Costly Invention Informs Compositional Signaling

by
Eissa Haydar[1]

A thesis submitted in partial fulfillment
of the requirements for the degree of
Bachelor of Arts with Honors
Department of Philosophy
in the University of Michigan
2024

Advisor: James M. Joyce (C. H. Langford Collegiate Professor of Philosophy)

Second Reader: Gordon Belot (Lawrence Sklar Collegiate Professor of Philosophy)

[1] I can be contacted at: ehaydar@umich.edu

Table of Contents

Acknowledgements

First and foremost, I would like to acknowledge Jim Joyce, who first sparked my interest in analytic philosophy, and who has been an advisor to me since. I met with him a year ago with an idea of a project I was considering working on, and with his support that became my undergraduate thesis. Much of who I am today has been influenced by Jim, and I derive a lot of inspiration from his outlook on philosophy. I could not have had a more well-suited thesis advisor, and I am very grateful. After declaring my honors major, the first step in the formal thesis process was to take Philosophy 401, for which I was lucky enough to have Sarah Buss leading. She, too, has been supportive of me throughout this project, and her challenge that I write clearer has continued to push me to think and rethink the ideas I present. She is also the reason I was reading about duck sex in a hostel in Spain. This class also afforded me the benefit of working with and learning from my peers, and I am especially grateful to Bennett Dombcik, Lucy Malcoun, and Irene Wang. My second reader is Gordon Belot, and during our independent study this semester, I have learned a lot about how to be a philosopher and academic, particularly one equipped with formal tools, and talking with him has helped me feel ready to begin graduate school. Upon beginning to work on the project itself, Jefferey Barrett first helped point me in the right direction and then helped me to think about what could be learned in such a project, and more specifically what I was looking to learn here. Matthew Harrison-Trainor, though I usually refer to him as my "math advisor," has played a significant role in guiding me through the academic world, and if not for his presentation of computability theory, I may not have found a home in formal epistemology. The philosophical questions therein are deeply intertwined with mathematical discovery, and I look forward to continuing my own exploration at Carnegie Mellon. My brothers, Ibrahim and Musa (in order of age not preference), have left helpful comments on this thesis, and Musa helped me while writing the code, but far more than that, they constantly challenge me to be a better and more confident philosopher and person. And on that note, I thank Abigail Peacock, Ammar Ahmad, Andrew Vincent, and Juan Ruiz for constantly inspiring me, making me think, and enriching my lives with their close friendships. Thank you also to my parents for supporting me. There are of course so many others to mention, and I feel strange not mentioning them, but at some point this must conclude and the paper must begin.

1 Introduction

"Names have no predetermined appropriateness. One forms agreement in order to name things. Once the agreement is set and has become custom, then they are called appropriate, and what differs from the agreed usage is called inappropriate."

_ _

"When a single word is sufficient to make oneself understood, then one uses a single word. When a single word is not sufficient to make oneself understood, then one combines words."

- Xunzi, "Correct Naming" (3rd century BCE, trans. Hutton, 2016)

The compositional structure of language seems to define it. Even the Stanford Encyclopedia of Philosophy's article on "Compositionality" (Szabó 2022) begins: "Anything that deserves to be called a language must contain meaningful expressions built up from other meaningful expressions." Using language involves combining simpler symbols in such a way that the meanings produced are functions of the meanings being combined. A basic example of this is "water bottle," which combines "water" and "bottle" to refer to a bottle which contains (or is meant to contain) water. This is opposed to the creation of a new, unrelated word for that concept, like "tob." We do invent new words, and of course all words were invented at some point, but we still compose using these new words. It seems absurd to imagine a single word with a meaning as complex as "the water bottle I dropped on my morning run yesterday." But in most cases, the reasons for which we might choose to compose extant terms, instead of inventing new ones, are unclear, and it seems difficult to point to general rules which inform this balancing act between composition and invention. Some obvious contenders are how "important" a concept is, maybe corresponding to how frequently it is used, and how difficult its meaning is to express. When the importance or difficulty is above some threshold, then invention may be more ideal. That is, there are cases in which composing seems more "expensive," and so we should invent; an example of this could be an important, oft-referred-to concept, difficult to describe using simpler vocabulary, like "water." And then there are cases in which inventing seems more "expensive," and so we should compose; why remember a new word just for that one water bottle I dropped? It seems obvious that we can't—at least, we certainly couldn't in all cases, given our cognitive constraints. So, in understanding this balancing act, perhaps it would be fruitful to ask what would happen if we loosened these constraints.

If we could invent a new word in all cases, with no costs in terms of, e.g., storing them and their meanings, perhaps we would never compose, and so we would not have language at all. We claim that in this case, there would be no evolutionary[2] pressure that would result in the compositional complexity, and perhaps richness, of language as we know it. This is because an agent, subject to some form of selection, should always opt to communicate as much as useful while incurring the least cost. Thus, with no bounds on either the information contained in individual words or the number of words an agent has access to, single words will always maximize payoffs; using more words to communicate will not increase the information communicated and will cost, at the very least, more time.[3] But we know such restrictions exist: we can only know and readily use so many words, and the creation, dissemination, and standardization of new terminology is costly. Thus, we seek to provide a formal argument for the notion that these limitations inform our tendency to compose.

Let us briefly state some questions which summarize what we are interested in, to narrow down our focus and thereby find and develop an appropriate toolkit. What drives the evolution of natural language towards a compositional structure? How complex do agents need to be in order to evolve compositional meaning (e.g. do they need to be rational)? If the invention of new words were "cheaper," so to speak, would we still have evolved to compose?

This last question suggests that we need to work in an environment where we can formalize the notion of cost. That is, in order to support the claim that the difficulty of invention may inform our tendency to compose, we must introduce costs for invention and composition, and make sure they are comparable; we should be able to "cheapen" one of the two and see what changes this affects. In this paper, we will formalize cost in the context of a game-theoretic model. Then, we will simulate the model and study it. In particular, we will look to signaling theory, the study of signaling systems and their evolution within signaling games, which can be seen as generalizing linguistic communication. There is already literature exploring what meaning, invention, and compositionality each look like within this context; we will seek to add a mechanism of invention to a compositional model in order to probe their simultaneous

[2] Evolution here refers both to biological and cultural evolution.
[3] In the words of Kevin Malone of *The Office*, "Why waste time say lot word when few word do trick?"

development. Our model will involve agents which need to communicate about different "states of nature," and we will be interested in identifying factors that drive such a model towards inventing as many signals as can be useful or inventing fewer signals which are then composed. The question we ask is: What will agents tend to do in situations where they must make trade-offs between costly linguistic invention and adding costly compositional structure to language? We are particularly interested in the case of increasing the cost of invention, and its effect on the evolution of compositional behavior. Furthermore, if compositional *meaning* can truly be said to arise during the life of the model, we will have also shown that the model's assumptions are sufficient for its evolution.

In this thesis, we will begin by justifying the use of models in the study of general philosophical issues in §2, where we argue that models are a type of *plausible account*. This is followed by a presentation of some history of the study of evolutionary models along with some useful terminology in §3.1. Then, we will motivate the use of models in studying human behavioral norms, such as justice, in §3.2. For the remainder of §3, we present some signaling games and investigate what meaning, compositionality, and invention can mean within them. In light of this setup, we will present a new signaling game, useful for our investigation, in §4, and summarize some of its behavior on simulation. Finally, we conclude with discussion in §5, where we also suggest some possible future work.

In short, the goal of this thesis is to explore the simultaneous development of the grammatical structure and vocabulary of natural language through the simplifying lens of the evolution of signaling games. We will argue that compositionality requires minimal assumptions to evolve (i.e. that it can evolve in a relatively simple set up), improving on recent studies. Then we will explore, in the context of the model we build, the extent to which the evolutionary pressure toward compositionality depends on how easy invention is.

2 Modeling as a Philosophical Practice

In exploring a question of the form: "how did x happen," we often find ourselves facing a barrier where what very well *could* be an empirical question cannot be approached empirically. Consider an example question: "How did life on the planet Earth evolve?" Surely to an all-grasping, long-lived viewer, this question could be fully answered in a finite number of descriptive statements, stating exactly what occurred on every relevant level, perhaps even including the causal relationships between each event. But we do not have access to all of this information, and, even if we did, this full descriptive account lacks much of what philosophers case about. For example, mere information about how life evolved fails to consider both why it evolved, if there is such a "why,"[4] and if there is some feature of our universe that necessitates its evolution.[5] Further, our experience of life does not necessarily constrain what life could be, and so even the question of "what is life?" may not be more well-informed by mere facts about what happened to have occurred in our case. Instead, exploring more general questions, such as "how could life have evolved?" or "how does life evolve?" might serve to further our understanding of what "life" is, and to formalize and clarify our intuitive ideas surrounding it. Finally, such questions, given our limitations in approaching a merely empirical account, form a far more reasonable project.

In this spirit, we recognize that questions about how language evolved *could* also be approached merely empirically, but given our interest in philosophical issues (such as: Is language somehow essential to human behavior, or innate to humans? What is "meaning"? Or, the questions we ask in §1) combined with our epistemic limitations, we instead ask "how could language evolve?" In approaching this question, we will look to *plausible accounts*. By "plausible accounts," we mean that any useful explanatory account must (a) be shown to actually, or at least convincingly, describe the phenomenon under question and (b) be plausible

[4] The existence of life might, for example, be claimed as evidence for something like intelligent design, serving as such a "why," but an account of life which explains its arisal in terms of random interactions subject to reinforcement, increasing in complexity over billions of years via evolution, deflates such a claim.

[5] These questions in turn help inform our understanding of, e.g, our place in the universe and the chance of life having evolved elsewhere.

given the empirical facts we do know.[6] Under this view, there may be an account which is at one point plausible, but through further exploration (perhaps done under the assumption that the account is true) we gather evidence which falsifies it, and so the account is no longer plausible. An example of this is Newtonian mechanics, which failed to predict the perihelion precession of the planet Mercury. A theory which plausibly (and accurately) described how planetary bodies moved later failed in application to a newly relevant (and indeed newly observed) circumstance, and physics had to account for that. Until issues such as the precession were uncovered, however, Newtonian mechanics met both (a) and (b) by accurately predicting the motion of studied objects assuming only a reasonable set of laws of nature.

Plausible accounts are inherently worth investigating and discussing, regardless of whether, for example, there are multiple contradicting plausible accounts, all serving to explain the same phenomena. The fact that, for example, language *could have* evolved in a particular way teaches us something about language, even if that may not be how it actually evolved. It may therefore be helpful to differentiate among a hierarchy of questions that can be asked about some phenomenon x:[7]

1) how could x have possibly happened?
2) how could x have possibly happened, with minimal assumptions?
3) how could x have plausibly happened?
4) how did x happen (actually)?

The first three questions are the domain of plausible accounts, whereas the fourth is an empirical question, rendering it difficult to investigate. Thus, models of any sort, no matter how useful, will quickly lose credibility upon approaching the fourth question. For example, the point-mass assumption, which (accurately) calculates motion under the premise that objects are fully concentrated at their centers of mass, is false, no matter how useful or accurate in modeling motion. We submit that all four questions, including the first, should be of interest to philosophers.

---

[6] In other words, (a) assumes the account's axioms and asks if reality coincides with its predictions and if what the account predicts amounts to the target of our investigation, and (b) asks if the theory's axioms are themselves reasonable (e.g. they don't assume too restrictive or complex of a set up so as to not be plausible). In an evolutionary model used to study meaning, (a) asks whether the model's predictions are realistic and if they amount to meaning, and (b) asks if the model's initial state and degrees of freedom are reasonable to assume.

[7] I thank Jeffery Barrett for helpful discussion and ideas about this.

To show that even the first question can be fruitful, suppose someone were to stipulate that natural language requires some strong condition, say rationality, in order to arise. If a model can show that meaning can evolve without assuming rational behavior, then either that stipulation was wrong, or the stipulator must clarify that natural language involves more than just evolved meaning; perhaps it needs the compositional structure we view in language, and it is this in turn what requires rationality. Then, a model can be (and has been)[8] created to show that a form of composition may too evolve given relatively minimal assumptions, and so the stipulator must once again clarify what natural language involves, or what they meant by composition and how what the model showed does not suffice. And so on. In this way, we can both formally narrow down what intuitive concepts (such as natural language) are, and better understand their relationships to other ideas such as meaning.[9] The use of simulated models can be thought of as a particular kind of thought experiment, one aided by the use of computer technology.

Of course, objections can be raised during this process, and discussions can be had about whether these models are actually plausible accounts or whether they somehow come imbued with what they seek to show evolve,[10] but throughout the history of philosophical modeling we have learned a lot, and, as the practice matures, we stand to make significant advances in our understandings of evolved phenomena.

---

[8] This is due to Barrett, Cochran, and Skyrms, and that result is strengthened in this thesis.

[9] Formalizations, and attempted formalizations, of intuitive concepts form a class of philosophical studies with a rich history and which deserves its own long-form discussion. To realize that these investigations not only can be but have been fruitful, one must look no further than the Church-Turing thesis of computability theory, for example. The next section begins with an important historical evolutionary model.

[10] This latter point will be discussed in §3.4.

3 Background

3.1 The Evolution of Strategy

In 1973, John Maynard Smith and George Price published “The Logic of Animal Conflict,” which would quickly cement itself as a cornerstone in the literature of behavioral evolution. Motivated by the counterintuitive nature of male animal combat, which tends to be ritualized or otherwise limited, they investigated whether a strategy that benefits the survival of a population, but less obviously the individual themself, can propagate through individual natural selection. To illustrate the issue, consider male snakes, which generally only wrestle their competitors without using their fangs; it seems that if the victorious males were to instead outright eliminate their competition, they would be selected for much more heavily. Conventional explanations had defaulted to group selection, where groups whose members exhibit some behavior might be selected for over differently behaving groups, thereby reinforcing the presence of that behavior in the overall population. Under this view, their central example could be explained by the idea that if the male animals in a group went all out, the survivors would likely be severely injured, and this might curb the viability of the group as a whole. Groups in which this practice did not occur as frequently would thus be favored. But to Maynard Smith and Price, this group-level account did not seem sufficient to explain the ubiquity of what they call “limited war” over what they call “total war.” In particular, there is significant skepticism about group selection as a mechanism of evolution, arising especially in the decades prior to the publication of the paper being discussed. This is because natural selection seems to occur at the genomic level. Unless changes are made in the genome, they, even if beneficial, will not be passed down to subsequent generations. Thus, explanations which rely (at least primarily) on individual selection are to be preferred. With this in mind, Maynard Smith and Price introduced the notion of an *evolutionarily stable strategy* (ESS), and argued through the use of simulation that limited war is evolutionarily stable, while total war is not. By using the set up modeled in the paper as a case study, they show that individual selection can indeed explain the success of limited war strategies.

The notion of an ESS hinges on there being different, interacting strategies. In informal terms, an ESS is a behavioral strategy that cannot be outperformed; i.e. no other strategy performs better when played against it.[11] They are considered "stable" because once an ESS arises in a population, it will never be at a disadvantage when interacting with any other strategy, and will therefore not be selected against. But a strategy is not an ESS *per se*; instead, a strategy is only an ESS relative to the other strategies featured in the population and how the members of the population match up.

For simplicity's sake, we will assume that individuals interact with random other members of the population (i.e. there is no bias such as spatial proximity).[12] The table below illustrates the possible interactions among three distinct strategies, A, B, and C (a similar table can be constructed for any set of interacting strategies). The entries are numerical values which generalize the idea of *fitness*, summarizing all payoffs and consequences of playing one strategy against another, with the following notational convention: for any B and A, $\text{fitness}_{B, A}$ is the amount of fitness an adoptee of strategy B gains in an interaction with an adoptee of strategy A. A player of a more successful strategy will gain at least as much fitness as its opponent, and so will be at least as successful in reproducing, causing their share of the population to grow at least as much as the opponent's by the next generation.

	A	B	C
A	$\text{fitness}_{A, A}$	$\text{fitness}_{A, B}$	$\text{fitness}_{A, C}$
B	$\text{fitness}_{B, A}$	$\text{fitness}_{B, B}$	$\text{fitness}_{B, C}$
C	$\text{fitness}_{C, A}$	$\text{fitness}_{C, B}$	$\text{fitness}_{C, C}$

An ESS is defined as a strategy against which no strategy performs better. For example, if $\text{fitness}_{B, A}$ and $\text{fitness}_{C, A}$ are both not greater than $\text{fitness}_{A, A}$ (equivalently, there is no value in A's column greater than $\text{fitness}_{A, A}$), then A is an ESS.

[11] In the literature, this is often worded as: if strategy S is an ESS, then no mutant strategy can invade a population of S players (i.e. successfully "take over" some portion of the population).
[12] This is a significant assumption, though it serves to vastly simplify the space of possibilities.

To motivate this definition, consider populations whose members primarily adopt one strategy, say A. If no other strategy has a higher payoff than A when paired against A (illustrated in the table: if $\text{fitness}_{B, A}$ and $\text{fitness}_{C, A}$ are both not greater than $\text{fitness}_{A, A}$) then A is an ESS. Any adoptee of a different strategy will be more likely to interact with an A player than anything else, since they form the majority, and if that strategy does worse against A than A itself does, it will have lower average fitness than A, which also mostly interacts with A, but performs better. This different strategy will thus not be selected for, as this difference in payoffs will be reflected in each strategy player's ability to reproduce and, in turn, the strategy's share of the next generation's population. Now suppose A is not an ESS, and another strategy does better than A against A. Then its adoptees will be selected for until it replaces A as the majority, and so it makes sense to think of A as unstable. In this way, populations will evolve so that a majority adopt an ESS, if there is at least one present.[13]

Having laid out this architecture, Maynard Smith and Price created a simulation involving three actions, corresponding to limited war tactics, total war tactics, and retreat. They specified the payoffs of successful actions and the probabilities of each action being successful, chosen in such a way they thought reasonable based on how costly certain behaviors seemed relative to each other.[14] They then constructed five behavioral strategies involving these actions, where one of them was a total war strategy and three of them were limited war strategies. These strategies included extremes: the total war strategy always used total war tactics, and one of the limited strategies never used total war tactics.[15] They simulated each pair of strategies interacting and retrieved the average payoffs for both strategists over many simulations. In the end, two limited war strategies clearly stood out as ESSs, and the total war strategy was not stable.[16]

[13] It may be the case that there are no evolutionarily stable strategies—consider a set of three strategies wherein each is tied with itself, performs better than a second, and performs worse than a third (in a way similar to the moves of rock-paper-scissors).

[14] They also claim to have tested different payoffs and probabilities, which they did not report; however, they claim that the outcomes of the simulation were the same while the parameters were close to the numbers they did report.

[15] In summary, the total war strategist only used total war tactics and continued doing so until their opponent retreated or until total war tactics were successfully used against them. One of the limited strategists only used limited war tactics and would retreat immediately upon any escalation, whereas the other two limited strategists would retaliate when their opponent escalated. What differentiated these latter two is that one would, with some low probability, escalate, but back down if their opponent escalated as well. The fifth strategy, which was of less interest to the authors, escalated by default, always played the opposite of whatever their opponent did, and retreated upon the second total war tactic used by their opponent.

[16] Notably, drastically increasing the probability of success for total war tactics made the total war strategy an ESS.

"Briefly," in the words of Maynard Smith and Price, "the reason that conflict limitation increases individual fitness is that retaliation behaviour decreases the fitness of [the total war strategists], while the existence of possible future mating opportunities reduces the loss from retreating uninjured."[17] Thus, a population with strategies analogous to these five will evolve so that the majority will adopt limited war strategies, even if most start off playing the total war strategy.

Though a simple model, these results show that, when considering strategies within a population, individual selection can be what drives the evolution of behaviors such as ritualized male combat in animal species. But Maynard Smith and Price also did something more subtle: they set a precedent. They made a strong case for the plausibility of an explanatory account (individual selection) by building a convincing model and showing its evolution to be in line with what the account predicts.

3.2 Two Ways of Splitting Cake

Maynard Smith and Price's method of studying the evolution of strategy through the generalizing lens of a model worked for animal behavior, but it seems that human behavior is far more complex, in that we generally assume human decision making to be, at least to some extent, rational. But humans tend to follow certain behavior-governing norms. So, is following these norms also rational, and why? In particular, we may explore whether the evolution of human strategy can be explained using the architecture set up in §3.1 by investigating whether the norms we follow are special among the set of possible strategies. We are far from the first to ask this—since Maynard Smith and Price's paper, other authors have sought to generalize the results to humans. For example, Brian Skyrms, in his 1996 book *Evolution of the Social Contract*, argues that the notion of the ESS might help explain the evolution of the concept of justice. This can be illustrated through two archetypal examples of well-studied social games which involve sharing resources. For both of the following games, the resource in question will be a cake, a deeply valuable resource in the eyes of two agents between whom it will be split. These examples are especially potent since resource sharing is something humans often do, and something which may involve rational decision-making. Will the results of studying the

[17] Maynard Smith and Price, 16.

evolution of strategy within these games agree with what we expect, and indeed observe, from rational human players?

Consider this first game, which we will call "Cut the Cake."[18] Here, both agents privately choose a percentage of the cake, which they will announce simultaneously. If the sum total is 100% or less, then both agents receive their stated amounts, while the remainder is discarded, and if the sum total is greater than 100%, then the entire cake is thrown away. Though both agents have motive to get as much cake as possible by stating 100%, it is clear that this will be deleterious to both. Since each agent must act independently, with no way to know what the other agent will do, it seems that they must play it safe by choosing a percentage that would result in a high probability of them getting cake. But the cake is valuable to both agents, so neither wants any going to waste. In general, a 50%-50% split of the cake seems like the most "just" strategy, or maybe the most efficient or prosocial one, but is this choice at all special in our game? Is there any reason it should be the rational choice, and, in particular, does it somehow maximize payoffs or fitness?

In game-theoretic terms, every combination of percentages which adds to 100% is actually a *Nash equilibrium*, meaning that neither agent can change their strategy (while holding the other's strategy constant) and perform better. Take a 70%-30% split—if either agent moves up or down, they play less optimally; by moving up even slightly they lose everything, and by moving down, they lose the amount they moved by, and that amount goes to waste instead. But are any of these waste-less strategies special in the context of an evolving population of interacting agents with fixed strategies? That is, consider a population where when agents interact, they play Cut the Cake, and the take-home amount corresponds to the fitness gained from the interaction. How should we expect such a population to evolve?

Let us consider one specific strategy: Suppose a majority of some randomly-interacting, reproducing population demanded 70%. Then most interactions will result in a payoff of 0%, and aberrant strategies playing 30% or less will do better than average, leading them to be selected for. Once those demanding 30% or less become a majority, the average payoff will be around 30%, so strategies playing more than 30% but less than 71% percent will have a higher average

---

[18] This game is as presented in Brian Skyrms's 1996 book *The Evolution of the Social Contract*.

payoff. This will keep oscillating back and forth until the majority play 50%. Playing 50% is evolutionarily stable, since when it is the majority, no aberrant strategy can do better.

The fact that the strategy of playing 50% is expected in the evolution of this system might be one reason why we humans have evolved to prefer it, even believing that it ought to be the strategy employed;[19] Skyrms suggests that our sense of justice could be explained as an evolved (though much more complex) strategy. That is, human agents who evolved a sense of justice, and in particular, a behavioral norm that results in splitting the cake correctly, will have higher average payoffs in such situations. This, combined with the fact that situations following this general pattern are quite common, results in an evolutionary pressure for individuals to follow the norm. Furthermore, when these individuals fail to act in accordance with these norms, they act in a way that is harmful not just to themselves, but to others, and so there is also a social pressure to act so that payoffs are maximized for the group. There are, however, games in which the socially correct choice is less clear, and where the evolutionary dynamics are less predictive.

As a second example of modeling human behavior game-theoretically, take the Ultimatum Game.[20] Here, the first agent announces a percentage of the cake which desire, and the second agent must then choose between taking the remainder and throwing it all away. Here, when the game is played a single time (the "one-shot" game), there is a clear rational strategy for the second agent, which is to take whatever they're offered. This will always result in a higher payoff for, and since the game is only being played once, there is no potential future benefit to throwing the cake away. And so it seems that humans ought to act in this way when playing the Ultimatum Game if they act rationally. This, however, is not what is observed: the first player will usually offer far more than they could get away with assuming rational behavior by the second player, and the second player often rejects such offers.[21] There are a number of possible, potentially interacting, explanations for this. Firstly, it may simply be that humans do not act rationally in such situations. Or, there are other objects of value to the players, and perhaps their strategies are not so stable; perhaps they value cake, but when presented with an insultingly low offer, their desire to preserve their honor prevails. There is, however, a third explanation, which

[19] In the sense that a different strategy would be unjust.
[20] I do not have a direct source for the version presented, though "On the rationality postulates underlying the theory of cooperative games" by John Harsanyi (1961) is an early account of it in literature.
[21] See "Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis" (2004) by Osterbeek, Sloof, and de Kuilen.

is that human players, by default, see themselves as part of an iterated game. That is, perhaps they only act incorrectly because they are in a one-shot game, and if we change the set up so that they repeatedly play the game, throwing the whole cake away may sometimes be rational.

Indeed, in the iterated game, the first agent might be motivated to make a "fair" offer, since the possibility of rejection for an "unfair" offer might, over many plays, push the average payoff of demanding a high percentage lower than the average payoff of stating a "reasonable" amount.[22] Thus, for an agent faced with a bad choice, being willing to hold out might be a better long-term strategy than merely taking what one can get, and so this behavior is not irrational. But even when the iterated game is played by humans, there does not seem to be any consistent notion of what a "fair" offer ought to be. Because of this, the Ultimatum Game has been used as a study of cross-cultural or geographical ideas of fairness, where different groups of people might have different average cutoffs for what they are willing to accept from the person making the offer.[23] But this prompts the following line of questioning. We saw in Cut the Cake that the strategy which tends to be played by humans is also the ideal strategy in an evolutionary context, in that it is an ESS. Why do human players not tend to consistently perform some strategy which is ideal in the Ultimatum Game, such as an ESS?

In an evolutionary context with randomly-interacting agents who play the ultimatum game, with their take-home percentage of cake representing their increase in fitness due to the interaction, there is no ESS.[24] *Prima facie*, it might seem that offering a 50-50 split is the "fair" choice, but let us consider a population where a majority will offer 50% of the cake in case they are the first agent, and they will only accept an offer of 50% or more in case they are the second agent. Call this strategy A. Then, the strategy which also offers 50% but accepts any offer made to them, call it B, has the same payoff against A as A, and higher payoffs against players who offer less than 50% than A, and so it may be introduced into the population. Then, once enough of these B-players are present in the population, greedy strategists, who offer low amounts, will be able to perform better than either A or B on average. And so, unlike Cut the Cake, it is unclear that a population will tend toward a particular norm for how this game should be played. For

[22] "Fair" and "reasonable" are in quotes because it is unclear what these words mean across all cases.
[23] Again, see Osterbeek, Sloof, and de Kuilen (2004).
[24] See, for instance, "The Indirect Evolutionary Approach to Explaining Fair Allocations" (1999) by Huck and Oechssler.

example, one might argue that the agent making the offer might "fairly" deserve a slightly higher payoff, given their position of authority.

These two games illustrate, and thereby provide a means of studying, human behavior. Though simple, they have natural analogies in human society and they can be played in experimental set ups to measure how people generally act. By studying them, we can learn about why we evolved our norms, such as fairness. This might extend to other facets of human behavior, including more general social contracts, and even, counterintuitively, injustice. For the purpose of our study, we will consider whether it extends to language.

3.3 Signaling Games

Having outlined motivations for the practice of studying games at all, we consider a new class of games which might help explain the evolution of both human speech and animal calls, along with far simpler phenomena such as chemical communication between cells: signaling games. In the simplest set up,[25] we will consider two agents, two states of nature, two signals, and two actions. For each state of nature, there is exactly one correct action, which, when performed given that state of nature, results in payoffs.

The set up of the game is as follows. One of the agents is the sender, and one is the receiver. The sender views the state of nature, and the receiver does not. The sender randomly chooses and sends one of the signals upon viewing the state of nature,[26] and the receiver randomly performs one of the actions upon receiving the signal. There is a payoff when the action of the receiver corresponds to the state of nature (i.e. the receiver acts appropriately, given the state of nature), and this payoff reinforces the correspondence between the state of nature and the chosen signal and between the signal and the chosen action. Having acted properly, the sender will be more likely to choose the signal it sent on this iteration the next time this same state of nature occurs, and the receiver will be more likely to act as it did this time the next time this signal is sent.

The form of reinforcement we will consider is modeled using balls in urns. The state of nature is always chosen randomly, with equal probability in the simplest set up,[27] and the sender

---

[25] Though originally introduced by David Lewis, I largely draw from Skyrms's 2010 book *Signals*.
[26] The word "chooses" will be used throughout. In this context, it means nothing more than "randomly selects."
[27] Changing this probability can drastically change the evolution of the signaling system.

has an urn for each state of nature. The balls in these urns correspond to the signals, and the game begins with one ball for each signal in each urn, and they are thus equally likely to start. The receiver, in turn, has an urn for each signal, and these contain balls for each action, also with each action initially being equally likely given either signal.[28] Whenever an action which matches the state of nature is chosen, a ball for the chosen action is added to the receiver's urn for the chosen signal, and a ball for the chosen signal is added to the sender's urn for the current state of nature. This urn model, often dubbed *Pólya's urns* for George Pólya, serves as a generalization of phenomena such as reinforcement learning or natural selection, where correct actions or beneficial mutations are, for whatever reason, more likely to occur again.

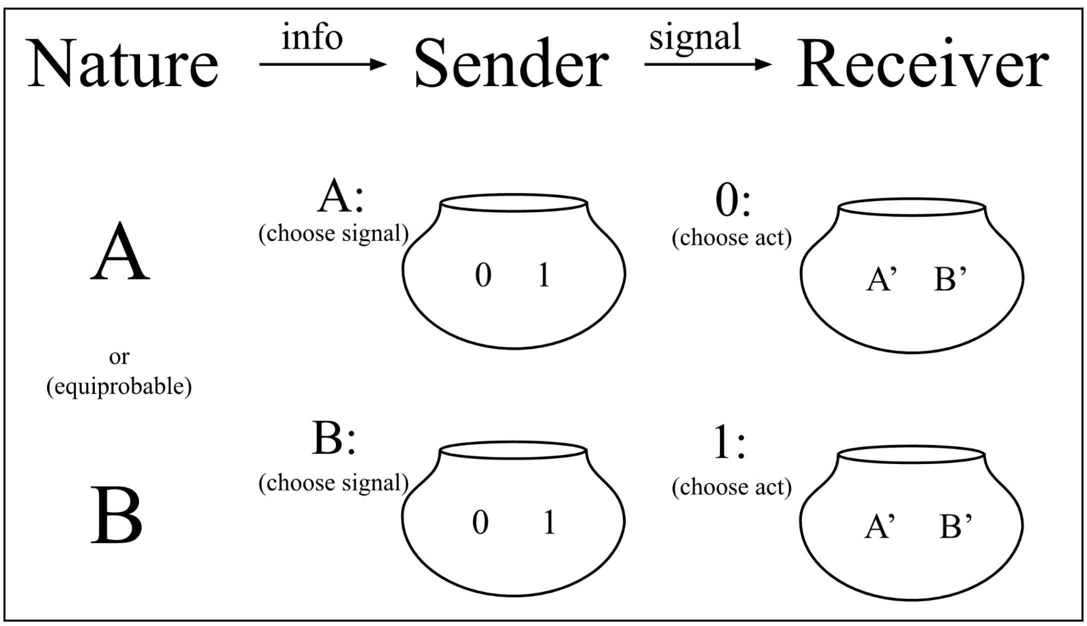

Figure 1: An illustration of the basic signaling game. If the state of nature is A, then the sender randomly chooses a signal, represented by a ball, from its A urn. If the sender chooses and sends the 0 signal, then the receiver randomly chooses an action from its 0 urn. If the state of nature was A and the chosen action was A', then the agents acted correctly, and reinforcement takes place (see Figure 2).

To illustrate an iteration of the game, suppose the states of nature are A and B, the possible signals are 1 and 0, and the actions are A' and B' (Figure 1). If, on a play, the state of nature is A, the sender chooses 0, and the receiver chooses A', then the receiver acts correctly

[28] Changing the number of balls corresponding to each signal/action initially present in their respective urns can also drastically change the evolution of the signaling system.

and reinforcement takes place: a 0 ball is added to the sender's A urn and an A' ball is added to the receiver's 0 urn (Figure 2). If the state of nature and the action do not match up on a play, then, in this simplest set up, nothing happens (i.e. there is no notion of cost or negative reinforcement).

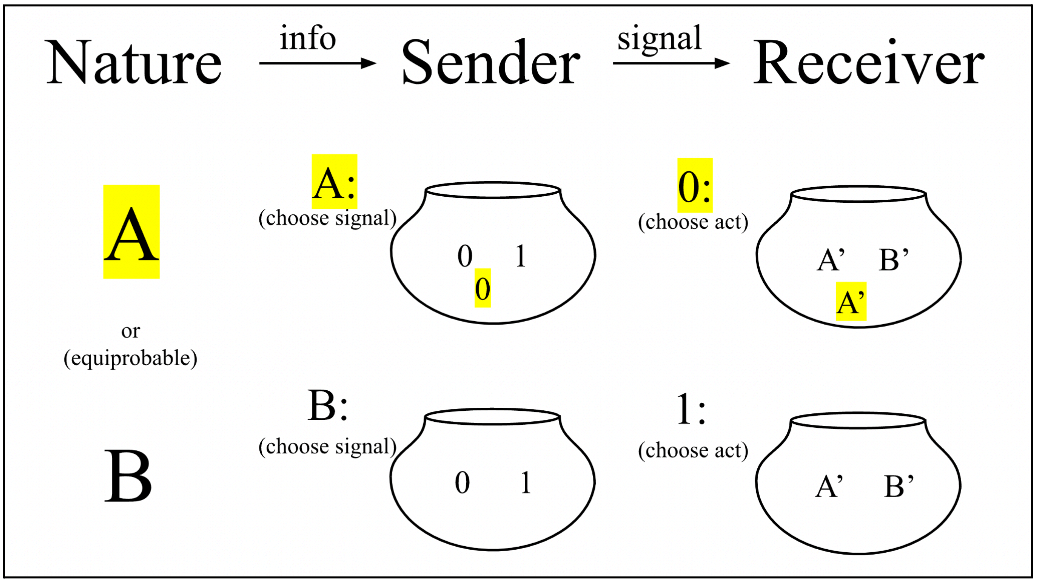

Figure 2: The basic game after one round of reinforcement. Here, the state of nature was A, the sender chose signal 0 from its A urn, and the receiver chose action A' from its 0 urn. Because A and A' match, the agents acted correctly, and reinforcement took place in the form of an additional 0 ball being added to the sender's A urn and an extra A' ball being added to the receiver's 0 urn.

There is much literature discussing this game and variations on it. One important result is that of Argiento et al. in their 2008 paper "Learning to signal: Analysis of a micro-level reinforcement model," where they provide an analytic proof that, with probability 1 (or, almost certainly), the game described above will always ultimately converge to a signaling system where one signal corresponds exactly to one of the states of nature and its corresponding action, and the other signal corresponds to the other state and action pair.

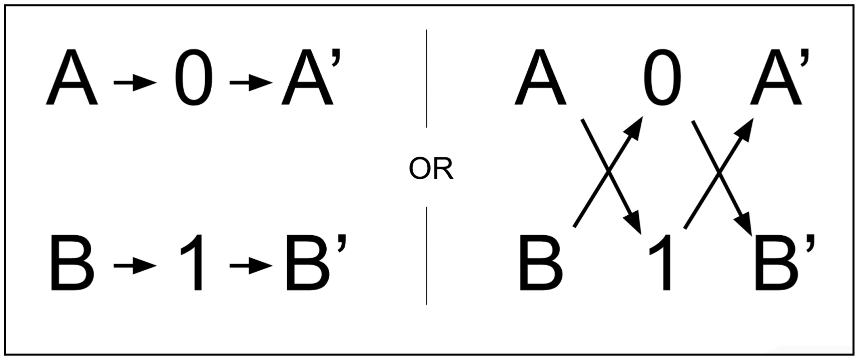

Figure 3: The possible convergences of the basic signaling game. The first case is where the signal A "means" 0 and B "means" 1, and this is reversed in the second case.

3.4 Language and Meaning

Does Argiento et al.'s proof that this game converges to having each signal correspond to a state of nature for both the sender and receiver show that meaning can form from random behavior subject to reinforcement? It seems that this sort of model of behavior is so vastly simplified that it might not, *prima facie*, apply to anything real. And, even if it did, it would still need to be argued that *meaning* really has formed in the process of converging, or, in particular, that what has formed during the evolution of the game really amounts to meaning. Thus, before delving into more complex signaling games, we must convince ourselves that this model of a basic sender-receiver game is a plausible account of evolved meaning, by showing that it meets conditions (a) and (b), outlined in §2. We briefly remind ourselves that condition (a) asks whether the account describes the phenomenon under question (here, meaning) and condition (b) asks whether the account is plausible given the empirical facts we know.

We will begin with condition (b). In particular, we ask what assumptions the model makes and whether they are realistic. Importantly, the formal, analytical proof by Argiento et al. which shows that the model will converge to one of the signaling systems shown in Figure 3 assumes the setup to be exactly the case of reinforcement with balls and urns, with perfectly random choice among the balls within the urns, and equiprobable states of nature. This is not realistic, and so not immediately persuasive. That is, more general axioms would make the account more convincingly *real*,[29] because the proof might literally apply to real-life scenarios,

[29] Or, plausible as an account of what *actually happened* in the natural evolution of what we call meaning. I.e. it might answer question (4) in §2.

thereby serving to explain how things literally happened to have evolved. But this is not within the scope of plausible accounts, and we remind ourselves that what these models and their evolution teach us is that the explanatory account *can* account for the phenomena we observe (i.e. assuming (a), to be discussed below, is met, this is one way meaning *could have* evolved). And all the assumed setup really presupposes is that behavior is selected randomly and that correct behavior is reinforced, and these two seem ubiquitous within nature. Even if a game may never organically occur exactly as presented, many phenomena that do naturally occur fall under the general pattern of a signaling game. Indeed, randomness exists throughout all taxa of life (such as DNA mutations) and in terms of reinforcement, what more could the evolution of signaling require than natural selection? It might be that any reinforcement and randomness, without being exactly as assumed in the proof, are not a *guarantee* of convergence to a signaling system, but it seems that a signaling system's evolution should at least be likely or possible, due to the guarantee in the ideal case assumed in the proof by Argiento et al. and the high-level similarities between the ideal case and reality.

Now consider criterion (a), and suppose a signaling system has already converged.[30] We then ask if *meaning* really evolved. Thus, we must first determine what it means to talk about the meaning of signals. We endorse, as suggested in literature, that there are two different notions of the evolved meaning: the *content* and the *information*. A good discussion of the notion of content can be found in Skyrms and Barrett's 2019 paper, titled "Propositional content in signals,"[31] where they link the notion of a signal's meaning to why it evolved, or the reason evolution maintains it (i.e. this is the *role* the signal plays—imagine an animal call that corresponds to a particular predator being present). This is what I am referring to as "content," and, notably, this necessitates some amount of interpretability, or understanding of the causes and effects at play, which is something we generally lack when studying evolution. (Consider the long-held belief that the appendix was vestigial; for most of human history, we had no understanding of its evolutionary role.) To consider the content of a signal, we need context about the world the signaling system evolved within, along with knowledge of the system's history.

[30] I take that the output is realistic to be trivial. Consider a simple case of two animals with two different, understood signals for two different natural states (perhaps monkeys with two different calls for two different predators).
[31] See their fourth section.

On the other hand, this is something the notion of information does not require, detaching itself altogether from the question of *why* a signal might exist. Following the information-theoretic approach of the third chapter of Brian Skyrms's 2010 *Signals*, we consider how the relevant probability distributions (here, the probabilities of the actions) change with and without the presence of the signal. That is to say, we consider the likelihood of any possible action after a signal is reinforced in one way and is sent, and compare it to the probabilities with no signal present. Then, we calculate the relative entropy, or Kullback–Leibler (KL) divergence (measured in bits), as a quantitative measure of "how much" information a signal contains. But a number tells us very little about the facts of the signaling system.[32] As an example, consider a signal which changes the receiver's probability of actions A' and B' from 0.5 and 0.5 (equally likely) to 0.1 and 0.9, respectively. This would result in a KL divergence of about 0.531 bits. This, in isolation, doesn't tell us very much. However, if the probabilities instead changed to 0.3 and 0.7, for example, the divergence would be approximately 0.119 bits. This tells us that the signal which changed the probabilities by a greater amount contained more information, and was thus able to convey more meaning to the receiver, though it still tells us nothing about the roles played by the signals, or how they should be interpreted. And so, in this set up where information is concerned with the probabilities of the signals and content has to do with the signals' evolutionary roles, it is clear that both evolve, and both seem to correspond to some of what "meaning" is.[33]

One objection that can be raised, briefly mentioned at the end of §2, is as follows. As Arigiento et al. proved, the basic signaling game always converges to the exact correspondence between states of nature and signals. Thus, even if we consider the convergence to amount to meaning, it may be the case that this meaning comes baked into the model itself. Maybe the meaning is always present in the model, perhaps as an implicitly assumed feature of the model, and therefore cannot be truly said to "arise" within the context of the model's evolution. But the previous discussion gives us a means of response. Initially, all the probabilities are equal, and so

[32] KL divergence isn't always a number, but returning positive infinity doesn't tell us very much either. Also, that won't happen here, since the initial presence of one of each ball means that no probability will ever be, after any finite amount of time, zero. This gives us a good sense that KL divergence is the correct measure of information in this context.

[33] It seems that even in common speech we use the word "means" to refer to both of these ideas. If a fire alarm rings, that both "means" there is a (possibility of) a fire, and it "means" that our behavior ought to reflect that (and so the probability distributions among our possible behaviors changes).

the KL divergence is the same with or without the un-reinforced signal. And so these signals both initially contain no information, and also play no particular or unique roles, and thus do not differ in content. Then, reinforcement carries each signal to a meaning, in an un-predetermined way. That is, that either signal can end up correlating to either state of nature (see Figure 3).[34] There is no meaning to start, but reinforcement drives the evolution of the meaning during the play of the game.

In general, we have a far from perfect understanding of signaling systems, and where exactly they apply. However, they serve as a reasonable description of how simple signaling behavior can be selected for in the presence of randomness and reinforcement. But what about something less simple? Would a signaling game be useful in studying complex human language, with its grammatical structure and extensive vocabulary?

3.5 Compositionality

As discussed in the introduction, compositionality seems to be a central feature of natural language. Thus, if we are to argue that signaling games model a generalization of language and its evolution, we must probe what compositionality would look like within the scope of a signaling game. In particular, if it were clear that compositionality could not evolve in a signaling game, that would make signaling games bad candidates for this study. In order to do this, we must first consider what sort of signaling game can be said to involve compositionality. The models described in this section follow Barrett, Cochran, and Skyrms's 2020 paper titled "On the Evolution of Compositional Language." We will begin with an example of a signaling system which, *prima facie*, looks like compositionality, and through an exploration of its failures, we come to a better sense of what such a signaling game must look like.

Suppose we wanted to introduce compositionality to a game like that described in §3.3. A naive solution might be to simply involve two senders, to introduce different signal combinations. For example, consider a set up with four states of nature, two senders with two signals each, and one receiver with four actions. Each time the senders view the state of nature, they each send one of their signals, and the receiver acts based on the combination of signals

---

[34] Of course, the actual correlation only occurs in the limit, but the system can get arbitrarily close to convergence in finite amounts of times, and the reinforcement occurs rather fast.

received (see the example in Figure 4). This seems like compositionality since the signals from each sender are indeed being composed on each play. However, the signals evolve no meaning independently of each other, since they are only ever sent in parallel, and so this model is in fact no different from one with a single sender with four signals (Figure 5). That is, we may view each double-signal as a single signal with probability 1/4. And there is clearly no compositionality in the latter case, which tells us that the former case fails in this regard as well. Thus, by coming up with an example that seems like compositionality but lacks something, we come to a benchmark for compositionality: that the composed signals are somehow functions of the meanings of the composing signals. Barrett, Cochran, and Skyrms realized this, and created what they call a “hierarchical model” (Figure 6) involving *contexts* which drive a more convincing account of compositionality. In particular, unlike the model shown in Figures 4 and 5, the hierarchical model’s senders’ signals will evolve meanings independent of the other senders, and this model cannot be represented using only a single sender.

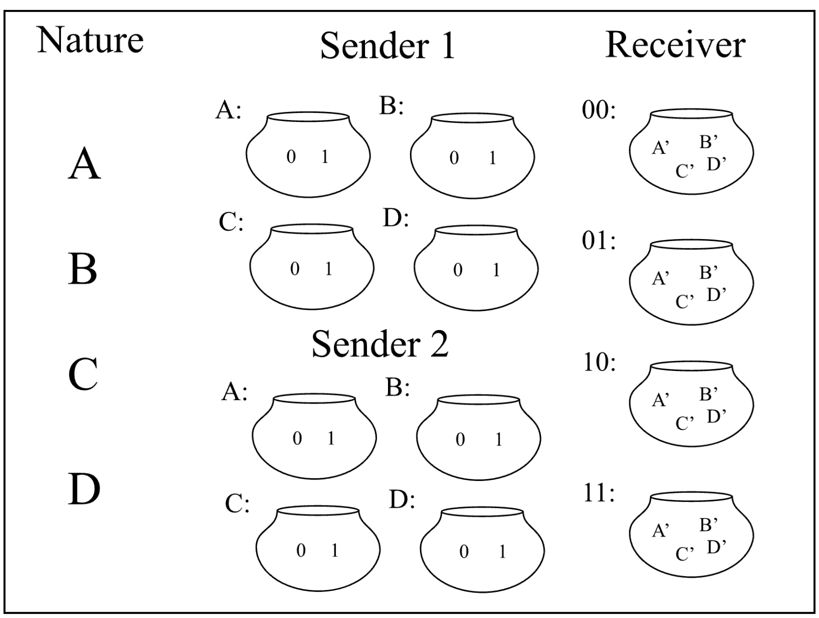

Figure 4: A game which involves composing signals but not compositionality. This game is played exactly as the basic signaling game, instead with two senders who each, simultaneously, send a signal upon viewing the state of nature. The receiver then chooses its action based on both signals.

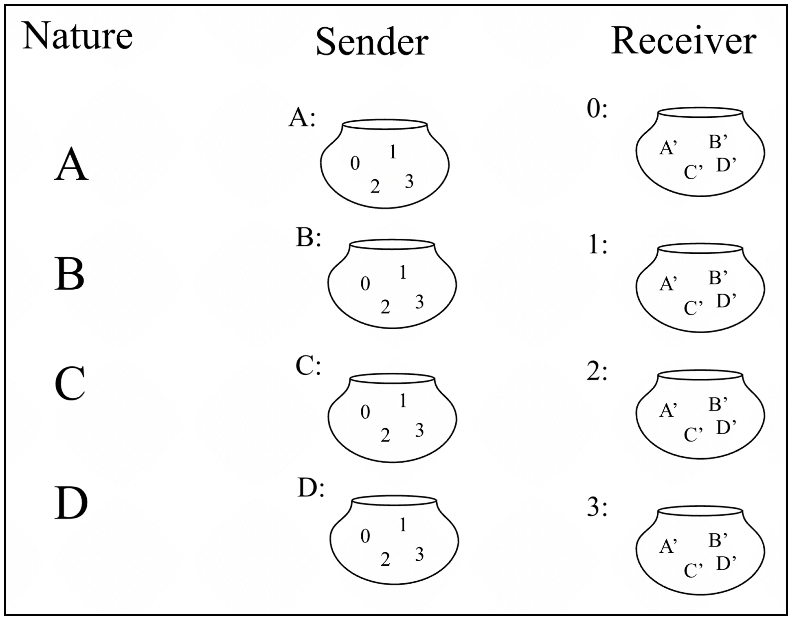

Figure 5: An equivalent game to that in Figure 4 represented with only a single sender. This shows us that the game in Figure 4 lacks compositionality.

The setup of their hierarchical model (found in their fourth section) is as follows. They consider four states of nature, each with two components (contexts); the illustrative example they use is color, animal, or both for contexts, with possible states of black cat, black dog, white cat, and white dog. In their set up, the basic receiver's action needs to match the state of nature and context by painting one of: black, white, cat, dog, black cat, black dog, white cat, or white dog on a blank canvas. There are five agents in their set up: an executive sender and receiver, two basic senders with two signals each, and a basic receiver. Just as before, every choice the agents make is represented by urns containing initially equinumerous balls representing each possible outcome. The executive agents do not represent behavioral roles in the game *per se*; instead, they represent the possibility for the basic agents to signal more or less depending on what is relevant, as will be explained below, and thus play "higher-level" (so to speak) roles in the game than the basic agents do. This is why the model is referred to as "hierarchical."

At the beginning of a play of the game, the state of nature and the context are chosen equiprobably at random, the two basic senders see the state, and the executive sender sees the context. Given the context, the executive sender determines which basic sender, or both, will send on this play, and the chosen basic sender(s) will pick the signal they'll send given the state of nature. The basic receiver then chooses an action based on the signals sent. If only one basic sender sends on a play, then the basic receiver randomly chooses one of the two urns corresponding to the signal that was sent (i.e. if the first basic sender has signals A and B and the second basic sender has signals C and D, then the basic receiver's urns are AC, BC, AD, BD. If only one basic sender sends, and they send, say, D, then the basic receiver chooses between AD and BD randomly). The basic receiver's actions need to be clarified by the executive receiver, who sees which basic senders sent and determines whether the basic receiver acts based on one context, and which, or both. Finally, the basic receiver acts based on their choice and that of the executive receiver. As an example of this, if the basic receiver chooses the action corresponding to white dog, and the executive receiver chooses that the action is to be interpreted as a color action, then the basic receiver will paint the canvas white.

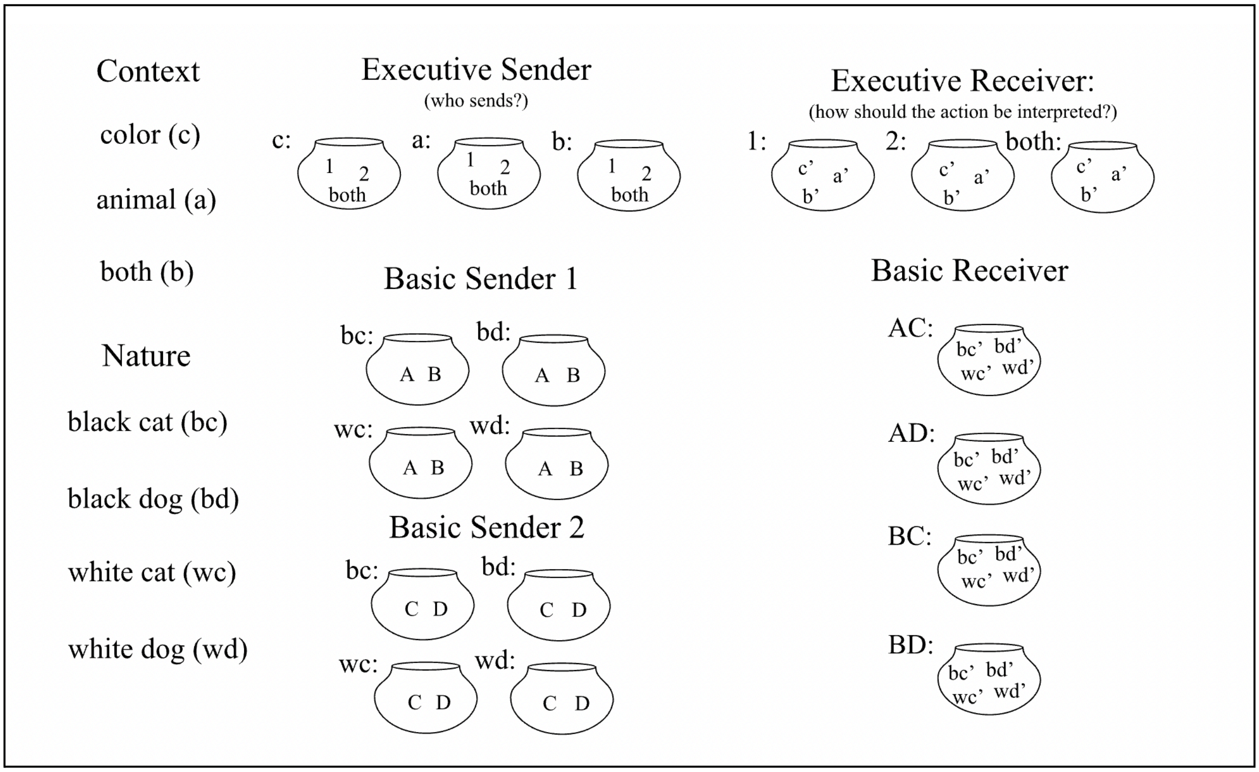

Figure 6: The initial state of the hierarchical model. The context and state of nature are chosen equiprobably at random. The executive sender chooses which basic senders will send on this iteration based on the context. The chosen basic senders choose which signals they will send given the state of nature. The basic receiver chooses an action based on the signal(s) sent. If only one signal is sent, the basic receiver will choose randomly between the urns corresponding to the sent signal. The executive receiver sees which basic senders sent and determines how the basic receiver's action should be interpreted. For example, if the basic receiver chooses wd' and the executive receiver chooses c', then the action performed is w', which is correct if the state of nature was wc or wd and the context was c.

If the action matches the state of nature and context, then there is a positive payoff. Barrett, Cochran, and Skyrms allow the payoff to be higher for the "both" context than for either of the simpler contexts.[35] In this model, there is also a cost to signaling, serving to generalize trade offs like time or cognitive costs that come with composing more than one signal. Since the payoffs are reinforcement balls added to urns, the costs are represented as lower payoffs via the removal of the very same balls which are added on successful plays. Given some signaling cost, the cost exacted on any play is the signaling cost times the number of signals sent. That is, if the

[35] Perhaps they presume that if more than one context is relevant, there is more at stake, and thus a higher payoff. This assumption will not be reproduced in the model we will present in §4. As we shall see, the authors state that the ratio that mattered the most was between signal cost and the single context payoff, so this higher payoff when both contexts are relevant will not tip the scales too much in favor of compositionality.

signaling cost is 0.5, and two signals are sent on a play, then the cost on that play will be 1. So, each time a ball is chosen from an urn, 1 of the chosen ball is removed from the respective urn, up to a limit of having no less than one of each ball in each urn at all times (as the system is at its start).[36] These two parameters, payoff and cost, need to be specified.

To understand the set up better, consider an example of a successful iteration (Figure 7). Following the authors, we will set the signaling cost to 0.5, the single context payoff to 1.5, and the "both" context payoff to 2. First, the context and state of nature are determined; say, color and black cat, so the run will be successful and payoffs will be attained if the basic receiver paints the canvas black. Next, the executive sender chooses who sends (ideally, the system will evolve so that whenever the context is color, the executive sender chooses a particular basic sender, when it is animal, it chooses the other, and when it is "both," it chooses both). Let's say the first sender is chosen, and they choose signal A. Now, the executive receiver sees who sent, so it needs to choose which context is relevant from its first sender urn while the basic receiver chooses its action. The basic receiver will need to choose between the urns corresponding to AC and AD (ideally in this case, these will be black cat and black dog, so that way A fixes black). Say it chooses the AC urn, and say that from the AC urn it chooses black dog. If the executive receiver chooses that the basic receiver should interpret their action as a color action, then the basic receiver will paint their canvas black, which is correct (we can imagine any other combination of choices from urns which might lead to an incorrect action). Due to cost, we must remove 0.5 (having sent only one signal) of: A balls from the first sender's black cat urn, first sender balls from the executive sender's color urn, black dog balls from the basic receiver's AC urn, and color balls from the executive receiver's first sender urn. Since the outcome was correct, we add 1.5 (as there was a single context) of these same balls to these same urns. Thus, the difference between the payoffs and the costs is +1.0. It is clear that in order for the system to actually reinforce correct behavior, the payoffs must be higher than the signaling costs.

[36] In particular, the signaling cost considered here can do no more than restart the game; ball-types cannot be negatively reinforced out of existence.

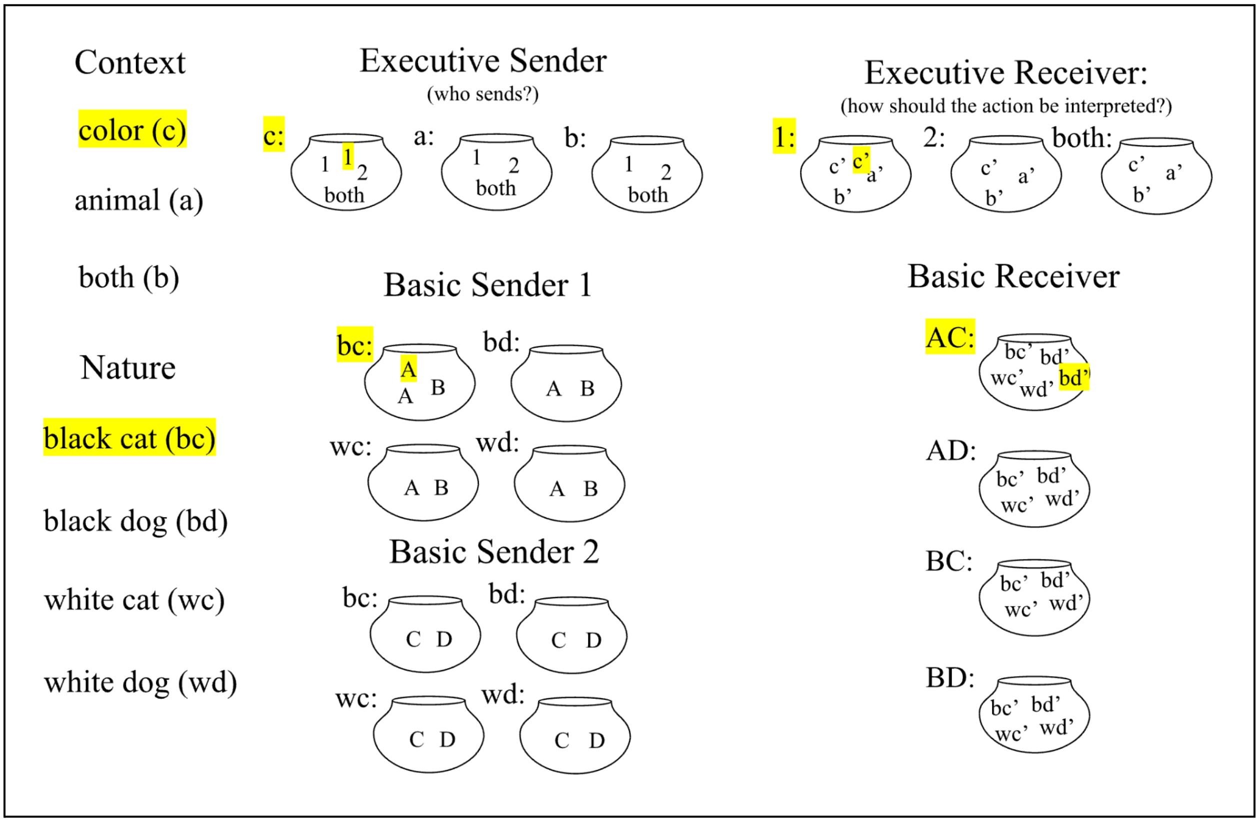

Figure 7: The hierarchical model after one round of reinforcement. Here, the context was color, the state of nature was black cat, and the correct action was therefore the one corresponding to black. The executive sender chose that the first basic sender should send, and the first basic sender chose to send signal A. The basic receiver then had to choose equiprobably between its AC and AD urns, since a full-length signal was not sent. It chose the AC urn, and within it selected the black dog action. Finally, the executive receiver saw that the first basic sender sent, and from its respective urn chose to interpret the basic receiver's action as a color action, resulting in the action corresponding to black. Thus, the action was correct, and all of the choices made were reinforced by the single state payoff of 1.5 minus the signaling cost of 0.5 times the number of signals sent (1), or 1.5-(0.5x1)=1. And so one of each chosen ball was added to its respective urn.

The authors tested a signal cost of 0.5 balls, a "both" context payoff of 2 balls, and a single context payoff of 1.5 balls, and found that the system almost always evolved to signal successfully and nearly-perfectly efficiently (perfect efficiency is exactly one signal for each color and each animal).[37] They found also that the success rate depended most on the ratio of signal cost to single context payoffs, specifically when the cost was sufficiently low, the successful systems were rarely ever nearly-perfectly efficient. This makes sense, and supports

[37] To see what "almost always" means, see their paper.

the claim that this is an accurate model of compositionality—if signaling cost (time, cognitive processes, etc.) was sufficiently low, it seems that signalers should always communicate everything they can to always ensure the best action, but when signaling cost is high, the sender should only send the signal that is (evolved to be) relevant to the context. Numerically, a lower cost results in less evolutionary pressure toward composition because the sum of the payoffs and costs of inefficiently sending two signals in a single context state will still be positive, and so inefficiency will also be selected for.

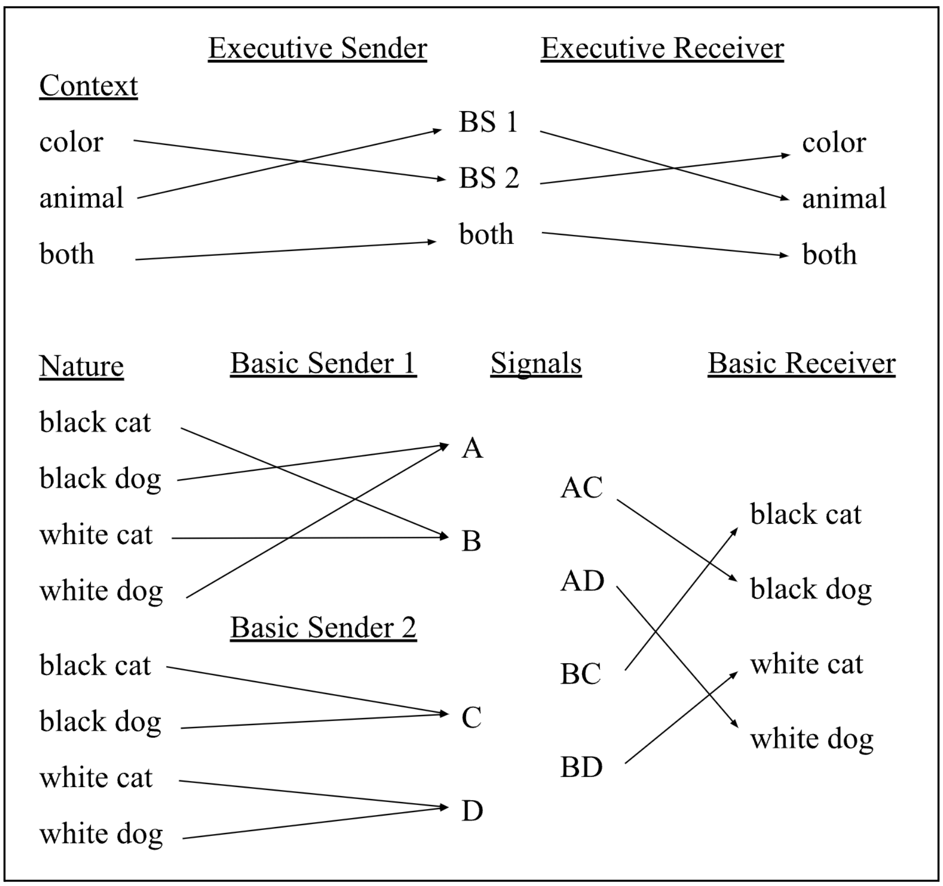

Figure 8: A flow chart depicting (one possible) perfectly-efficient composition. Here, the first basic sender communicates the animal, as per the executive sender, with its signal A picking out dog and signal B picking out cat. The second basic sender communicates the color, and its signal C picks out black and signal D picks out white. The executive and basic receivers interpret and act accordingly.

It remains to check that the meanings of the composed signals are actually functions of the meanings of the individual signals, as this is a benchmark requirement for compositionality.

Consider the run described above; when the only signal sent was A, the basic receiver had to choose between the AC and AD urns. Thus, A signals whatever AC and AD evolve to have in common. In Figure 8, A picks out dog, and indeed AC and AD correspond to black dog and white dog, respectively. However, it may be the case that AC and AD do not have any state of nature in common (imagine one evolves to mean black cat and the other evolves to mean white dog). This complicates the question of what A means,[38] but in cases like these, as Barrett, Cochran, and Skryms say: "composition on this type of hierarchical model is sometimes significantly more subtle than simple conjunction."[39] They give the example of "only," which is rarely used on its own, but when paired with "child" may communicate a precise idea. Thus, if the second basic sender distinguishes among the colors, for example, then the combined signals may differentiate the four states of nature even if they do not hold precise meanings on their own.

Obviously, this model, like each we've discussed, is quite simple—there is only one form of compositionality, and there is very little to communicate. That being said, it does have the desired property of the meaning of the composed signals being a function of the meanings of the composing signals, and it behaves as we expect, with high signaling cost driving the evolutionary pressure toward compositionality.

One possible objection to the set up of the model, however, is that it begins with exactly as many senders as contexts (and therefore as many combinations of senders as combinations of contexts) and exactly as many signals per sender as states of nature per context. Thus, it may be argued that though compositionality consistently evolves, there may be no other evolutionary loci for the model to converge to. In nature, it seems unlikely that these numbers will conveniently align. We will seek to strengthen their result by removing this assumption, replacing it with an initial state lacking any basic senders and signals, and instead assuming a mechanism with which these may be invented. We can then study the evolution of a model which can have arbitrary signals and combinations, and thereby measure the relationship between invention cost and compositional signaling, as is the goal of this paper.[40] However, in order to

---

[38] And this is also a less ideal signaling system since it can do no better than random for the context the model evolves to associate with the first basic sender.
[39] Barrett, Cochran, and Skryms, p. 918
[40] The model, of our own creation, will be discussed in §4. See Figure 12 for the initial state of that model.

introduce this complication, we first now probe what invention looks like in the context of a signaling game. This will also serve the purpose of showing these signaling games to be useful in the study of another major feature of human language, which is a vocabulary that grows and evolves with the language.

3.6 Invention

In considering the development of a language's vocabulary in terms of the invention of new signals, we will step back to the basic signaling game described in §3.3 and add this new layer of complexity on its own. Then, we will consider a model combining both composition and invention in the next section. An intuitive implementation of invention would be for the sender to be able to either send one of its signals as usual or introduce a new signal. An investigation into one way this can look can be found in "Inventing New Signals" by Alexander, Skyrms, and Zabell, published in 2011.

The paper describes two different formulations of the same probabilistic invention process: the Chinese restaurant process and Hoppe-Pólya urns. The former is as follows. Imagine a restaurant with infinitely-many tables, each with infinitely-many seats, and a phantom guest sitting at the first. Now, infinitely-many guests enter, one at a time, and they will only sit at a table with another person (including the phantom guest) seated at it. The probability that they will sit at any given table is equal to the number of guests at that table divided by the total number of seated guests. Thus, they are most likely to sit at the table with the most guests, second-most likely to sit at the table with the second-most guests, etc. Thus, the first guest sits at the first table. But, the phantom guest does not enjoy the company of others, so they move to the next table whenever someone sits next to it. The second guest then has an equal probability of sitting at either of the first two tables, and if they sit with the phantom guest, the phantom guest moves to the third table. This continues *ad infinitum*.

To map this to an urn process, let the restaurant be an urn. The table the phantom guest sits at is represented with a special ball, which we will call the "invention ball," and this is initially the only thing in the urn. Each draw from the urn represents another guest choosing a seat. So, when the invention ball is chosen (as will always happen on the first draw), this is like a new guest sitting with the phantom guest (which always happens to the first guest)—the phantom guest moves tables, and a new type of ball, representing the table the phantom guest just left, will be added to the game. There will now be one of this new ball-type in the urn, corresponding to the one guest now seated at that table; when it is selected, then another ball of that type is added to the urn, and this is like a new guest sitting at a table with another guest.

In this way, if each type represents a signal, a sender may invent a new signal by choosing a ball corresponding to the phantom guest, which we will call the invention ball. The sender will have an invention ball in each of its urns (i.e. its urn for each state of nature). Once invented, a new signal will be added to each of the sender's urns, as the sender is allowed to send any of its signals given any state of nature. To curb the rate at which signals are added to the game, the authors of the paper built the model so that if the invention ball is chosen, a new signal would only be added if that run of the signaling game was also successful and reinforcement took place. That is, when the invention ball is chosen, and a new signal is sent, the basic receiver acts as it does with any yet-to-be-reinforced signal, and chooses an action at random. When it chooses the wrong action, and the run fails, that new signal is discarded, but when a run introducing a new signal is successful, that new signal is added to the sender's urns (along with a second copy, since the success reinforces that signal; see Figures 9 and 10 for an illustration).

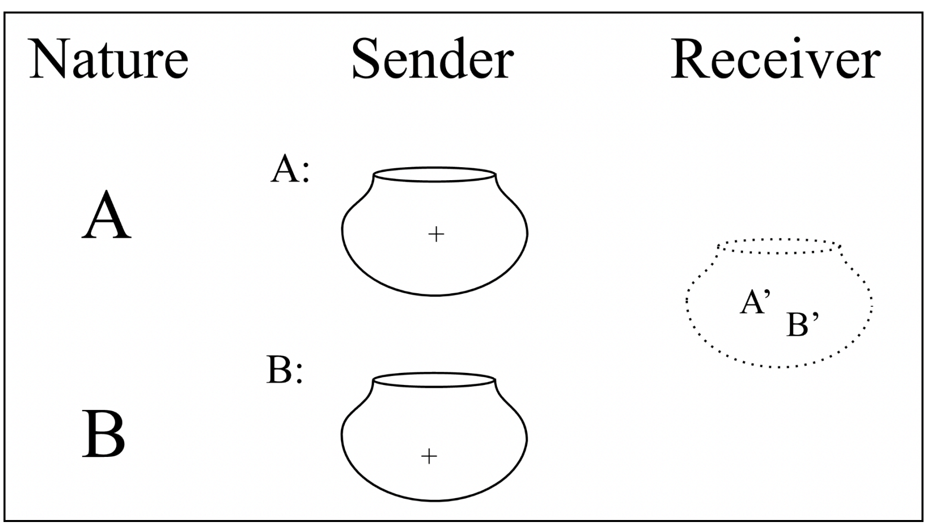

Figure 9: The initial state of the signaling game with invention. There are yet to be any signals present in the game, and so the receiver has no urns.

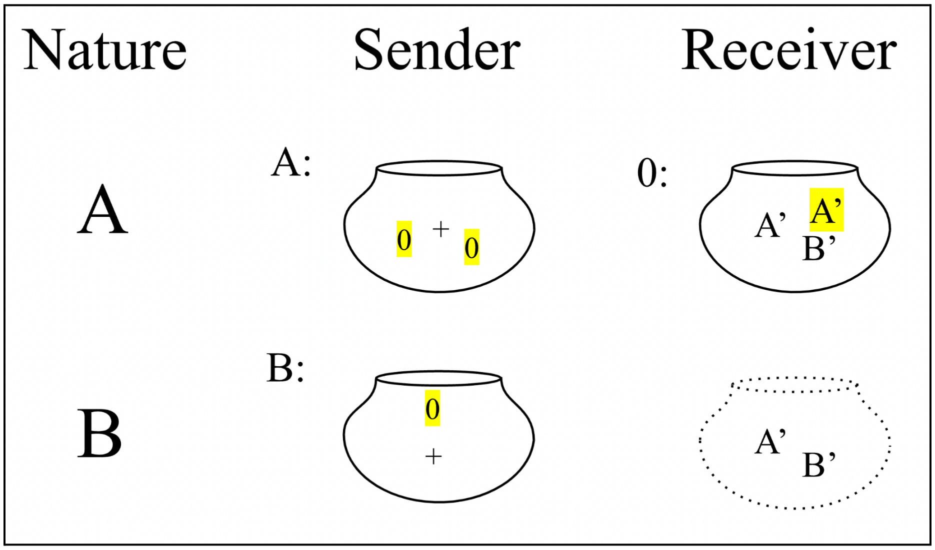

Figure 10: The invention signaling game after a successful round wherein the state of nature was A, the invention ball was chosen, and the receiver acted correctly by choosing A'. Thus, the 0 ball was added to **both** of the sender's urns, as it is being introduced as a signal, and a second copy was added to the sender's A urn since reinforcement takes place. The receiver's 0 urn was also reinforced with an A' ball.

A new consideration that now arises is the presence of synonyms. Indeed, in infinite time, we will certainly have infinitely-many signals, and so for any state of nature, we can have arbitrarily-many signals which converge to correlate exactly to it. But the authors found that for a three state, three act signaling game with one sender and one receiver, only between five and twenty-five signals would form in one thousand trials of one hundred thousand iterations (runs through the game). Furthermore, it was generally the case that for any state of nature, only one or two strong synonyms (i.e. signals reliably chosen given the state of nature) would form. Thus, the presence of synonyms is not overly deleterious to the model, and the authors also introduced two notions of forgetting in their tenth section, to deal with runaway synonym production.[41]

Synonyms can be useful, however. An issue that generally arises in signaling games with three (or more) states, signals, and actions without invention is that of *partial pooling equilibria*, where information is not communicated efficiently by the signals (i.e. a state of nature might correspond to more than one signal or a signal to more than one action; see Figure 11 for an

[41] See their paper for a more detailed discussion on the invention model itself and the forms of forgetting they considered and implemented.

example),[42] but this did not occur when invention was present. Alexander, Skyrms, and Zabell propose that this was due to invention resulting in synonyms, and they support this claim by showing that the initial presence of extra signals helps deal with partial-pooling in systems without invention.[43] Given that we know synonyms (or more generally, redundancies) exist, the model discussed here gives us a potential explanation; if more efficient signaling results in higher payoffs, agents with reasonable numbers of synonyms should be better selected for.[44]

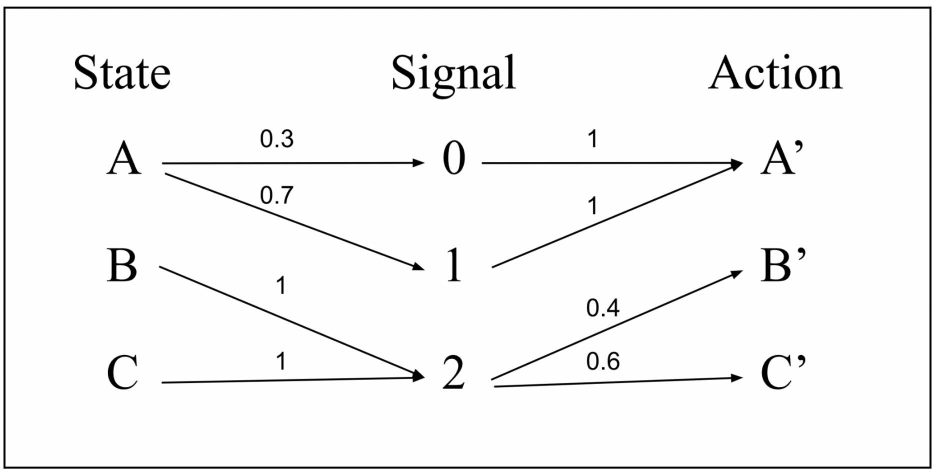

Figure 11: An example of a partial pooling equilibrium. A three signal, state, and action game can get stuck with this inefficiency, but this problem never arises with invention. Here, the state A corresponds to two signals, whereas states B and C share one signal. Then, there is information lost, and the signaling system is not efficient, as no signals differentiate states B and C.

Now, given independent models of both compositionality and invention, we may now ask: how would a signaling game with both invention and compositionality look and evolve, what changes in parameters (such as invention cost, signaling cost, single context payoff, and "both" context payoff) will drive its evolution, and would the model's behavior on simulation match our intuitions about how language works?

---

[42] For a more in-depth discussion of partial pooling equilibria due to unequal probabilities among the states of nature or due to the presence of multiple states of nature and actions (i.e. greater than two), see the fifth chapter of *Signals* by Brian Skyrms.

[43] The details are not particularly important here; what matters is that the introduction of invention solved a general inefficiency for even slightly more complicated signaling systems. This problem could have been solved in another way, such as an initial presence of more signals than states of nature, but invention is a more natural assumption, given, for example, the existence of genetic mutation which can result in never-seen-before genes and respective expressions. For their discussion on this issue and its resolution, see their section 7.

[44] What makes a number of synonyms reasonable likely depends entirely on the context—a balance must be struck between having enough to benefit but not too many to reasonable store and communicate with.

4 Our Model

4.1 The Initial State of the Model

The goal of this thesis is to probe the relationship between invention and composition in a model which can be said to have both. Structurally, our model will be hierarchical in the same sense as the model of Barrett, Cochran, and Skyrms, discussed in §3.5, with a key difference: just as the invention model of Alexander, Skyrms, and Zabell, discussed in §3.6, begins without signals and is able to invent them via an invention ball in its sender's urns, our model will begin without signals and indeed without basic senders. In its initial state, the game will only have an executive sender, an executive receiver, and a basic receiver. For simplicity, we will begin with the same states of nature, contexts, and actions as in the presented hierarchical model: states of black cat, black dog, white cat, and white dog; contexts of color, animal, and both; and, actions of painting on a canvas.[45] We will build the model such that we can change the numbers of states and contexts it in order to consider the evolution of a signaling system which needs to communicate about a more or less complex world. We can also change the probabilities of the states of nature and contexts, but within our scope, they will all be equiprobable. Though we will not test these changes, we note that differences in these parameters may result in changes in the model's evolution.

The basic receiver's actions correspond to the states of nature—there is a one-to-one correspondence between states of nature and actions such that the action is correct if and only if the corresponding state of nature is present. The action chosen by the basic receiver needs to be contextualized by the executive receiver as in the following examples: if black cat or black dog is chosen by the basic receiver and the executive receiver chooses color, then the canvas will be painted black, if white dog or black dog is chosen along with the animal context, then a dog will be painted, and if white cat and the "both" context are chosen, then a white cat will be painted. The idea is that if an animal is painted without a color, the only information about the state of nature that will be acted upon is that of the animal. It is not the case that a random color is chosen. This distinction is necessary because if the both contexts are relevant but only the animal

[45] These aren't relevant to the model or its evolution, but make things easier to talk about. The only important detail is that there are two contexts each corresponding to two mutually-exclusive states of nature.

is communicated, this is considered a failure and there will be no payoff; that is, there is no probability of accidentally having the correct color and acting correctly. A random color being chosen does not seem like an unreasonable choice, but in the model we will build, a "both" context with no color interpretation per the executive receiver is equivalent to the wrong color being chosen. Thus, we can think of an animal action as being the respective animal being drawn on the canvas in a neutral color, say blue.[46]

In the model, each random choice (with the exception of the state of nature and context whose probabilities will not change during the life of the game) is modeled with Pólya urns, where each option is represented as balls within an urn, initially containing one of each kind. Whenever the culminating action is correct, reinforcement acts by adding more of each chosen ball to its respective urn (the exact number added will depend on the payoffs which will be specified later on).

The game begins with a state of nature and context being chosen at random. The executive sender views the context, and the basic senders, if any have been invented, view the state of nature. The executive sender has an urn for each context, and each of these urns begins with only an invention ball, which invents basic senders. Here, the invention of basic senders is tantamount to the addition of new compositional, or grammatical, structures in language. Then, words, represented by signals, must themselves be invented in order to be conjugated within these contexts. Thus, the basic senders also start with only an invention ball, which invents signals. Like Alexander, Skyrms, and Zabell, we will build the model so that invention only takes place when reinforcement takes place.[47] In other words, when invention takes place in a round of the game which fails because an incorrect action is performed, the invention ball will simply be put back into its urn. But if the correct action is chosen after an invention ball is selected, then invention takes place; that is, if the executive (resp. basic) sender invented a basic sender (resp. signal), then the newly invented basic sender (resp. signal) will be added to all of

---

[46] This issue is unfortunately magnified by this particular example since it seems strange to consider a drawn animal that has no color, but since this has already been used in literature, we will continue using it. Interestingly, this same issue is not present with color since we can easily imagine a colored canvas with no animal on it.

[47] Each time invention takes place, the chance of success is the same, since all agents who act given the newly invented (basic sender/signal) will act randomly. For the two context, two states per context set up, the chance of successfully inventing a basic sender having chosen the executive sender's invention ball is 1/3 (there are two single context and one "both" context), and the chance of successfully inventing a signal having chosen a basic sender's invention ball is 1/4 (there are four states of nature).

the executive (resp. basic) sender's urns, along with the payoffs from this successful run. For the sake of simplicity, the executive sender will only be able to invent a new basic sender if there are no already invented basic senders who have yet to invent any signals. Besides inventing basic senders, the executive sender is able to choose any combination of basic senders that have been invented to send on that play. Then, on each play, the executive sender's choice is, in turn, the relevant urn of the executive receiver; that is, the executive receiver chooses its interpretation of the action based on which basic senders sent. If there are 3 basic senders labeled 1, 2, and 3, the executive sender's ball-types will be 1, 2, 3, 12, 13, 23, 123, or invention. Note that the numbers merely denote the basic senders chosen on that run, and so the order does not matter. For example, if 23 is chosen, then basic senders 2 and 3 will send signals during this play of the game.

The role of the basic senders is exactly as in the classic signaling game, where they choose a signal from urns corresponding to the state of natures, with the practical differences that they initially have no signals, but can invent, and they may or may not send a signal on some plays, per the executive sender's decision. The basic receiver has an urn for every possible full-length combination of signals, and so whenever some basic sender is not chosen to send on a play, the basic receiver chooses equiprobably between the signals for that sender.[48] This generalizes the notion introduced in the hierarchical composition model of Barrett, Cochran, and Skyrms, where, for example, if there are two signals for two signalers, A or B and C or D, then the possible combinations would be AC, AD, BC, BD, and if only D was sent, for example, the basic receiver would choose among the AD and BD urns with equal probability. To illustrate how this would work in our model, imagine then that we have 4 basic senders, each of which sends 0, 1, or 2, and so the basic receiver's urns are the 4-digit strings composed of 0, 1, and 2. Then, suppose that the first, second, and fourth senders send on some iteration, sending 1, 1, and 0, respectively. That means that the basic receiver would choose among its 1100, 1110, and 1120 urns, with equal probability. Within these urns, the basic receiver chooses one of the actions. An interpretation of this could be the basic receiver "guessing" what the correct state of nature is, having fixed whatever information was conveyed by the signals it did receive.

[48] An alternative formulation of this is that the basic receiver chooses equiprobably among the urns which agree with the signals which were sent.

Finally, the executive receiver is exactly as in the hierarchical compositional model without invention; it has urns for each combination of senders, and chooses, from the relevant urn, the context the basic sender's action should be interpreted against. If invention is chosen by the executive sender, then the executive receiver will always choose among the contexts equiprobably (as reinforcement will not yet have taken place at all for this new choice of a basic sender). The executive receiver's choice can be thought of as interpreting what the basic receiver actually needs to do given the information it received. In particular, if the basic receiver guesses that the state of nature is black cat, and the executive receiver determines the animal context to be relevant, then the basic receiver furthermore guesses that the only thing its action needs to take into account is the animal, and so a cat is painted. Thus, if the state of nature were actually white cat or black cat and the actual context was animal, this run would have resulted in a successful action and paid off.

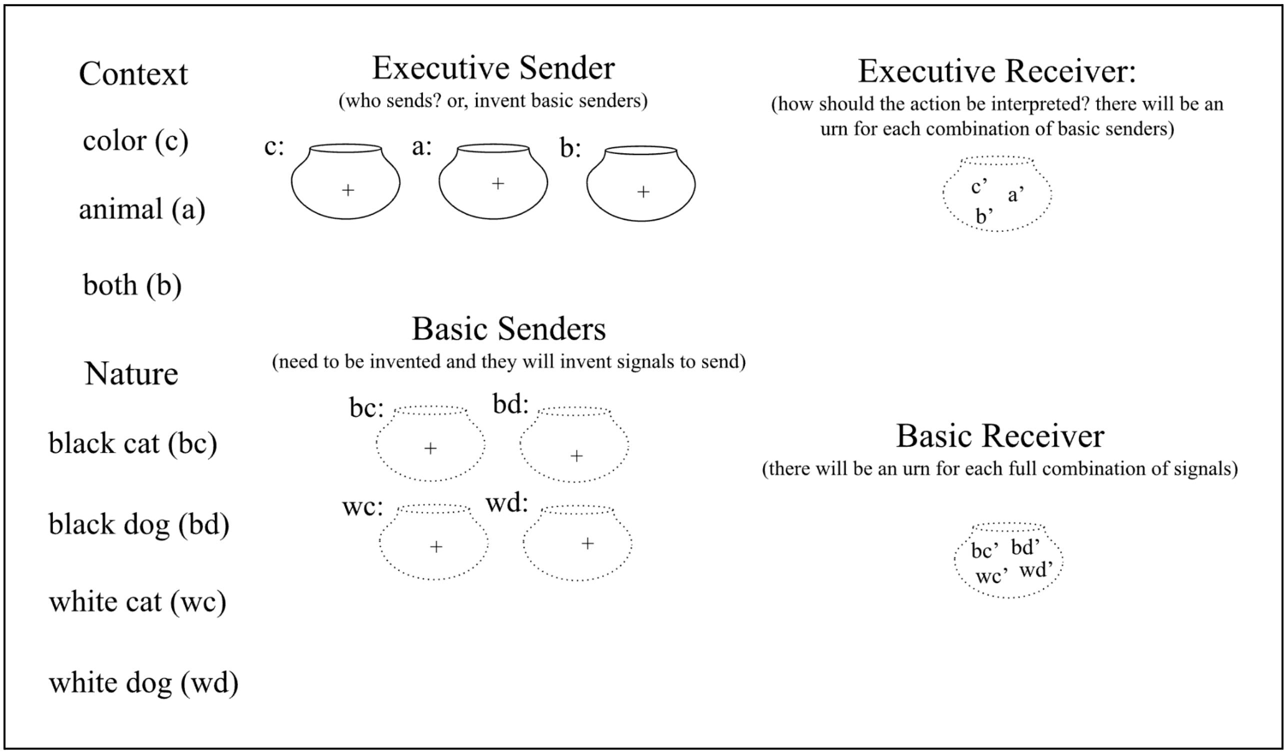

Figure 12: The initial state of our model. Structurally, it is like the hierarchical model in §3.5 (see Figure 6), with the addition of invention balls, labeled +, and it starts without basic senders and signals. On each play, a context and state of nature are chosen equiprobably at random. The executive sender views the context and chooses a combination of already-invented basic senders or invents a new basic sender. The basic senders chosen by the executive sender view the state of nature and send or invent signals. The basic receiver sees the signals sent and chooses an action; it randomly chooses a signal from each basic sender who did not send. Finally, the executive receiver sees which senders sent on this run, and chooses how the basic receiver’s action is to be interpreted.

4.2 The Evolution of the Model

Having laid out the model’s structure and initial state (§4.1, following the hierarchical model of §3.5), we must now describe how it iterates; we have already described the mechanisms of reinforcement and invention (following §3.3 and §3.6, respectively). Briefly, reinforcement occurs within Pólya urns, where choices, represented by the selection of balls from the urns, undergo the following reinforcement mechanism: when a correct action is chosen, any balls chosen on that run are put back along with an extra copy. Invention takes place via special “invention balls,” which, when chosen on successful runs, introduce new ball-types to the relevant urns. It remains for us to specify the reinforcement payoffs and costs. Given the goals of the thesis, it is important that the costs of invention and composition are comparable. That is, we should be able to change the parameters in such a way that it makes sense to say something like "invention is cheaper," so we may then probe the effects of such a difference.

The payoffs must take place on the level of the Pólya urns, as this is where reinforcement takes place. Thus, we will retain the standard reinforcement mechanism, where additional copies of the chosen balls are added to their respective urns upon correct actions, with a few slight changes. Firstly, note that the basic receiver has urns for each full-length signal, and so whenever a new basic sender invents a new signal, the basic receiver has an entirely new set of urns to choose from, and all previously relevant urns (for shorter length signals) will never again be used.[49] What this means is that it will always be the case that, no matter how reinforced a basic receiver is for some particular signals, it will need to restart. Thus, the basic receiver has less time to reinforce correct behavior than any of the other agents. To counteract this, we reinforce the basic sender's urns multiplicatively, where, instead of adding a single of the chosen ball in the case of a correct action, the total number of that ball is instead multiplied by 1.1. In this way, assuming the other agents have already begun converging onto specific choices, the basic receiver will be able to catch back up quickly, so to speak. For the rest of the agents, we simply add an extra of the chosen ball after a correct action is produced. Unlike the hierarchical composition model, we will not reinforce more after a successful "both" context than after individual contexts. In Barrett, Cochran, and Skyrms's model, this was necessary because this made the extra signal cost induced by composing more worthwhile during the "both" context. We will not be reusing their notion of cost, and so adding more balls on these plays would not have the same effect.[50]

To understand why we must create and implement a different notion of cost, notice that Barrett, Cochran, and Skyrms modeled cost as signal cost, where there was some specific cost to signaling, and the payoffs on any run would be reduced by the number of signals sent times the

[49] Rather than using new urns, we could theoretically have new urns somehow depend on the rest of the previous urns, but there is no obvious way of doing this that wouldn't essentially prevent new basic senders' signals from having meanings not yet evolved toward.

[50] In particular, we could have made the magnitude of reinforcement correspond to the number of relevant contexts, but this would not have changed the evolution of the model. This is because, since the payoffs do not serve to cancel out costs (as will be explained shortly), inefficient, but sometimes correct, behavior will still be positively reinforced, as in the case of the hierarchical model with lower signaling cost evolving to compose efficiently less frequently, as discussed in §3.5. The goal of the reinforcement here is only to drive the agents toward correct behavior, and then the costs will determine which sorts of signaling systems will be most successful.

specified cost. For example, with a cost of 0.5, sending two signals on an unsuccessful run results in a payoff of -1. One caveat they provided this mechanism of cost was that it could not change the possible outcomes of any play, only their probabilities. Thus, if cost brought the number of some type of ball below one, the number of that ball is instead reset to one. They found that when the cost was very low, the system would rarely compose. This agrees with our intuitions about how compositional language should work—if there is no harm in sending more, we might as well always send all potentially relevant information. There is no pressure for composition to be selected for.

It seems clear that signals are costly, and in more than one way. For example, an agent sending many signals should incur more cost, since they must form and send a more complex combination of signals, which takes (at the very least) time and some kind of computational power, and the receiver must interpret that more complex combination, also taking time and computational power. But, only considering signaling cost, the game described in the previous section might always evolve toward a signaling system where each context is exactly represented by one basic sender and each state of nature by one signal. It seems that the system should always avoid composing, regardless of context, since it can invent a unique signal for each state and context, and if a single signal can always communicate all the relevant information, the agents will maximize received payoffs by never composing. Thus, we must generalize signaling cost to a notion which can also accommodate a notion of cost for inventing signals. The intuition we seek to match in our model is that we know that we humans sometimes invent and sometimes compose, and so our model should be able to evolve toward primarily doing either[51] or not reach efficient signaling at all (this last possibility will also help to defend the model against the objection that the meaning which evolves within it is baked into it).

4.3 Our Notion of Cost

We will develop a notion of cost which applies to the entire signaling game. This will provide us with a new interpretation of the game being played: we will run the simulation some large number of times, and we will view each result in the context of all the results obtained under the same parameters. Thinking of this total as a population, those simulations who achieve

[51] To be formalized shortly.

the most efficient signaling can be thought of as having accrued the greatest amount of fitness. Thus if, for example, perfectly efficient composition (see Figure 8 in §3.5), as tends to evolve in the hierarchical model, can relatively consistently evolve under a certain set of parameters, this can be interpreted as it possibly arising in the population. If we can argue that it is an ESS, then its arising will further represent its eventual displacing of other, less efficient, signaling strategies.

This will be dealt with soon. For now, we consider one possible notion of cost, and develop it. We will constrain our evolving agents with limited resources (including time, cognitive capacity, etc.), and require that they act before those resources diminish. As such, we will provide each simulation with some number of *i*, which will serve to generalize this notion of resource, and each time composition or invention takes place, some portion of the *i* will be removed. By fixing a high enough *i* such that efficient signaling can evolve and then changing the respective costs, we can compare how those costs affect the signaling strategies that evolve within *i*. These costs will be multiplicative.[52] For example, we could choose that invention incurs a cost of 0.9 (90% of the *i* will remain after invention takes place) and composition incurs a cost of 0.99 (99%). So, if five signals were to be sent on one play of a simulation with these costs, and *i* was at 1000 at the start, it would be at 1000×0.99^5 which is approximately 951 after signaling. If, instead, invention had taken place on that iteration, *i* would be at 900.

In our set up, the "amount" of convergence towards a signaling system that has taken place by the time the *i* diminish will be taken to be irrelevant; we will only consider the most populous ball-type of each urn, without caring about exactly how populous they are. That is, imagine some basic sender, call it A, has invented four signals, say 1, 2, 3, and 4, by the time the *i* diminish. If this A's black cat urn has four 1 balls, ten 2 balls, one-thousand 3 balls, and thirty-nine 4 balls, then we say that A sends 3 given black cat. The same is true if it only had forty 3 balls in the black cat urn, but if it had thirty-eight 3 balls, then it sends 4 given black cat. In case of a tie, the first to be invented will be chosen. Given how Pólya urns iterate, the most reinforced ball-type is statistically the most likely to continue being reinforced. Also, we could interpret this as telling the simulation to choose the signaling system it deems best having spend the *i* it did evolving, and this will be the most reinforced.

[52] To deal with the fact that multiplicative costs might never cause *i* to go to 0, we also add in a small chance of *i* randomly decrementing by one on any iteration of a simulation.

A second notion of cost which we will also implement is as follows: there is no particular reason to have exactly one invention ball for both the executive sender and all of the basic senders. By decreasing this number, we increase the (metaphorical) bar which must be met in order for invention to take place. For example, it makes sense for it to be more likely for a signal to be invented than for a basic sender to be invented; it seems that the natural conditions that must be met for a signal to evolve given a signaler ought to occur more frequently than the natural conditions that must be met for a signaler to evolve. Thus, we may have fewer invention balls in the executive sender's urn than invention balls in the basic sender's urns, and both of these values may be less than one. For example, by having 0.05 invention balls in the executive sender's urn, once it has invented a single basic sender, the probability of it doing so again is far lower than it would have been should we have put, say, 1 invention ball instead. Furthermore, if we have 0.05 invention balls in the executive sender's urn but 0.1 invention balls in the basic senders' urns, then we can think of the invention of signals as more accessible for the basic senders than the invention of basic senders is for the executive sender. By decreasing the number of balls, we make invention more difficult, and, in a sense, more costly. With this, we have finished describing the model and its evolution.[53]

4.4 Interpreting the Model's Results

Firstly, we must distinguish between successes and failures for the model's simulations, and then consider different kinds of successes, and their relationships to each other and to the model parameters. An obvious case of failure is when the signaling system fails to act correctly at the point when the *i* run out, i.e. if the wrong action is performed given some state of nature and context, such as painting the canvas white when the state of nature is black cat and the relevant context is animal. However, acting correctly should not be enough for the signaling system to count as having succeeded. What we want is efficient signaling, and I submit that there are two possible, efficient ESSs which may evolve in this system. In order to compare two different, but both correctly-signaling outcomes, it will clearly not suffice to only consider whether they signal correctly. Thus, we will consider the structure of the systems themselves,

[53] The Code Appendix contains Python implementations of our model and every other signaling game described in this paper, along with the code we used to gather our results.

and thereby determine which systems are most successful, and therefore evolutionarily stable. In particular, even if all of the communicated information is correct, a signaling system may be less efficient by needing to maintain more basic senders; we can think of basic senders as somehow resource intensive, and that population members with evolved signaling systems that use fewer basic senders thus have greater fitness, as they can allocate more of those resources (such as time, energy, etc.) toward reproduction.

That being said, another clear case of failure is having only one basic sender in play. This is because the executive receiver decides the context against which the basic receiver interprets their action given the basic senders who sent, and so with only one basic sender present, the system can do no better than random (and it can decidedly do worse by reinforcing one of the contexts). The executive receiver needs to choose the correct context in order for the run to be successful, and with only one basic sender, there is no information about the context communicated between the agents of the game.

So we need more than one basic sender and, immediately, it may seem that the only ESS would then be what I will call "perfect invention," (Figure 13) where for each possible context there is a single, unique basic sender whose signals can differentiate between the relevant states (i.e. there is one basic sender who can differentiate between white and black, a second who can differentiate between cat and dog, and a third who can differentiate between all four states). This being the only ESS, however, implicitly assumes that it is more efficient to use the least number of basic senders on any iteration, or one for each context. But perfect invention requires the use of three basic senders in total. If it were at least as efficient to instead have the lowest total number of basic senders, what I will call "perfect composition" (see Figure 8 in §3.5) would be an ESS. Perfect composition is the case where there are two basic senders, each differentiating between the states of one context, which differentiate between all four states when composed. That is, perfect composition occurs when one basic sender uses different signals for different colors and is chosen by the executive sender when the context is color, another basic sender has different signals for different animals and is chosen by the executive sender when the context is animal, and when the context is "both," both of these basic senders (and no more) send, and their composed signal picks out the correct action. Here, there are only two basic senders in play, at the tradeoff of using two basic senders during the "both" context. Intuitively, it seems that in an environment where composition is more costly, perfect invention is to be favored, whereas when invention is more costly, perfect composition should be ideal.

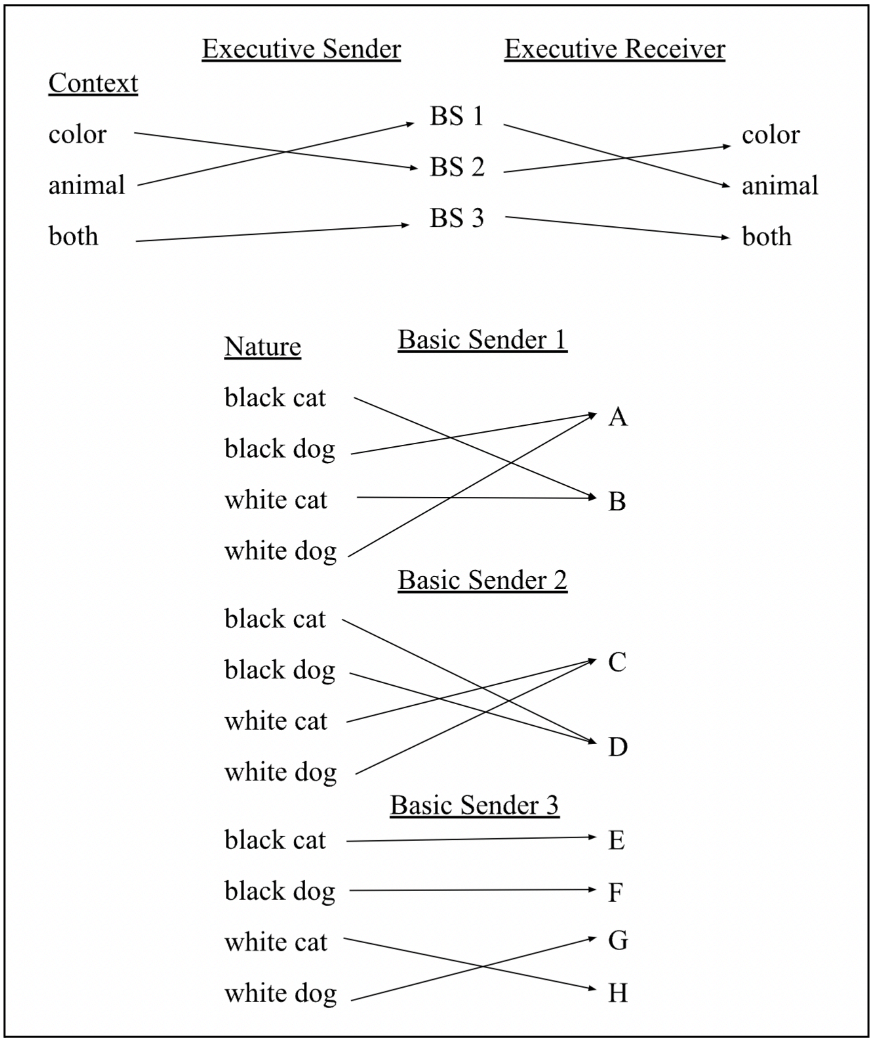

Figure 13: An example of signals sent in perfect invention. Here, basic sender 1 is chosen when the context is animal, and it sends signal A given dog and signal B given cat. When the context is color, basic sender 2 is chosen, and it sends signal C given white and signal D given black. When both of the contexts are relevant, then basic sender 3 sends, and it has a signal for each state of nature—E for black cat, F for black dog, G for white dog, and H for white cat.

There is, however, an inherent issue with perfect invention in this model. Consider the example in Figure 13. Since the basic receiver must choose among the full-length signals, the basic receiver cannot evolve to always choose black cat when E is sent *and* always choose one of black dog or white dog when A is sent, for example. If it could, then suppose the state of nature is black cat and both contexts are relevant. Then the executive sender would select basic sender 3 to send, and it would send E. Then, it would randomly select among the signals of basic senders 1 and 2, so it will choose among ACE, ADE, BCE, and BDE equiprobably, and all of these are contradictory except for BDE. Thus, we expect perfect invention to signal less successfully and thus gain less fitness, than other, potentially less efficient, signaling systems, and so it should rarely evolve in this model. However, we cannot solve this issue by simply letting the basic receiver consider partial signals instead of full length-signals, since this mechanism is what allows the signals' meanings to compose to the meaning of the composed signals, as discussed in §3.5.

One final note is that, occasionally, the model would invent so much that, on simulation, it would run for longer than feasibly testable as the numerous inventions were too computationally expensive. This can be considered a failure for two reasons, but it will form its own category distinct from failures. Firstly, we can think of this result as the simulation having invented so much that it had no signaling system ready by the time *i* ran out. Secondly, whenever simulations like this would eventually halt, the evolved signaling system was usually extremely inefficient, where many basic senders would be composed even for each single context. This occurred rarely (as will be seen in §4.5), and only when invention of basic senders was sufficiently easy (in terms of the number of invention balls present), and was solved by merely cutting off individual simulations after running for a certain amount of time. Thus we add the rule that, on simulation, the model must evolve a signaling system by the time *i* runs out or within five minutes, whichever comes first. However, since the data from these simulations are technically unknown, and since this depends on the computer on which these simulations are run, these will be considered their own, distinct form of failure, which we can call "time out." And so the simulation has four possible outcomes: perfect invention, perfect composition, failure (including any form of inefficiency or incorrect signaling), and time out.

4.5 Results

Setting $i=10^9$, we ran the two context, two states per context 1,000 times under a variety of configurations, changing the number of invention of basic sender balls the executive sender had in their urns (BAS), the number of invention of signal balls the basic senders had in their urns (SIG), the multiplicative invention cost (INV), and the multiplicative composition cost (COM). Since these are multiplicative and less than one, a lower number is a higher cost. In the following table, we report some tested parameter values, along with the number of perfect inventions (I), perfect compositions (C), failures (F), and time outs (E) that occurred during the 1,000 simulations with those parameters.

Parameters				Results			
BAS	SIG	INV	COM	I	C	F	E
0.05	0.1	0.90	0.99	0	20	980	0
0.05	0.1	0.99	0.99	0	18	982	0
0.05	0.1	0.90	0.90	0	4	996	0
0.05	1.0	0.90	0.99	0	39	961	0
0.1	0.1	0.90	0.99	0	15	985	0
0.2	1.0	0.99	0.99	0	17	970	13
1.0	1.0	0.90	0.99	0	1	986	13
1.0	1.0	0.90	0.90	0	2	997	1

As expected, perfect invention did not evolve frequently (or indeed at all) on simulation. Perfect composition's frequency depended on all of the parameters, but especially the number of invention balls, i.e. how accessible invention was within the signaling system. In particular, perfect composition occurred most frequently when signals, but not basic senders, were easy to invent. Making basic senders easier to invent always decreased the frequency of perfect composition, whereas making signals easier to invent tended to increase the frequency of perfect composition. Lowering composition cost increased the frequency of perfect composition, as did increasing invention cost, though the multiplicative cost of invention had less of an effect than the other parameters.

5 Discussion

5.1 Evolutionary Pressure

Firstly, it is notable, though expected, that perfect invention did not evolve on simulation. As discussed in the previous section, there is an inherent issue with the model's reinforcement mechanism that discourages perfect invention—the model cannot evolve to have basic senders whose signals exactly correspond to conflicting states of nature, as, even when only one is sent, the basic sender may choose an urn corresponding to both signals. Less damningly, perfect invention also has a particularly difficult requirement to meet. Having a separate basic sender to differentiate (with its signals) between the states of each single context is not that difficult; indeed it is also required for perfect composition to evolve, and this occurred. The further requirement, however, of having a third agent whose signals differentiate between all four states of nature requires that a basic sender be invented who themselves invents at least four signals which are then reinforced in the correct way. Perhaps this is not reasonably attainable in the model, especially given the fragile nature of the basic receiver's reinforcement and the impossibility of correctly reinforcing the basic sender for perfect invention. Even given these issues, we do expect that perfect invention should, with some probability, occur in this model, due to the fact that it is possible within the framework and since chancy, sometimes-incorrect behavior will still be positively-reinforced without negative reinforcement.[54] It is therefore possible that some of the failures were somehow "close" to perfect invention and were cut off.

Furthermore, we remind ourselves that many of the failures will be signaling correctly, but inefficiently, and the successes we measure exist in the context of a population of signaling strategies with positive fitness. We are interpreting the proportion of successes as the proportion of a population of one thousand randomly-acting but subject to reinforcement signalers which evolve to adhere to one of the strategies we have argued to be an ESS. Within this framework, we can consider what might occur in the next generation. Assuming that our agents' offspring are biased to at least begin signaling as their progenitors do, the greater the likelihood for an agent to randomly evolve to use an ESS suggests a greater speed at which the population should

[54] I have also seen it occur during testing.

tend toward it. Thus, the differences in values among the different parameters' resultant perfect compositions are significant; they represent how "easy" it is for perfect composition to arise under those circumstances, so to speak. Though we did not measure how many of the failures are correctly signaling, it is clear that perfect composition, being an ESS, will be selected for the most among its population members. And so, with a higher proportion of the population (say 39/1000 as in the fourth set of parameters as opposed to 2/1000 in the last), the population will also move towards being a majority perfect composers more rapidly. Thus, it is clear that when this value is substantially higher, there is a greater evolutionary pressure for the population to evolve perfect composition, and our results tell us that this pressure was correlated with it being difficult to invent basic senders, it being easier for signals to be invented than for basic senders to be invented, composition not being too expensive, and, slightly less so, invention not being too cheap.

There is also the issue of perfect composition seemingly rarely evolving. Even in the most ideal case, it is not immediately clear that 39/1000 is a significant proportion. However, there are two considerations here. Firstly, we are interpreting the results on the scales of populations. Thus, this 39/1000 represents not only the possibility of perfect composition evolving given a relatively simple model, assuming only reinforcement and invention, but it further represents it evolving consistently. This means that the results actually suggest that this simple form of composition evolves fairly easily. Secondly, there is a disanalogy to the hierarchical composition model of Barrett, Cochran, and Skyrms, where perfectly-efficient composition evolved 97.5% of the time.[55] As discussed earlier, they had a mechanism of negative reinforcement which itself drove the evolution towards compositionality. When they sufficiently decreased signal cost, inefficient signaling was far more prevalent, as it still resulted in positive payoffs. In our model, there is no negative reinforcement driving the results, and we are merely asking whether there is a relationship between the relative costs of invention and composition and the resultant signaling systems. Since there is no negative reinforcement, inefficient signaling, which is chancy and often incorrect, still has on-average positive payoffs, and will thus be reinforced. It is possible that implementing a notion of negative reinforcement would drive the model toward greater proportions of perfect signaling, whether only composition or both

[55] Barrett, Cochran, and Skyrms, p. 917.

composition and invention, much more frequently. Furthermore, many of the failures that occurred in our model were not possible in their model, where the signaling system which evolved could not, for example, be inefficient by utilizing too many basic senders.

5.2 The Scope of this Work

As introduced in §1, we created this model with two primary goals in mind. One was to show that compositionality evolves with more natural assumptions than those made in the hierarchical model of Barrett, Cochran, and Skyrms. This was accomplished by removing the exact numerical alignment between the contexts and basic senders and between the signals and states per context, instead introducing a mechanism of invention, finding that perfect composition evolved consistently within this framework, and showing that this evolution is dynamic and depends on the parameters. The other goal was to provide formal justification for the claim that the evolutionary pressure toward compositionality depends on how accessible and costly invention is, and this is made clear by the results of the simulation. Thus, this project can be thought of as providing a formal justification for claims such as "composition evolves easily" and "evolutionary tendency toward linguistic composition correlates with how expensive invention is." The latter point is especially important since, as argued in §1, it seems obvious that if we had no restrictions on the words we use, we would gain nothing from language.

Importantly, however, the goal of the model is *merely* to show that 1) the evolution of compositional meaning assumes very little and 2) there is a relationship between the costs of inventing and composing in signaling (as seems intuitive in language). This is not to say that human language can be fully explained by models such as the one presented here. Human language is complex; take the examples of portmanteaus and anacronyms.[56] These cases are simultaneously invention and composition. The portmanteau "brunch," which both composes "breakfast" and "lunch" and is itself a new word, and the anacronym "laser," which is now a word in its own right but used to be an acronym for "light amplification by stimulated emission of radiation," do not fit cleanly into the binary I have presented. Even if the evolution of human language began in the same way as simpler signaling systems, it has since grown in complexity, and there is an explanatory gap of sorts. To even posit that natural language *can* fully be

[56] I thank Abigail Ehrhardt for suggesting these cases during an in-class presentation.

explained in terms of these sorts of interactions would be wrong, and even if it could, we might not necessarily ever be able to reduce it fully, and even if we could, that full reduction wouldn't necessarily teach us about the relationship between these simple games and natural language, any more than a map of all the neurons in the brain, if even possible, wouldn't necessarily teach us any more about consciousness.[57]

This is an issue with this genre of philosophical work in general. For example, perhaps the games presented in §3.2 help us intuit why we evolved norms such as justice, but our cultural evolution has grown complex enough that it would seem strange to try to motivate an example as arbitrary as saying "bless you" upon hearing a sneeze (something present across many cultures) on the level of evolutionary payoffs and natural selection using the tools of §3.

Having restricted the scope of the work, it is now much easier to approach the question of whether the model presented is a plausible account. Given that the model's set up assumes no more than the union of the assumptions of invention signaling game and the compositional signaling game, and given that the meaning that arises within it is defined in the same way, its status as a plausible account holds if and only if those two models are themselves plausible accounts. And we have already justified the assumptions of these models and discussed the presence of meaning, and in particular compositional meaning.

5.3 Future Work

The results of this paper could be strengthened. This could be done by, for example, measuring the proportion of failures that still signaled correctly, and somehow checking whether they were close to perfect invention or perfect composition. Also, as mentioned earlier, adding a negative reinforcement mechanism might make the model more realistic, as inefficiencies tend to be costly for signalers, whether they be cells or humans, and this might increase the evolved proportion of perfect signalers in the population. What exactly this mechanism would look like, however, is unclear. Since perfect invention never evolved, and since there seemed to be an inherent contradiction between the model's set up and its evolution, a further exploration of the conditions where it would evolve might shed light on why it failed to here. Finally, the

[57] This latter example I borrow from an interview with Matthew Harrison-Trainor (https://today.uic.edu/mathematician-matthew-harrison-trainor-receives-sloan-research-fellowship/).

underlying questions about compositionality and its relationship with invention could be explored using other sets of tools. An evolutionary model provides us with one facet of it, but there are also, for example, linguistic, cultural, and cognitive considerations not taken into account here.

References

Barrett, Jeffrey A., Calvin Cochran, and Brian Skyrms. “On the Evolution of Compositional Language.” Philosophy of Science 87, no. 5 (2020): 910–20. doi:10.1086/710367.

Harsanyi, John C. “On the Rationality Postulates Underlying the Theory of Cooperative Games.” The Journal of Conflict Resolution 5, no. 2 (1961): 179–96. http://www.jstor.org/stable/172785.

Huck, Steffen, and Jörg Oechssler. “The Indirect Evolutionary Approach to Explaining Fair Allocations.” Games and Economic Behavior 28, no. 1 (July 1999): 13–24. https://doi.org/10.1006/game.1998.0691.

McKenzie Alexander, J., Skyrms, B. & Zabell, S.L. Inventing New Signals. Dyn Games Appl 2, 129–145 (2012). https://doi.org/10.1007/s13235-011-0027-2

Oosterbeek, H., Sloof, R. & van de Kuilen, G. “Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis”. Experimental Economics 7, 171–188 (2004). https://doi.org/10.1023/B:EXEC.0000026978.14316.74

Raffaele Argiento, Robin Pemantle, Brian Skyrms, Stanislav Volkov, Learning to signal: Analysis of a micro-level reinforcement model, Stochastic Processes and their Applications, Volume 119, Issue 2, 2009, Pages 373-390, ISSN 0304-4149, https://doi.org/10.1016/j.spa.2008.02.014.

Skyrms, Brian (2010). *Signals: Evolution, Learning, and Information*. Oxford, GB: Oxford University Press.

Skyrms, Brian. “Sex and Justice”. Chapter. In *Evolution of the Social Contract*, 2nd ed., 1–22. Cambridge: Cambridge University Press, 2014. doi:10.1017/CBO9781139924825.003.

Skyrms, Brian & Barrett, Jeffrey A. (2019). “Propositional content in signals”. Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences 74 (C):34-39.

Smith, J., Price, G. “The Logic of Animal Conflict”. Nature 246, 15–18 (1973). https://doi.org/10.1038/246015a0

Szabó, Zoltán Gendler, “Compositionality”, The Stanford Encyclopedia of Philosophy (Fall 2022

Edition), Edward N. Zalta & Uri Nodelman (eds.), https://plato.stanford.edu/archives/fall2022/entries/compositionality/

Xunzi: The Complete Text. Trans. by Eric L. Hutton. Princeton, NJ: Princeton University Press, (2016).

Code Appendix

What follows is Python (NumPy) code, with my own comments which may help with readability, though the comments were written for myself. Each subsection is titled after the game whose code it is. Though the literature often lacks this level of transparency, I believe it is imperative for the development and sharing of ideas within the academic community; the code helps answer questions about the model I might have overlooked in my discussion, and it saves future contributors from needing to reinvent the wheel, so to speak. The code is directly copied from Visual Studio Code. The final subsection is the code I used to script (i.e. automatically run and gather results from) our model.

Basic Signaling Game

```
import numpy as np
import matplotlib.pyplot as plt
import time

rng = np.random.default_rng()

#storing globals here, with initial values:
states_of_nature = [0,1] #if 0th state of nature, sender chooses 0th urn, etc.

sending_urns = [[0,1],[0,1]]

receiving_urns = [[0,1],[0,1]]

#run the signaling game for this many iterations
for t in range(10**4):
   #keep track of every choice so we can reinforce afterward (not state of nature)
   #choose a state of nature
   nature = rng.choice(states_of_nature)

   #choose a signal, conditional on the state of nature
   signal = rng.choice(sending_urns[nature])

   #choose an action, conditional on the signal sent
   action = rng.choice(receiving_urns[signal])
```

```
    #check if the action was correct
    if action == nature:
        #reinforce if it was
        sending_urns[nature].append(signal)
        receiving_urns[signal].append(action)

#after running, check how well reinforced everything is
counter = 0
for i in sending_urns:
    val = np.sum(i) / len(i)
    print("sending urn " + str(counter) + ": " + str(val))
    counter += 1

counter = 0
for i in receiving_urns:
    val = np.sum(i) / len(i)
    print("receiving urn " + str(counter) + ": " + str(val))
    counter += 1

#theoretically now, could run the program itself many times and find the proportion of
runs of the game
#with above a certain amount of reinforcement (i.e. the percentages are sufficiently
high)

#for pausing: time.sleep(secs)
```

Hierarchical Composition Game

```
import numpy as np
from itertools import chain, combinations, product
import copy

'''
NOTE: on reinforcement urns
floating point values this time
ex: urn=[4.5, 1.0]
    signal=rng.choice([signals], p=urn/urn.sum())

each index represents a ball type and each value represents the number of balls of
that type
```

```
'''

rng = np.random.default_rng()
def norm(inp):
   return [float(i)/sum(inp) for i in inp]

#globals
num_contexts = 2
states_per = 2
iterations = 10**5
signal_cost = -0.5
def payoff(context):
   return 1.5 + 0.5*(len(context)-1)
   #the exact value here is flexible and is meant to match the 2 state, 2 context
   #game in 'on the evolution of compositional language'


'''
   here, each bit of the state of nature represents one context.
   so, two contexts of two states each will be ['00', '01', '10', '11']
   for simplicity, I will begin with each context having the same number of states
'''
temp_string = ""
for i in range(states_per):
   temp_string += str(i)
states_of_nature = [''.join(x) for x in product(temp_string, repeat=num_contexts)]

'''
   notably, the number of relevant contexts is more than the number of contexts,
   since any combination of contexts is itself a relevant context.
   numerically this will be #P(S)-1 (the size of the power set minus the empty set)
   and we want to identify the bit number of the context.

   e.g. with 2 contexts we want: {{0}, {1}, {0,1}}
        with 3                 : {{0}, {1}, {2}, {0,1}, {0,2}, {1,2}, {0,1,2}}

   then, we will loop through the relevant contexts to see if everything relevant is
correct.
   the payoff will depend on the number of correct
```

```
   note that I got the following code from a stackexchange post and it returns a list
of tuples,
   which are similar to lists but are not mutable
'''
s = list(range(num_contexts))
contexts = list(chain.from_iterable(combinations(s, r) for r in range(1, len(s)+1)))

'''
   we need as many basic senders as contexts and as many signals per sender as states
per context
   and we need an urn with each signal for each state of nature

   example:
   sender[0] should be the first sender
   sender[0]['00'] should be the first sender's urn for the '00' state of nature
'''
basic_senders = []
basic_sender = {}
initial_basic_sender_urn = []
for i in range(states_per):
   initial_basic_sender_urn.append(1.0)
   #we add one ball to each "slot" of the urn
for i in states_of_nature:
   basic_sender[i] = initial_basic_sender_urn.copy()
for i in range(num_contexts):
   basic_senders.append(copy.deepcopy(basic_sender))

'''
   there are only one of each: executive sender, basic receiver, and executive
receiver

   - executive sender needs to choose from among the combinations of basic senders
(i.e. P(S)) given the context
   - basic receiver needs to choose from among the actions (exactly the states of
nature) given the signals sent
       - the basic receiver has an urn for each full combination of signals and will
choose equiprobably among the
         ones which match the sent signals
   - executive receiver chooses among the interpretations (exactly the contexts) given
the combination of basic senders
     chosen by the executive sender
```

```
   - the end state is "correct" if the interpretation matches the actual context and
the action matches the actual state at
     the contextually relevant indices
'''

#the combinations of senders are exactly the combinations of contexts (...)
executive_sender = {}
executive_sender_urn = []
for i in contexts:
   executive_sender_urn.append(1.0)
for i in contexts:
   executive_sender[i] = executive_sender_urn.copy()

executive_receiver = {}
executive_receiver_urn = []
for i in contexts:
   executive_receiver_urn.append(1.0)
for i in contexts:
   executive_receiver[i] = executive_receiver_urn.copy()

basic_receiver = {}
basic_receiver_urn = []
for i in states_of_nature:
   basic_receiver_urn.append(1.0)
for i in states_of_nature:
   basic_receiver[i] = basic_receiver_urn.copy()

#play
for t in range(iterations):
   #choose the state of nature and context equiprobably
   state = rng.choice(states_of_nature)
   context = contexts[rng.choice(range(len(contexts)))]

   #choose who sends based on the context
   sending = rng.choice(list(enumerate(executive_sender[context])),
p=norm(executive_sender[context]))
   '''
   here, sending[1] is the value we will want to reinforce and apply signaling cost to
and
   sending[0] is the index in contexts representing the senders we are sending
   so, contexts[int(sending[0])] is the actual list of senders we chose
```

```
    '''

    #choose the sent signals, keeping track of which to reinforce
    signal = ""
    signal_indices = []
    for i in range(num_contexts): #this will go in numerical order so I should be able
to treat it index by index of the signal
        if i in contexts[int(sending[0])]:
            chosen = rng.choice(list(enumerate(basic_senders[i][state])),
p=norm(basic_senders[i][state]))
            signal += str(int(chosen[0]))
            signal_indices.append([i, int(chosen[0])])
        else:
            signal += str(rng.choice(states_per))
            #if index wasn't chosen, choose signal randomly

    #choose the action, which will later need to be contextualized by the index of
states_of_nature it represents
    action = rng.choice(list(enumerate(basic_receiver[signal])),
p=norm(basic_receiver[signal]))

    #choose the interpretation based on the senders who sent
    interpretation = rng.choice(
        list(enumerate(executive_receiver[contexts[int(sending[0])]])),
p=norm(executive_receiver[contexts[int(sending[0])]])
    )

    print(interpretation[0])

    #reinforcement and signal costs
    add = len(contexts[int(sending[0])]) * signal_cost

    correct_type = contexts[int(interpretation[0])] == context
    correct_act = True

    for i in context:
        if states_of_nature[int(action[0])][i] != state[i]:
            correct_act = False

    if correct_type and correct_act:
        add += payoff(context)
```

```
   executive_sender[context][int(sending[0])] =
max(executive_sender[context][int(sending[0])] + add, 1.)
   basic_receiver[signal][int(action[0])] = max(basic_receiver[signal][int(action[0])]
+ add, 1.)
   executive_receiver[contexts[int(sending[0])]][int(interpretation[0])] = max(
       executive_receiver[contexts[int(sending[0])]][int(interpretation[0])] + add, 1.
   )
   for i in signal_indices:
       basic_senders[i[0]][state][i[1]] = max(basic_senders[i[0]][state][i[1]] + add,
1.)

#print the outcome
print("\nexecutive sender:" + str(executive_sender))
print("\nbasic senders:" + str(basic_senders))
print("\nexecutive receiver:" + str(executive_receiver))
print("\nbasic receiver:" + str(basic_receiver))
```

Invention Signaling Game

```
import numpy as np
import matplotlib.pyplot as plt
import time
from collections import Counter

rng = np.random.default_rng()

#storing globals here, with initial values:
number_of_states = 5

states_of_nature = []
sending_urns = []

for i in range(number_of_states):
   states_of_nature.append(i) #if 0th state of nature, sender chooses 0th urn, etc.
   sending_urns.append([-1,])
   #-1 represents the invention ball
   #any invented ball is added to both urns with an extra reinforcement ball for the
chosen urn

#keep track of how many balls have been added
```

```
to_add = 0

#since the states of nature are fixed, the urns that are added
new_urn = states_of_nature.copy()

#this begins empty as there are no signals
receiving_urns = []

#run the signaling game for this many iterations
for t in range(10**5):
   #keep track of every choice so we can reinforce afterward (not state of nature)
   #choose a state of nature
   nature = rng.choice(states_of_nature)

   #choose a signal, conditional on the state of nature
   signal = rng.choice(sending_urns[nature])

   if signal == -1:
       #choose an action equiprobably
       action = rng.choice(states_of_nature)

       #check if the action was correct
       if action == nature:
           #add new signal, urn
           for x in sending_urns:
               x.append(to_add)
           receiving_urns.append(new_urn.copy())

           #reinforce
           sending_urns[nature].append(to_add)

           #iterate
           to_add += 1
   else:
       #(same as normal signaling game)
       #choose an action, conditional on the signal sent
       action = rng.choice(receiving_urns[signal])

       #check if the action was correct
       if action == nature:
           #reinforce if it was
```

```
            sending_urns[nature].append(signal)
            receiving_urns[signal].append(action)

#check how many signals have been invented
signals = set(sending_urns[0])
print(signals)

#count occurences of each signal in each sending urn
ctr = 0
for urn in sending_urns:
   print("\nurn " + str(ctr) + ":")
   for i in signals:
       prop = urn.count(i) / len(urn)
       print("signal " + str(i) + ": " + str(prop))
   ctr += 1
```

Our Game

```
import numpy as np
from itertools import chain, combinations, product
import copy
import random

rng = np.random.default_rng()

import sys
print ('argument list', sys.argv)
if len(sys.argv) != 8:
   print("not enough args")
   exit()

#args are: num_contexts states_per t invention_cost composition_cost
invention_of_basic_sender_balls invention_of_signal_balls

#HELPER FUNCTIONS#
def norm(inp):
   return [float(i)/sum(inp) for i in inp]

#this decoder takes an executive sender's urn's index value to the combination of
basic senders it represents
```

```
#these numbers will also be the prior for the executive receiver (as opposed to the
combinations themselves).
def decoder(num, senders):
   i = 1
   while 2**i-2 < num:
       i += 1
   senders.append(i-1)
   num -= 2**(i-1)
   if (num >= 0):
       decoder(num, senders)
   return(senders)

#and this is because I forgot about the invention ball in the executive sender's urns
def fake_decoder(num, senders):
   return decoder(num-1, senders)

#this is by Musa :)
def get_perms(length, locked_index, locked_val):
   res = []
   get_perms_helper(length, locked_index, locked_val, res, [])
   return res

def get_perms_helper(length, locked_index, locked_val, res, curr_perm):
   if length == 0:
       res.append(curr_perm.copy())
       return

   # check if this index is locked
   # if so, use the desired value and proceed
   if locked_index == 0:
       curr_perm.append(locked_val)
       # generate the rest of the recursive perms as usual
       get_perms_helper(length - 1, locked_index - 1, locked_val, res, curr_perm)
       curr_perm.pop()

   # and if we are using the locked value, we don't want to use any other values
   else:
       # inclusive of the max value at this index
       # len(curr_perm) gives us the index of the value we are about to place in the
current perm
       for j in range(1, len(basic_senders[len(curr_perm)][states_of_nature[0]])):
```

```
            curr_perm.append(j)
            get_perms_helper(length - 1, locked_index - 1, locked_val, res, curr_perm)
            curr_perm.pop()


#GLOBALS#
num_contexts = int(sys.argv[1])
#we will remain under the assumption that each context has the same number of
corresponding state values.
states_per = int(sys.argv[2])
t = 10**int(sys.argv[3])

#the invention cost will apply to both invention of basic senders and signals
invention_cost = float(sys.argv[4]) / 100
#this will be the cost per signal sent
composition_cost = float(sys.argv[5]) / 100

invention_of_basic_sender_balls = float(sys.argv[6])
invention_of_signal_balls = float(sys.argv[7])

random_removal = 1 / 10000

#SET UP#
s = list(range(num_contexts))
contexts = list(chain.from_iterable(combinations(s, r) for r in range(1, len(s)+1)))
#these tuples the possible combinations of contexts
temp_string = ""
for i in range(states_per):
   temp_string += str(i)
states_of_nature = [''.join(x) for x in product(temp_string, repeat=num_contexts)]
#these strings correspond to the states of nature

'''
here, we will set up the following:
- executive sender with an urn for each context and context combination, each
containing an invention ball
   - whenever basic senders are invented, balls for each possible combination are
added (double the current number minus one)
- basic sender with an urn for each state of nature, each containing an invention ball
   - note that to begin with there will be none, but these will be deep copied when
invented by the executive sender
```

```
   - there will be array of all extant basic senders
- executive receiver who will begin with no urns, but whose urns will be added when
basic senders are invented.
   - these urns will choose among the interpretations which are exactly the contexts
- basic receiver who will begin with no urns, but whose urns will be added when basic
senders invent signals
   - these urns will choose among the actions which are exactly the states of nature
   - there will be one urn for each full-length signal
   - these urns will be in a dictionary where the ith index of a string corresponds to
the signal sent by the ith basic sender
   - whenever a basic sender doesn't send, the basic receiver will choose randomly
from among its extant signals
   - whenever a signal is invented, fix that signal, run through every full length
signal containing it, add urns to the dictionary
'''

#when adding new balls will have to loop through context and append to each
executive_sender = {}
executive_sender_urn = [invention_of_basic_sender_balls]
for i in contexts:
   executive_sender[i] = executive_sender_urn.copy()

basic_senders = [] #initially empty. make sure to do a deepcopy when populating
basic_sender = {}
initial_basic_sender_urn = [invention_of_signal_balls]
for i in states_of_nature:
   basic_sender[i] = initial_basic_sender_urn.copy()

executive_receiver = [0] #the 0 here is a dummy value so we can use the executive
sender's chosen index which will never be 0
executive_receiver_urn = []
for i in contexts:
   executive_receiver_urn.append(1.0)

basic_receiver = {}
basic_receiver_urn = []
for i in states_of_nature:
   basic_receiver_urn.append(1.0)

#PLAY#
while t > 0:
```

```
    #random chance that we just remove a t, to deal with getting stuck near 0
    if random.random() < random_removal:
        t -= 1

    #notably, in this model, different iterations will cost different amounts of t
    #as the solution for making invention and composition costs comparable.
    #thus, we must iterate t manually.

    #choose the state of nature and context equiprobably
    state = rng.choice(states_of_nature)
    context_num = rng.choice(range(len(contexts)))
    context = contexts[context_num]

    #choose who sends based on the context
    sending = rng.choice(list(enumerate(executive_sender[context])),
p=norm(executive_sender[context]))
    #reminder: choice list enumerate returns a tuple, [0] is index and [1] is value.

    if int(sending[0]) == 0:
        #inventing a basic sender takes place. no one sends on this play

        #inventing happens probabilistically
        #in 'inventing new signals' its 1/4 since the basic receiver initially has all
actions equiprobable given a signal
        #here, it will be the chance of getting the interpretation correct (if it was
both the interpretation and action correct then the
        #probability would be quite low and since the basic sender has no signals to
start considering the action wouldn't make sense)
        if random.random() < 1 / len(contexts):
            #invention can only take place if there are no basic senders who have no
signals
            should_invent = True
            for i in basic_senders:
                if len(i[state]) == 1:
                    should_invent = False
            if should_invent:
                #invention takes place
                for j in range(2**(len(basic_senders))):
                    for k in contexts:
                        executive_sender[k].append(1.0)
                    executive_receiver.append(executive_receiver_urn.copy())
```

```
                #reinforcement
                executive_sender[context][2**(len(basic_senders))] += 1.0
                executive_receiver[2**(len(basic_senders))][context_num] += 1.0
                basic_senders.append(copy.deepcopy(basic_sender))
                t *= invention_cost
            else:
                #no invention cost when invention doesn't take place, but costs an
iteration
                t -= 1
        else:
            t -= 1
        continue
    else:
        signal = ""
        signal_indices = []
        inventing_signals = []

        for i in range(len(basic_senders)):
            if i in fake_decoder(int(sending[0]), []):
                chosen = rng.choice(list(enumerate(basic_senders[i][state])),
p=norm(basic_senders[i][state]))
                if int(chosen[0]) == 0:
                    #invention takes place
                    inventing_signals.append(i)
                    #if length is n then there are n-1 signals and the invention ball
and the next signal to add is n.
                    signal += str(len(basic_senders[i][state]))
                    #note that we havent added the new signal to the basic sender, yet
                    t *= invention_cost
                else:
                    signal += str(int(chosen[0]))
                    signal_indices.append([i, int(chosen[0])])
                    t *= composition_cost
            else:
                #if a sender wasn't chosen and it has a signal choose a signal from its
options
                if len(basic_senders[i][state]) > 1:
                    signal += str(rng.choice(range(1, len(basic_senders[i][state]))))

        #if invention takes place we don't want to send the signal to the basic
receiver who doesn't (yet) have an urn for this signal
```

```
        if not len(inventing_signals) and signal in basic_receiver:
            action = rng.choice(list(enumerate(basic_receiver[signal])),
p=norm(basic_receiver[signal]))
        else:
            action = rng.choice(states_of_nature)

        interpretation = rng.choice(
            list(enumerate(executive_receiver[int(sending[0])])),
p=norm(executive_receiver[int(sending[0])])
        )

        correct_type = contexts[int(interpretation[0])] == context
        correct_act = True

        for i in context:
            if states_of_nature[int(action[0])][i] != state[i]:
                correct_act = False

        if correct_type and correct_act:
            for i in inventing_signals:
                #add the new signal to the basic sender
                for j in states_of_nature:
                    basic_senders[i][j].append(1.0)

                #check how long full length signals are by seeing if the most recently
added basic sender has signals
                #it should be the only basic sender that can not have signals by the
stipulation when adding new basic senders
                new_signal_length = len(basic_senders)
                if len(basic_senders[len(basic_senders) - 1]) == 1:
                    new_signal_length -= 1

                for j in get_perms(new_signal_length, i, len(basic_senders[i][state]) -
1):
                    #minus 1 because zero indexed, we've already added the new signal
above
                    new_signal = "".join(str(e) for e in j)
                    basic_receiver[new_signal] = basic_receiver_urn.copy()

                signal_indices.append([i, len(basic_senders[i][state]) - 1])
```

```
            executive_sender[context][int(sending[0])] += 1.0
            for i in signal_indices:
                basic_senders[i[0]][state][i[1]] += 1.0
            executive_receiver[int(sending[0])][int(interpretation[0])] += 1.0
            if signal in basic_receiver:
                #this is to prevent NaN with multiplicative reinforecement
                basic_receiver[signal][int(action[0])] =
min(basic_receiver[signal][int(action[0])] * 1.1, 10*9)

#Here, we will print either c (for correct composition), i (for correct invention), or
f (for failure)
#I'm only going to write this code for the 2 context, 2 states per context setup, but
it's very easily generalizable
#However, since the paper will only explicitly discuss this case, we will ignore the
fact that the model is capable of more (for now)
if num_contexts == 2 and states_per == 2:
   senders = []
   for i in contexts:
       senders.append(fake_decoder(np.argmax(executive_sender[i]), []))
   if (len(senders[0]) > 1 or len(senders[1]) > 1) or len(senders[2]) > 2:
       print("f")
       exit()
   else:
       if (senders[0][0] != senders[1][0]) and ((senders[0][0] in senders[2]) and
(senders[1][0] in senders[2])):
           #check if senders[0][0] differentiates between the first context states
           first_zero = []
           first_zero.append(np.argmax(basic_senders[senders[0][0]]["00"]))
           first_zero.append(np.argmax(basic_senders[senders[0][0]]["01"]))
           first_one = []
           first_one.append(np.argmax(basic_senders[senders[0][0]]["10"]))
           first_one.append(np.argmax(basic_senders[senders[0][0]]["11"]))
           diff_first = True
           for i in first_zero:
               if i in first_one:
                   diff_first = False
           #check if senders[1][0] differentiates between the second context states
           second_zero = []
           second_zero.append(np.argmax(basic_senders[senders[1][0]]["00"]))
           second_zero.append(np.argmax(basic_senders[senders[1][0]]["10"]))
           second_one = []
```

```
            second_one.append(np.argmax(basic_senders[senders[1][0]]["01"]))
            second_one.append(np.argmax(basic_senders[senders[1][0]]["11"]))
            diff_second = True
            for i in second_zero:
                if i in second_one:
                    diff_second = False
            if (diff_first and diff_second):
                print("c")
                exit()
            else:
                print("f")
                exit()
        elif len(senders[2]) == 1 and (((senders[0][0] != senders[1][0]) and
(senders[0][0] != senders[2][0])) and senders[1][0] != senders[2][0]):
            #check if senders[0][0] differentiates between the first context state
            first_zero = []
            first_zero.append(np.argmax(basic_senders[senders[0][0]]["00"]))
            first_zero.append(np.argmax(basic_senders[senders[0][0]]["01"]))
            first_one = []
            first_one.append(np.argmax(basic_senders[senders[0][0]]["10"]))
            first_one.append(np.argmax(basic_senders[senders[0][0]]["11"]))
            diff_first = True
            for i in first_zero:
                if i in first_one:
                    diff_first = False
            #check if senders[1][0] differentiates between the second context state
            second_zero = []
            second_zero.append(np.argmax(basic_senders[senders[1][0]]["00"]))
            second_zero.append(np.argmax(basic_senders[senders[1][0]]["10"]))
            second_one = []
            second_one.append(np.argmax(basic_senders[senders[1][0]]["01"]))
            second_one.append(np.argmax(basic_senders[senders[1][0]]["11"]))
            diff_second = True
            for i in second_zero:
                if i in second_one:
                    diff_second = False
            #check if senders[2][0] fully differentiates between the states
            third = []
            for i in states_of_nature:
                third.append(np.argmax(basic_senders[senders[2][0]][i]))
            diff_third = True
```

```
            for i in enumerate(third):
                for j in enumerate(third, i[0] + 1):
                    if i[1] == j[1]:
                        diff_third = False
            if ((diff_first) and (diff_second)) and diff_third:
                print("i")
                exit()
            else:
                print("f")
                exit()
        else:
            print("f")
            exit()

"""
#PRINT RESULTS#
#print("\nexecutive sender:" + str(executive_sender))
for i in contexts:
   print(str(i) + " " + str(fake_decoder(np.argmax(executive_sender[i]), [])))
for i in range(len(basic_senders)):
   print("\n")
   for j in states_of_nature:
       print("sender: " + str(i) + " state: " + str(j) + ", signal: " +
str(np.argmax(basic_senders[i][j])))
#for i in enumerate(basic_senders):
#    print("\nbasic sender " + str(i[0]) + ": " + str(i[1]))
#print("\nexecutive receiver:" + str(executive_receiver))
#print("\nbasic receiver:" + str(basic_receiver))
"""
```

Scripter

```
import time
import subprocess

file = open("log.txt", "w")

times = 1000
input_vals = ['python3', 'our_game.py', '2', '2', '9', '90', '99', '0.05', '0.1']
timeout = 300
```

```
#args are: num_contexts states_per t invention_cost composition_cost
invention_of_basic_sender_balls invention_of_signal_balls

f = 0
c = 0
i = 0
e = 0
for k in range(times):
   outp = 'e'
   try:
       proc = subprocess.run(
           args = input_vals,
           timeout = timeout,
           stdout = subprocess.PIPE)
       outp_str = proc.stdout.decode("utf-8")
       outp = outp_str[outp_str.find('\n') + 1]
       if outp == 'f':
           f += 1
       elif outp == 'c':
           c += 1
       elif outp == 'i':
           i += 1
       else:
           print("error")
       file.write("s " + str(k) + ": " + time.strftime("%H:%M:%S", time.localtime()) +
"\n")
       file.flush()
   except subprocess.TimeoutExpired:
       e += 1
       file.write("e " + str(k) + ": " + time.strftime("%H:%M:%S", time.localtime()) +
"\n")
       file.flush()

inp = " "
for k in range(7):
   inp += input_vals[k + 2] + " "

string1 = "args:" + inp + "iters: " + str(times) + " timeout: " + str(timeout)
string2 = ", f = " + str(float(f) / times) + ", c = " + str(float(c) / times) + ", i =
" + str(float(i) / times) + ", e = " + str(float(e) / times)
print(string1 + string2)
```