of signal she received. This interpretation together with the ball the receiver drew determines how the receiver will act. The executive receiver is equipped with a *color-sender* urn, an *animal-sender* urn, and a *both* urn. Each of these initially contains a *color* ball, an *animal* ball, and a *both* ball. The ball drawn by the executive receiver determines what type of act the receiver takes as salient given the signal(s) she received. Specifically, if the executive receiver draws a *both* ball, then the basic receiver simply performs the action corresponding to the ball she drew, but if the executive receiver draws a *color* or *animal* ball, then the basic receiver performs the action corresponding to the color or animal indicated on the ball she drew. The balls drawn by executive receiver and the basic receiver, then, determine whether the receiver's action is *black*, *white*, *dog*, *cat*, *black dog*, *white dog*, *black cat*, or *white cat*.

Again, the agents are successful on a particular play of the game if and only if (1) the receiver performs the correct action given the current context and (2) the senders only send the signals required for success given the context. The second condition requires the senders to be as efficient as possible given the current context. Namely, they are only successful if a signal from the color sender is sent only when the context requires color and a signal from the animal sender is sent only when the context requires animal. One might think of this condition as imposing a significant cost on inefficient signaling even if it leads to the right action. It is this condition that drives the evolution of a compositional language.

If a play of the game is successful, then each agent who was involved in that particular play returns the ball she drew to the urn from which she drew it and adds another a ball of the same type to that urn. Otherwise, each agent simply returns the ball she drew to the urn from which she drew it.

On simulation, the agents in the special composition game nearly always evolve a successful and optimally efficient compositional language. See table 1. As they evolve to satisfy the demands of transmitting information but not transmitting too much, the meanings of the individual terms and the composite expressions coevolve. Specifically, each of the color sender's terms comes to represent a color, each of the animal sender's terms comes to represent an animal, and functional composition comes to represent the conjunction of the two properties.

Simulation results:

# runs	**# plays per run**	**mean overall success rate**
10,000	1,000,000	0.9743
1,000	10,000,000	0.9898
100	100,000,000	0.9975

The executive sender always learns to correctly read the context and the executive receiver always learns to correctly interpret the type of signal sent. Indeed, in this game, reinforcement learning guarantees that they will learn to coordinate between the current context and how the signal is interpreted by the receiver. Suppose that the context requires only color. The agents are only successful if the executive sender instructs only the color sender to send her signal and the executive receiver instructs the receiver to perform a color-type action. So, on the reinforcement dynamics, the executive sender's *color* urn can only ever have *color-sender* balls added to it and the executive receiver's *color-sender* urn can only have *color* balls added to it. And similarly for contexts that require only animal or both for success. This means that the executive sender and executive receiver always evolve the same coordinated conventions for representing the context and interpreting the type of signal sent on the special composition game.[5]

The color sender, animal sender, and receiver, however, have significant freedom in what conventions they may evolve. One almost always (in 0.97 of runs) finds that the agents evolve a maximally efficient and successful compositional language on simulation, but one sees every permutation of successful conventions that allows for this. One such convention is depicted in figure 2.

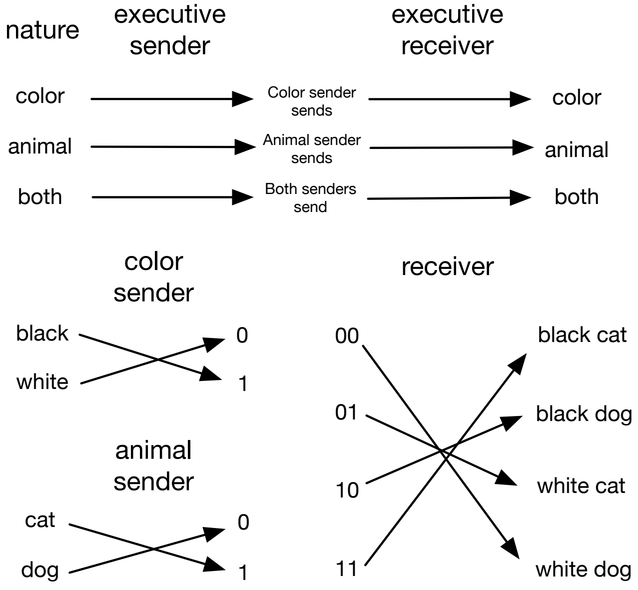

---

[5] But, as we will see, this is not the case for the role-free and general composition games.

3 The role-free composition game

Preassigning each basic sender to one of the two properties imposes significant structure on the special composition game up front, and, as we saw, it is this structure that guarantees that the executive sender and executive receiver always learn to coordinate the representation of the context and interpretation of the signal type. In the role-free composition game the basic senders have no preassigned representational roles. Rather, they are both simply shown the full basic state of nature.

Since each basic sender lacks the expressive resources on his own to represent the full state, they must evolve representational roles in order to be successful. Further, the executive sender must learn what these emerging roles are as they evolve in order to know what signal to send for each context. And the executive receiver must learn what the signal types are as they evolve in order to correctly interpret the signal. If the agents are successful, they will have evolved a compositional language where the representational and interpretational roles of the agents coevolved with the meanings of terms and the significance of the composition of terms.

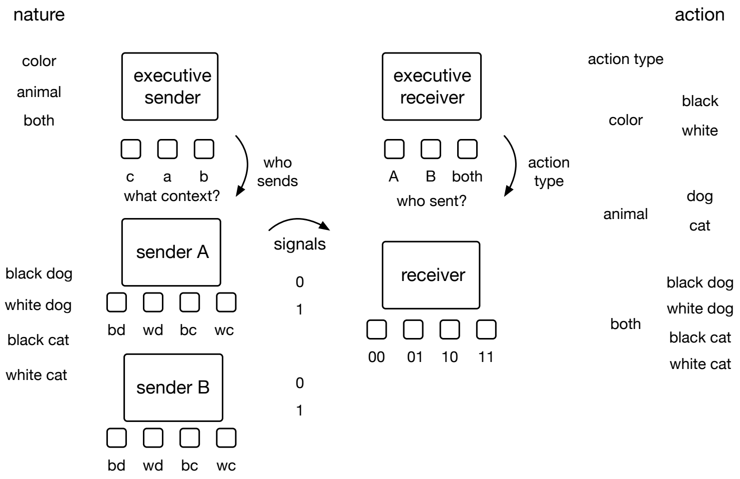

The role-free composition game is very much like the special compositional game, but there are two important differences. First, since there are no preassigned roles, there is, at least initially, no color sender or animal sender. Rather, on each play of the game, the two basic senders, sender A and sender B, are both shown one of the four possible basic states *black dog*, *white dog*, *black cat*, or *white cat*. Each sender is equipped with an urn corresponding to each of these states and each urn begins with a 0-ball and a 1-ball. The balls, again, indicate the possible signals. Since each sender still has only two possible signals for each of the four states, neither is able to convey the full state alone. To be successful, they must somehow partition the state of nature in a way which lets them give complementary information about the state of nature to the receiver.

The second difference from the special composition game concerns the conditions for success. In the basic composition game, a play was successful if and only if the receiver performs the correct action given the state of nature and the senders send precisely the type of information demanded by the current context. Since there is no pre-established color sender or animal sender, this second condition does not even make sense in the role-free composition game. The weaker efficiency condition for the role-free game, then, is that both senders send a signal if and only if the current context requires maximal information. So the full conditions for a successful play are (1) the receiver performs the correct action given the current context and (2) both senders sent a signal if and only if the current context requires both color and animal. Again, we suppose that all agents learn by simple reinforcement.

As can be seen in table 2, the agents do very well overall evolving a compositional language on this more subtle game. But the table just reports the mean cumulative success rate over all the runs. This misses two important aspects of the behavior of the game. First, most of the time (0.736) the composite system evolves nearly perfect (better than 0.98 cumulative success rate) signaling. And second, the composite system sometimes gets stuck in one of several suboptimal pooling equilibria that exhibit significantly worse cumulative success rates. We will discuss each of these behaviors in turn.

Simulation results:

# runs	**# plays per run**	**mean overall success rate**
1,000	1,000,000	0.9072
1,000	10,000,000	0.9332
500	100,000,000	0.9478

When the agents are optimally successful in the role-free composition game, the senders evolve representational roles and the executive sender learns which roles each sender has evolved. The meanings of the individual terms and the meanings of the composite expressions coevolve to allow the executive sender to satisfy the expressive demands required by each possible context. As the agents learn when basic terms and complex expressions are needed, their success and failure in the use of the basic expressions influences the meanings of the complex expressions, and their success and failure in the use of the complex expressions influences the meanings of the basic expressions. In this way, the semantic function of composition coevolves with the meanings of the more basic terms.

The evolution of functional composition here, as in the special composition game, is driven by the contextual demands. To be successful the agents must evolve a system that allows them to communicate precisely the information required for successful action by the receiver on the current play and no more. On simple reinforcement learning, the agents are not directly punished for sending more information than necessary, but they are also not rewarded. The simulation results show that this is typically enough to evolve a compositional language that allows for optimal signaling.

Unlike the special composition game, there is no single canonical convention that evolves between the executive sender's representation of the context and the the executive receiver's interpretation of the signal type in the role-free game. Rather, when optimal signaling evolves, there are two possible conventions that the system might exhibit, each corresponding to a particular assignment of color and animal roles to the two senders (see the figures below).

Since there are no pre-established roles for the senders, when the system evolves optimal signaling, sender-A sometimes evolves to be the color sender and sender B the animal sender and sometimes it is the other way around. Importantly, color and animal properties must be respected by the partition the basic senders evolve in order for the composite system to evolve optimal signaling. While there are other partitions of the four color-animal states that the senders might evolve that would allow them to communicate which of the four basic states of nature was observed, since the contextual demands are either color, animal, or both, the only way for the executive sender to always be able to satisfy the *partition of demands* is for there eventually to be a color sender and an animal sender available. So for the composite system to be successful, the two senders must adjust to *this* constraint as they evolve how they partition nature.

Consider a basic signaling game with four states, two senders each

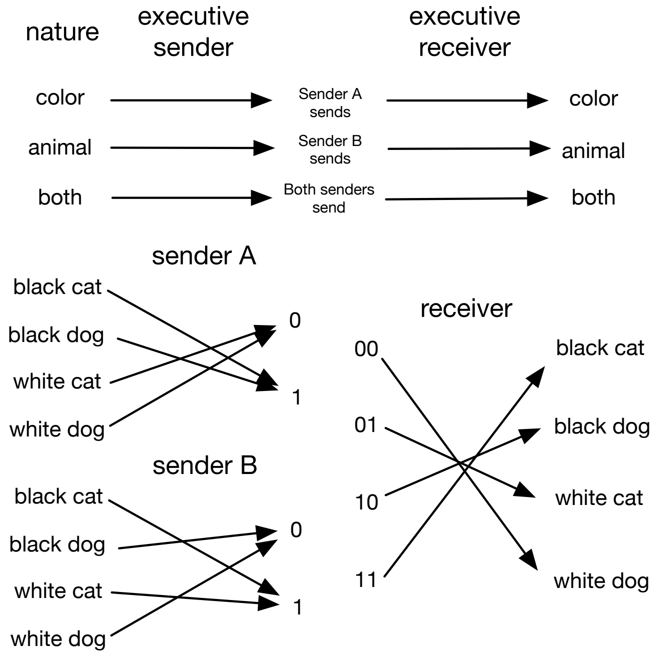

with two possible signals, a single receiver, and no context that must be satisfied. Here both senders just always send a signal. On simulation, this system typically evolves optimal signaling with simple reinforcement learning and no favored partitions of nature. The system is equally likely to evolve either of the two possible optimally effective partitions and either of the two possible assignments of each to the two senders [4, 5].[6] In the role-free composition game, in contrast, there is a single favored partition imposed by the game's contextual demands. More specifically, since the executive sender must be able to express precisely these properties both individually and together, *color* and *animal* might be thought of as "natural kinds" given the de facto payoffs for different contexts. But, of course, different contextual demands would individuate different "nature kinds" in this weak sense of the notion.

When the system evolves optimal signaling, one sender takes the role of representing *color* and the other takes the role of representing *animal* and composition represents the simple conjunction of these two properties. But on the role-free composition game the agents some-

---

[6] Here about 0.73 of the runs have a success rate higher than 0.80 on simulation.

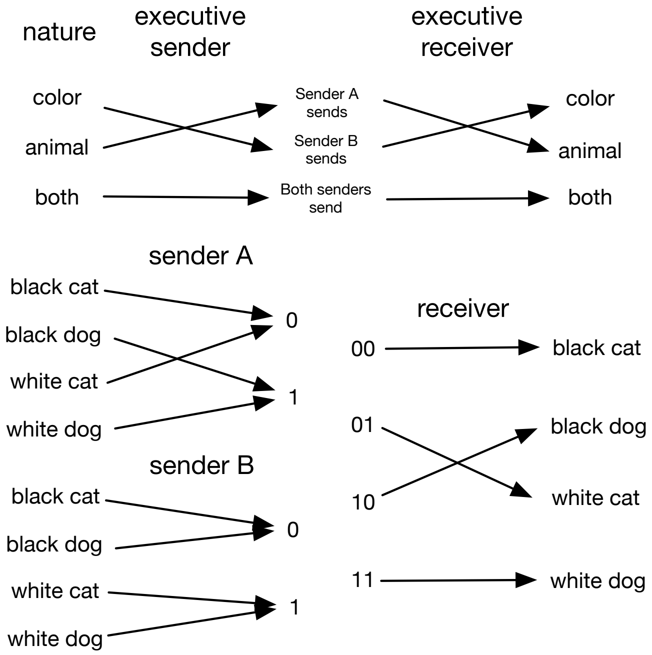

times fail to evolve an optimal signaling system. When this happens, the language they evolve and how it works can be significantly more subtle.

Consider the simulated run of the game that resulted in the player mappings in figure 5. Here sender B has evolved to be a color sender. Color-animal pairs in which the color is *white* are mapped to 1 and those where the color is *black* are mapped to 0. Sender A, however, is neither a color sender nor an animal sender. She sends both *white dog* and *black cat* to 0, so she she cannot by herself distinguish between either color or animal. If the context requires color, the executive sender uses sender B. Such plays are always successful. If the context requires animal, then if the executive sender uses sender A, the agents will only succeed half of the time. But when the context requires both color and animal, the agents are *always* successful. For *black dog* they send 10, for *white dog* they send 01, for *black cat* they send 00, and for *white cat* they send 11. Since the agents are always successful when just animal is required and when both color and animal are required, and since they are successful half the time when just color is required, this leads to an expected success rate of $\frac{1}{3} + \frac{1}{3} + \frac{1}{2} \cdot \frac{1}{3} \approx 0.833$, which is approximately the observed success rate on runs where this structure of dispositions evolves. Such runs occur about 0.11 of the time. [7]

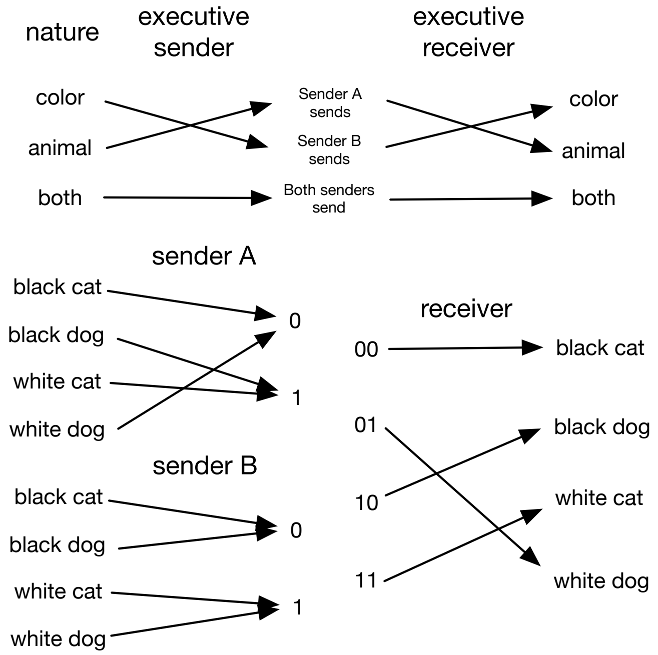

Composition in this case does not involve the simple conjunction of color and animal properties. While sender B's terms refer to *black* and *white*, sender A's terms do not individuate between colors or animals. Rather, the referents of A's terms *cross-cut* both animal and color to individuate between (*white dog* or *black cat*) and (*white cat* or *black dog*). Given the salient contexts presented by nature, A's terms are entirely useless by themselves. This failure in efficiently marshaling the representational resources of the agents has the consequence that if the context requires just animal, the agents are helpless and cannot do better than chance.

But sender A's terms are useful *when composed* with sender B's terms. When the context requires both animal and color, B's term communicates a color and, together with A's term, also selects an animal. It is not that A's terms are meaningless. Rather, it is that they are never useful on their own given the contextual demands. Here

---

[7] Less interesting things, such as the basic agents evolving the opposite partition to that used by the executive agents or the executive receiver never choosing a certain context like color or animal, happen at about 0.05 frequency.

functional composition evolves to produce useful expressions from terms that are not always useful alone.

In natural language we routinely employ expressions that play important semantic roles in combination with other expressions but are relatively useless on their own. English language adverbs, adjectives, and pronouns often behave this way. While the term *only* is rarely useful on its own, expressions like *only child* or *only decaf* have precise meanings that may allow one to usefully characterize states of nature and, hence, facilitate successful action.

4 The general composition game

The *general composition game* is precisely like the role-free game except for how actions are reinforced. In the games we have considered so far the agents are only successful if they both choose the correct act and utilize the minimum number of signals necessary for the given context. In contrast, the general composition game rewards successful action however it might be achieved but imposes a fixed cost to each signal. And, importantly, like the role-free game, the basic senders have no pre-established representational roles.

The general composition game introduces three new parameters to the role-free game: simple-context payoff, complex-context payoff, and signal cost. Now when the agents choose the correct act, regardless of the number of signals sent, they get a basic payoff. If the context for that play was simple (*color* or *animal*), then the reward is the simple-context payoff. If the context was *both*, agents receive the complex-context payoff. These payoffs might be equal in some situations, but one might also imagine situations where the rewards for success are greater in a complex context where more information is required. If the receiver's action is not successful, the basic payoff is zero.

Reinforcements are a function of both basic payoffs and signaling costs. Specifically, the fixed cost of each signal is subtracted from the basic payoff and the corresponding number of balls are added to or subtracted from the agents' urns.[8] If the agents are unsuccessful, then the basic payoff is zero but since they still have signaling costs to pay, they will lose balls from the urns that led to the failed action.[9] Agents

---

[8] Note that this may involve adding or subtracting fractions of balls depending on the specific payoffs and costs involved.

[9] If a punishment would ever push the number of balls of a particular type to something

are thus reinforced the strongest when they use the least number of signals to achieve the greatest reward, and they are punished if the signal costs they incur are higher than their basic payoff.[10]

Suppose that the simple-context payoff is 1.5, the complex-context payoff is 2, and signal cost is 0.5.[11] One might *measure* the agents' overall success rate by just tracking the proportion of plays in which they chose the correct act. In order to give a better sense of how the agents' dispositions evolve and for the purposes of comparison to the results of the simple and role-free games, we will measure their success on the stricter condition that they (1) chose the correct act and (2) send the minimum number of signals needed to convey that act. Importantly, note that, regardless of how successful their actions may be, the agents have only evolved *to use* a compositional language if they in fact exhibit this sort efficiency.

On this measure, after 1,000 runs of 10^8 plays each, the mean overall cumulative success rate on simulation was found to be 0.9942. Of these runs, 0.975 exhibited near perfect and efficient signaling. The agents crosscut nature (i.e. develop suboptimal roles) in only 0.022 of runs. In addition, agents approach the perfectly efficient signaling relatively quickly, with rates comparable to those in the special signaling game, as can be seen in the table below:

Simulation results for the general composition game:

# runs	**# plays per run**	**mean overall success rate**
1,000	1,000,000	0.9851
1,000	10,000,000	0.9905
1,000	100,000,000	0.9942

Setting the simple-context payoff and the complex-context payoff both to 2 and keeping a signal cost of 0.5 produces a slightly lower overall success rate of 0.9733 with 0.936 of runs near perfect play. The fact that the agents are somewhat less successful in this case seems to have less to do with the relationship between the relative payoffs of the simple and complex contexts and more to do with the relationship between the basic payoffs and signal cost. When the basic payoffs are increased and/or the signal cost is decreased, the agents do not do

---

less than one, we set the number of balls of that type to one.

[10] There is ample evidence of signaling in biological and human interactions.

[11] This produces payoffs in the general game that are similar to the payoffs in the role-free game above.

as well at evolving a successful a compositional language. This can be seen by the fact that when the simple-context payoff is 1.5, the complex-context payoff is 2, and signal cost is 0.3 the agents evolve nearly perfect signaling 0.61 of the time; and when the signal cost is 0.1 they evolve near perfect signaling only 0.26 of the time.

The relationship between the basic payoffs and the signal cost drives the evolution of a successful compositional language. On the first several plays of the game, where no signals or acts are salient, agents will usually fail but occasionally succeed by chance. While these early successes can become a foothold on which to build a successful compositional language, there is also a chance that they will start the agents down a path towards a suboptimal equilibrium. One way that this can happen is if their action is successful when the context is *both* and one basic sender is cross-cutting nature. While a cross-cutting equilibrium will succeed if the context is *both*, it may not when the cross-cutting agent is chosen to send her signal alone since cross-cutting signals are useless in the simple contexts of *color* or *animal*. The fact that signals have significant costs means that when a cross-cutting sender's signal results in a failure in a simple context, the agents subtract balls from their urns and increase the likelihood of a more successful compositional language evolving, one that works well both for simple and complex contexts.[12]

5 Discussion

The three models for the evolution of compositional language considered here show how basic terms and functional composition might coevolve in a hierarchical signaling game under simple reinforcement with efficiency conditions on success and, in the case of the general composition game, under simple reinforcement learning with costly signaling.

When the agents are successful in the general composition game, the basic senders adopt representational roles and evolve signaling conventions appropriate to each role, the executive sender learns which roles each sender has adopted and how to use her signals to represent the current context, the executive receiver learns how to interpret the

---

[12] Cross-cutting equilibria are possible and occur on about 0.02 of the runs with the original parameter settings. If the context is repeatedly *both*, the agents' urns may fill too quickly for ball withdrawals to steer them away from the cross-cutting equilibria later.

type of expression sent, and the basic receiver learns how to interpret the specific content of each expression type. The meanings of the individual terms and the composite expressions *coevolve* to satisfy the expressive demands required by the salient contexts. In this sense, the semantic function of composition coevolves with the meanings of the individual terms.

The evolution of functional composition is required to satisfy the contextual demands under the efficiency constraints. If the agents were always reinforced for sending maximal information, then there would be no need for the individual terms to evolve their own meanings.

When the agents evolve an optimally successful system in the general composition game, both the individual terms and the composite expressions are useful. But we have also shown how terms might evolve that are only useful when composed with other meaningful terms.[13]

---

[13] We would like to thank in particular Travis LaCroix and Josh Armstrong for their insightful comments on an earlier draft of this paper.

References

[1] Argiento, R. R. Pemantle, B. Skyrms and S. Volkov [2009] 'Learning to Signal: Analysis of a Micro-Level Reinforcement Model,'*Stochastic Processes and Their Applications* 119(2): 373–390.

[2] Arnold, Kate, and Klaus Zuberbühler (2006) "The Alarm-Calling System of Adult Male Putty-Nosed Monkeys, Cercopithecus Nictitans Martini," *Animal Behaviour* 72: 643–653.

[3] Arnold, Kate and Klaus Zuberbühler (2008) "Meaningful Call Combinations in a Non-Human Primate," *Current Biology* 18(5): R202–3.

[4] Barrett, J. A. [2006], "Numerical simulations of the Lewis Signaling Game: Learning Strategies, Pooling Equilibria, and the Evolution of Grammar" *Institute for Mathematical Behavioral Sciences.* Paper 54. http://repositories.cdlib.org/imbs/54

[5] Barrett, J. A. [2007] 'Dynamic Partitioning and the Conventionality of Kinds,' *Philosophy of Science* 74: 527–546.

[6] Barrett, J. A. [2013] 'The Evolution of Simple Rule-Following' *Biological Theory* 8(2): 142–150.

[7] Barrett, J. A. and B. Skyrms [2016]: 'Self-Assembling Games,' *British Journal for the Philosophy of Science.* First published online 13 September 2015. doi: 10.1093/bjps/axv043.

[8] Barrett, J. A., B. Skyrms and A. Mohseni [2017] 'Self-Assembling Networks,' forthcoming in *British Journal for the Philosophy of Science.*

[9] Barrett, J. A. and K. Zollman [2009] 'The Role of Forgetting in the Evolution and Learning of Language,' *Journal of Experimental and Theoretical Artificial Intelligence* 21(4): 293–309.

[10] Erev, I. and A. E. Roth [1998] 'Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria' *American Economic Review* 88: 848–81.

[11] Franke, M. (2014) 'Creative compositionality from reinforcement,' in E. A. Cartmill, S. Roberts, H. Lyn, and H. Cornish (Eds.), *The evolution of language (Proceedings ofEvoLang 10).* Singapore: World Scientific, pp. 82–89.

[12] Frederiksen, J. K. and C. N. Slobodchikoff (2007) ‘Referential Specificity in the Alarm Calls of the Black-Tailed Prairie Dog,’ *Ethology, Ecology & Evolution* 19: 87–99.

[13] Herrnstein, R. J. [1970] ‘On the Law of Effect,’ *Journal of the Experimental Analysis of Behavior* 13: 243–266.

[14] Hofbauer, J. and S. Huttegger [2008] ‘Feasibility of Communication in Binary Signaling Games,’ *Journal of Theoretical Biology* 254(4): 843–849.

[15] Hu, Y., B. Skyrms, P. Tarrès [2011] ‘Reinforcement Learning in a Signaling Game,’ arXiv:1103.5818 [math.PR].

[16] Huttegger, S., B. Skyrms, P. Tarrès, and E. Wagner [2014] ‘Some Dynamics of Signaling Games,’ *Proceedings of the National Academy of Sciences* 111(S3): 10873–10880.

[17] Lewis, D. [1969] *Convention.* Cambridge, MA: Harvard University Press.

[18] Manser, M., R. Seyfarth and D. Cheney [2002] ‘Suricate Alarm Calls Signal Predator Class and Urgency,’ *Trends in Cognitive Science* 6(2): 55–57.

[19] McGregor, P. [2005] *Animal Communication Networks.* Cambridge: Cambridge University Press.

[20] Ouattaraa, K., A. Lemassona, K. Zuberbühler [2009] ‘Campbell’s Monkeys Concatenate Vocalizations into Context-Specific Call Sequences’ *Proceedings of the National Academy of Science* 106(51): 2202622031, doi: 10.1073/pnas.0908118106.

[21] Roth, A. E. and I. Erev [1995] ‘Learning in Extensive Form Games: Experimental Data and Simple Dynamical Models in the Immediate Term,’ *Games and Economic Behavior* 8:164–212.

[22] Searcy, W. A. and S. Nowicki [2006] *The Evolution of Animal Communication: Reliability and Deception in Signaling Systems.* Princeton: Princeton University Press.

[23] Sherman, P. W. [1977] ‘Nepotism and the Evolution of Alarm Calls,’ *Science* 197: 1246–1253.

[24] Skyrms, B. [2010] *Signals: Evolution, Learning, & Information.* New York: Oxford University Press.

[25] Skyrms, B. [2009] ‘Evolution of Signaling Systems with Multiple Senders and Receivers,’ *Philosophical Transactions of the Royal Society* B 27 364 (1518): 771–779.

[26] Skyrms, B. [2006] ‘Signals,’ *Philosophy of Science* 75(5): 489–500.

[27] Suzuki, Toshitaka N., David Wheatcroft, and Michael Griesser (2016) ‘Experimental Evidence for Compositional Syntax in Bird Calls,’ *Nature Communications* 7, Article number: 10986. doi:10.1038/ncomms10986

[28] Maynard Smith, J. [1965]. ‘The evolution of alarm calls,’ *American Naturalist* 99: 59–63.

[29] Slobodchikoff, C. N., A. Paseka and J. L. Verdolin (2009) ‘Prairie Dog Alarm Calls Encode Labels about Predator Colors,’ *Animal Cognition* 12: 435–439.

[30] Steinert-Threlkeld, Shane (2016) ‘Compositional Signaling in a Complex World,’ *Journal of Logic, Language and Information* 25(3–4): 379–397.

[31] Zahavi, A. and A. Zavahi [1997] *The Handicap Principle: A Missing Piece of Darwins Puzzle.* Oxford: Oxford University Press.