Signaling in an Unknown World

Rafael Ventura

Abstract

This paper proposes a sender-receiver model to explain two large-scale patterns observed in natural languages: Zipf's inverse power law relating the frequency of word use and word rank, and the negative correlation between the frequency of word use and rate of lexical change. Computer simulations show that the model recreates Zipf's inverse power law and the negative correlation between signal frequency and rate of change, provided that agents balance the rates with which they invent new signals and forget old ones. Results are robust across a wide range of parameter values and structural assumptions, such as different forgetting rules and forgetting rates. Analysis of the model further suggests that Zipf's law relating word frequency and rank arises because of language-external factors and that frequent signals change less because frequent signals are less subject to drift than rare ones. The paper concludes with some brief considerations on model-based and data-driven approaches in philosophy.

1 Introduction

Lewis (1969) introduced sender-receiver models to argue that agents with no common language could imbue signals with conventional meaning and thus give rise to communication. His models were simple and elegant. But they also made very stringent assumptions: interactions were thought to take place between agents that are perfectly rational and have common knowledge of the entire game. For all these assumptions, his main result proved remarkably robust. Extending Lewis' original framework to evolutionary and learning theory, Skyrms (2010) established that communication also arises in the absence of full rationality and common knowledge if natural selection or reinforcement learning are at work. Similar models went on to show that communication emerges across a wide range of conditions, including diverse learning rules (Huttegger et al., 2014), different levels of cooperation (Wagner, 2011; Zollman et al., 2013), and various forms of population structure (Wagner, 2009; Zollman, 2005).

This is progress. Whereas early models made strict demands on the cognitive capacity of agents, more recent incarnations à la Skyrms account for the origin of communication under very different assumptions. But there are reasons to worry about how much progress has been made. Recent gains in complexity have been driven primarily by a concern with the formal structure of these models, often to the detriment of empirical considerations. This is progress but of a type that is mostly "internal", as Levy (2011) puts it in a related context. In addition to results that are invariant across parameter values and structural assumptions, part of what it takes to build a good model is to relate assumptions and results to the real world.

Strides in this direction have already been made. Recent examples in philosophy include experiments by Bruner et al. (2014) and Rubin et al. (2019). Drawing on methods from behavioral economics, these studies test the sender-receiver framework by offering participants incentives to play the roles of sender and receiver in the lab. There is also a sizable and much older literature in experimental economics that goes at least as far back as Blume et al. (1998, 2001). Experiments testing sender-receiver models of conventional communication have also been conducted with non-human animals—for an early study and a recent review, see Silk et al. (2000) and Seyfarth and Cheney (2017). In all these fields, results have been largely consistent with model predictions. This is clearly valuable. By bridging the gap between reality and the lofty idealizations that permeate modeling exercises, experimental studies help models advance.

But the lab is also not the real world. Experiments on signaling may therefore not be representative of what goes on in nature. If the sender-receiver framework is to make empirical progress, then it is also important to assess how well its assumptions and predictions hold up against observational data. To this end, I present here a sender-receiver model that provides a possible explanation of two striking patterns found in natural languages: Zipf's inverse power law relating the frequency of signal use and signal rank (i.e., the position that a signal occupies in a frequency table), and the negative correlation between the frequency of signal use and rate of signal change (Lieberman et al., 2007; Moreno-Sanchez et al., 2016; Pagel et al., 2007; Zipf, 1949). Analysis of the model suggests that real-world communication systems obey Zipf's inverse power law because of language-external factors. The model further suggests that frequent signals change less because frequent signals are less subject to drift than rare ones. Although sender-receiver models may not be an adequate representation of natural languages (LaCroix, 2019), these results also indicate that some of the patterns observed in full-fledged human languages emerge in simpler communication systems as well.

The sender-receiver model in this paper present draws on Alexander's (2014) sender-receiver model of a dynamic state space. In his model, sender and receiver must learn to communicate at the same time that they discover new states. When the number of states increases without bound and at a constant rate, Alexander shows that communication cannot easily emerge and persist. In the present model, however, the number of states increases without bound but at a decreasing rate. This is because sender and receiver discover new states according to Hoppe's (1984) urn model—a model that has long been used to represent discovery processes (Huttegger, 2017; Skyrms, 2010; Zabell, 1992). The model also assumes that sender and receiver invent new signals and forget old ones. In this setting, I show that sender and receiver quickly learn to communicate as long as they balance the rates with which they invent new signals and forget old ones. I then show how the model makes contact with observational findings in linguistics: it recreates Zipf's inverse power law relating the frequency of signal use and signal rank, and the negative correlation between the frequency of signal use and rate of change.

The paper proceeds as follows. Section 2 begins by introducing a sender-receiver model in which states are discovered at a decreasing rate. Section 3 then presents results showing that agents learn to communicate with reinforcement learning provided that they balance the rates with which they invent new signals and forget old ones. Results also indicate that the model recreates Zipf's inverse power law, as well as the correlation between signal frequency and rate of change. In Section 4, I discuss these findings and offer possible explanations for the two patterns. Section 5 concludes with some methodological considerations about model-based and data-driven approaches in philosophy.

2 Model

In its basic form, a sender-receiver model consists of nature, a sender, and a receiver. Nature can be in one of several states given by the set $S = \{s_1, ..., s_n\}$. The sender first observes the state of nature and then chooses a signal from $M = \{m_1, ..., m_n\}$, the set of signals. The receiver does not observe the state but observes the signal that has been sent. Upon receiving a signal, the receiver picks an act from the set $A = \{a_1, ..., a_n\}$. In each state, there is a unique act that is successful in yielding a positive payoff to both agents; all other acts result in a payoff of zero. On their own, the sender cannot act and the receiver cannot observe nature. So agents must coordinate on a communication system if they are to maximize their payoffs.

The basic model assumes that nature occupies a constant number of known states. This is a plausible assumption to make in some cases. Sometimes we have to communicate in a world that we already know a lot about, so that states of nature are known in advance. However, this is not realistic in other scenarios. The need for communication also arises when we do not know in advance what or even how many states of the world there are. Sometimes we have to communicate in a new or yet unexplored environment. In such cases, we must first learn what states nature can occupy. Formally, this is analogous to the "sampling of species problem"—the question of how to infer the total number of species in an ecosystem from the number of species sampled so far (Fisher et al., 1943).

Zabell (1992) and Huttegger (2017) offer a precise mathematical representation and detailed treatment of this problem. Their basic idea is that we can represent the discovery of new species or categories by Hoppe's (1984) urn model. In this model, we suppose that there is initially an urn containing a single black ball. This is the "mutator" ball. We pick a ball at random from the urn. Whenever we pick the black ball, we return to the urn the black ball together with a ball of a novel color. This represents the discovery of a new species or category. Whenever we pick a ball of a color other than black, we return the chosen ball to the urn and add another ball of the same color to that urn. This represents the reinforcement of a category that is already known. Over time, the urn becomes populated with balls of many different colors. Some colors are more common in the urn, representing categories that are more likely to be observed than others. The black ball is never destroyed. So there is always a small probability that we pick the black ball and discover a new category.

In a signaling context, Skyrms (2010, pp. 118-135) proposes a similar model to represent the discovery of new states, signals, or actions. Although Skyrms uses this model to represent the invention of new signals, I assume here the same mechanism for the discovery of new states. That is, sender and receiver discover states of nature according to Hoppe's urn model. The model is thus similar to Alexander's (2014) model in that the number of states increases without bound. But unlike Alexander's model, new states arrive at a decreasing rate. In particular, I assume that no state is known in the beginning. Nature initially chooses a state from $S = \{s_0\}$, where s_0 is the "mutator" state. Whenever nature picks this state, sender and receiver discover a new state. The newly discovered state is added to the set of states that nature can choose from in the future. The sender can then observe this state at any subsequent point in time, and the receiver can draw an act corresponding to that state from the set of available actions. Whenever nature picks a state other than the mutator state, the sender observes the state, and the state is reinforced.

Why should we think that nature discovers new states and reinforces them when they are observed? The short answer is that we should not. To avoid confusion, it is therefore worth emphasizing that it is sender and receiver who discover new states and reinforce them whenever they are observed. The decision to model nature in this way is just a formal tool to represent more easily what sender and receiver know. In fact, we could suppose that each agent makes their own discoveries about nature. But this would require us to keep track of two separate discovery processes: one discovery process for the sender, and another for the receiver. It would also increase the complexity of the coordination problem as agents would have to coordinate not only signals and acts, but also what states of nature they know. Although it would be interesting to explore such a dual-discovery process in more sophisticated versions of the model, for simplicity I assume here that sender and receiver experience one and the same discovery process. To say that nature discovers and reinforces states is thus a shorthand for saying that sender and receiver discover new states in synchrony and reinforce existing ones at the same pace. Given that this is a game of cooperation, the assumption that agents discover and learn in tandem is not all that unrealistic.

It is also important to clarify why, in the current model, sender and receiver discover new states and reinforce them according to Hoppe's urn model. Clearly, Hoppe's urn is not the only possible way to represent a process of discovery. But for a number of reasons, it is a good starting point to represent discovery processes in the real world. For one, there is a long tradition of using Hoppe's urn to model discovery processes in a decision- and game-theoretical context that goes back to Zabell (1992), Skyrms (2010), and Huttegger (2017). Second, Hoppe's urn is a plausible model in that it is easy for agents to discover new states when they know little about the world and in that agents are more likely to observe a particular state again once they have already observed it. For example, children run up against many states that are new to them as they explore the world—states that they are very likely to encounter again. Adults, on the other hand, are less likely to make new discoveries. Or consider the real-life scenarios that first motivated models of discovery: in the sampling of species, the number of species discovered in a given ecosystem grows rapidly at first but the rate of discovery decreases with the number of samples (Gotelli and Colwell, 2001). Finally, I show below that this model of discovery gives rise to state frequencies that approximate a Zipfian distribution. I also show that signal frequencies tend to follow state frequencies. Although this may not be the only way to get a Zipfian distribution of signals, Hoppe's urn therefore provides a plausible representation of discovery processes.

In nature, conventional systems of communication are often and in large part acquired through learning. Given the goal of relating sender-receiver models to real-world data, I follow Skyrms (2010) in assuming that agents learn according to Roth and Erev's (1995) model of differential reinforcement. In the simplest version of this model, the probability with which an agent chooses to perform an action changes in proportion to the accumulated past reward of that action. The agent starts out at time t by giving weight $q_i(t) = 1$ to an action i. Upon choosing an action, the weight given to that action is updated to $q_i(t+1) = q_i(t) + 1$ if the action is rewarded; the weight does not change otherwise. Regardless of whether the action is rewarded, the weight given to other actions remains unchanged. The probability with which the agent chooses a particular action i is then given by dividing $q_i(t)$ by the sum of $q_j(t)$ for all actions j.

But language users also invent new expressions and forget old ones. To represent the invention of signals, I follow Alexander et al.'s (2012) model. This model is also a version of Hoppe's urn but it is Hoppe's urn supplemented with differential reinforcement. At first, the sender has no signals and so must choose a signal from the set $M = \{m_0\}$. The signal m_0 is the "inventor" signal. If the sender chooses the inventor signal, the sender invents a new signal. Upon receiving the newly invented signal, the receiver picks an act. For simplicity, I assume that there are exactly as many acts as states: whenever a state is discovered, a new act is added to the receiver's repertoire of acts. I also assume that there is a single successful act with a positive payoff in every state. If the sender invents a new signal and the receiver picks the successful act, then the sender adds the new signal to the set of available signals. The sender also reinforces the new signal, and the receiver reinforces the corresponding act. If the sender invents a new signal and the act is unsuccessful, no signal is added to the set of signals. When the sender sends a known signal and the act is successful, sender and receiver again reinforce their choices; sender and receiver do not reinforce their choices if the act is unsuccessful. The inventor signal is never destroyed but also never reinforced. For every state, there is thus always a positive and yet decreasing probability that the sender invents a new signal.

It is again helpful to think of an urn. Initially, the sender has no urns since there are no states. Whenever the sender discovers a state, the sender creates an urn corresponding to that state. As an urn represents a known state of nature, the sender eventually acquires an urn for each state of nature that has been discovered. The sender then observes the state of nature, picks the urn corresponding to that state, and draws a ball at random from that urn. If the sender picks the black ball (i.e., the inventor signal), the sender invents a new signal. Whenever the new signal leads to a successful act, the sender returns the black ball and adds a ball of a novel color to the urn. If the urn already contains balls other than the black, there is also a chance that the sender picks a colored ball and sends the corresponding signal to the receiver. Whenever the signal leads to a successful act, the sender reinforces that signal by returning the original ball and adding another ball of the same color to the urn.

The receiver also starts the game with no urns. Whenever the sender invents a new signal, the receiver creates a new urn. The receiver's urns correspond to the sender's signals. In each one of the receiver's urns, there are balls of different colors. Each color represents a different type of act. Upon receiving a signal (newly invented or otherwise), the receiver chooses the corresponding urn and then picks a ball at random. If the color of the ball matches the color of the ball that nature chose, then the act is successful. Whenever the act is successful, the receiver returns the ball and adds another ball of the same color to the urn. If the act is not successful, then the receiver returns the balls to the original urn but does not add another ball.

To represent the forgetting of signals, I assume with Alexander et al. (2012) that at the end of every round there is a certain probability (f) that a signal is forgotten. Once the decision has been made to forget a signal, the sender chooses to forget a signal according to a forgetting rule. Alexander et al. (2012) consider two forgetting rules: Forgetting A, and Forgetting B. According to Forgetting A, we first pick a sender's urn at random, then select a ball from that urn at random, and finally discard that ball. This partially decreases the reinforcement level for that ball color, which corresponds to a signal type. According to Forgetting B, we first pick a sender's urn at random, then select a ball color at random, and finally discard one ball of that color from the urn. This again partially decreases the reinforcement level for that ball color. Alexander et al. (2012) find that Forgetting B improves the efficiency of signaling between sender and receiver. But Forgetting A does not promote efficiency, as it leads to the proliferation of signals that are often left unused.

As agents should not be expected to store in memory signals that are hardly ever sent, I ignore Forgetting A and consider instead Forgetting B. I propose two variants of Forgetting B: Random Forgetting (B-RF), and Reinforcement-Dependent Forgetting (B-RDF). As both B-RF and B-RDF are versions of Forgetting B, we first choose a sender's urn at random, then select a ball color, and finally discard one ball of that color from the urn. But the two variants diverge as to how ball colors are chosen. According to B-RF, we choose a ball color at random. This means that all ball colors—i.e., all signal types—are equally likely to be chosen regardless of how many balls of that color there are in the urn. According to B-RDF, we choose a signal to remove with probability inversely proportional to its reinforcement level in the chosen urn. That is, signals with a high level of reinforcement are less likely to be chosen than signals with a low level of reinforcement. According to both rules, a signal is completely removed from the sender's repertoire and the corresponding receiver urn is destroyed whenever its reinforcement level reaches zero in all states. In either case, the inventor signal is never destroyed and it is thus always possible for the sender to invent a new signal.

Other forgetting rules are in principle possible (Barrett and Zollman, 2009). To assess whether results are robust to changes in the forgetting rule, I also consider another type of forgetting: Forgetting C. As with Forgetting B, a signal is forgotten with a certain probability (f) at the end of every round. But according to Forgetting C, sender and receiver completely remove a signal from their repertoire once the decision has been made to forget a signal. Completely removing signals from their repertoire may seem too extreme, especially in comparison to the fact that Forgetting B destroys a signal only if its reinforcement level reaches zero in all states. But I choose to implement Forgetting C in this manner because we are interested in what happens when communication breaks down. That is, we want to know whether results are robust to variations in the forgetting rule that are so extreme that no communication takes place. There are also two variants of Forgetting C: Random Forgetting (C-RF), and Reinforcement-Dependent Forgetting (C-RDF). According to C-RF, the sender chooses a signal to remove at random; according to C-RDF, the sender chooses a signal to remove with probability inversely proportional to its reinforcement level in the chosen urn. In both cases, once a signal is marked for destruction, it is removed from the sender's and the receiver's repertoire. The inventor signal is never destroyed, so it is always possible to invent new signals.

A few observations are in order before proceeding. First, it should be noted that Forgetting C is an extreme type of forgetting in that it prunes signals very severely. One consequence of this is that, unlike Forgetting A, Forgetting C does not lead to the proliferation of signals that are often left unused. Another consequence is that Forgetting C can even render communication impossible. As I show in the next section, this is important because it allows us to compare results that obtain with communication to those that obtain without it.

Second, the contrast between RF and RDF in both their B and C forms is especially consequential for the purposes of this paper. This is because B-RDF and C-RDF implement rules that are the reinforcement-learning analogs of frequency-dependent selection. With frequency-dependent selection, the probability that an individual retains a copy of a trait depends on the frequency of that trait. With RDF, the probability that sender and receiver retain a copy of a signal also depends on the frequency of that signal since that probability is proportional to the reinforcement level of that signal. In this respect, RDF is therefore analogous to frequency-dependent selection. In the case of RF, on the other hand, the probability that sender and receiver retain a copy of a signal does not depend on the reinforcement level and so does not depend on the frequency of that signal. With RF, the probability that sender and receiver retain a copy of a signal is distributed uniformly across signals. RF is therefore analogous to random drift. So whereas B-RDF and C-RDF represent the case of frequency-dependent selection in a reinforcement-learning setting, B-RF and C-RF represent a neutral baseline.

Third, it is important to keep in mind that despite representing invention, forgetting, and learning in a dynamic state space, the present model contains many idealizations. For one, it makes simplifying assumptions about how agents discover new states: both sender and receiver always know the same states, the rate with which they discover new states decreases over time, and observing a state makes it more likely for the sender to observe the same state again. For another, the model assumes a one-to-one correspondence between states and acts since the receiver always has as many acts to choose from as there are states of nature. It also assumes that there is a single act in every state that yields a positive payoff. Finally, it assumes that agents have no conflicting interests as they always receive the same payoff. The following results should therefore be considered with these caveats in mind.

3 Results

It can be shown analytically that in the model studied here the number of states increases without bound and at a decreasing rate. Consistent with this, computer simulations show that the total number of states is a strictly increasing function of time, that the rate of discovery increases very fast at first, and that the rate slows down with time (Figure 1). The rate of discovery decreases because even though the probability of nature choosing the mutator state is never zero, nature chooses from an increasingly large set of states. This decreases the probability with which nature chooses the mutator state, thus lowering the rate of discovery. Nature also reinforces states other than the mutator state after observing them, which increases the probability of nature choosing known states and decreases the rate of discovery.

As for the number of signals, results differ slightly depending on the type of forgetting. With Forgetting B, the number of signals currently in use increases at first but eventually ceases to increase and plateaus at different levels depending on the forgetting rate. For lower rates of forgetting ($f = 0.1$), agents are able to maintain a higher number of signals than for intermediate and higher rates ($f = 0.3$ and $f = 0.9$). These results hold for B-RF and B-RDF. With Forgetting C, on the other hand, similar results hold only for intermediate forgetting rates. For lower forgetting rates ($f = 0.00001$), more signals are invented than forgotten. As a result, the number of signals in use strictly increases over time. For higher forgetting rates ($f = 0.1$), the number of signals does not increase monotonically over time since agents forget old signals before they can invent new ones. Given that Forgetting C prunes signals much more severely than Forgetting B, we consider different forgetting rates—lower for the former, higher for the latter—to ensure that results can still be compared across the two types of forgetting.

But regardless of forgetting type, forgetting rate, and number of signals, the number of states increases over time. This increases the complexity of the

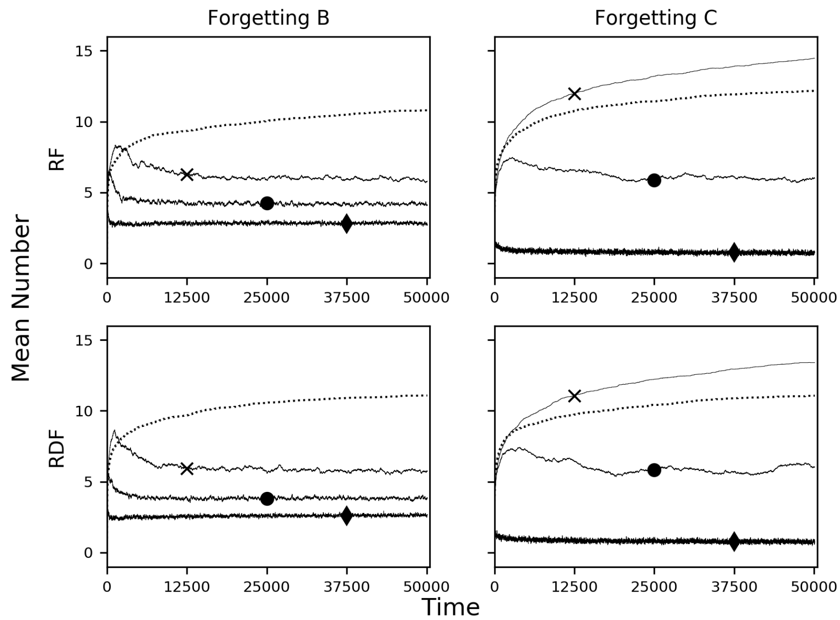

Figure 1: Mean number of states (*dotted lines*) and signals (*solid lines*). Left panel: Forgetting B with forgetting rules B-RF (*top*) and B-RDF (*bottom*) and forgetting rates equal to 0.1 (*cross*), and 0.3 (*circle*), 0.9 (*diamond*). Right panel: Forgetting C with forgetting rules C-RF (*top*) and C-RDF (*bottom*) and forgetting rates equal to 0.00001 (*cross*), and 0.001 (*circle*), 0.1 (*diamond*). Results are average of 100 runs, each with 50,000 iterations.

coordination problem that sender and receiver now face. As a result, sender and receiver must keep a sufficiently high number of signals in their repertoire if they are to communicate successfully. For a given forgetting rate and forgetting type, the number of signals that sender and receiver maintain should therefore affect their ability to communicate in the face of an increasingly complex state space.

This is what we in fact observe. With Forgetting B, sender and receiver already attain a high success rate after about $5,000$ iterations (Figure 2). Success rate is given here by the expected payoff of the agents. Although the success rate is at first visibly higher when the forgetting rate is lower or intermediate ($f = 0.1$ and $f = 0.3$), the success rate reaches a high level for higher forgetting rates as well ($f = 0.9$). The success rate is also high with Forgetting C if the forgetting rate is not very high. With Forgetting C and a higher forgetting rate ($f = 0.1$), signals are forgotten all too often and communication collapses. Notice also that the success rate remains largely unaffected by the forgetting rule.

That is, agents achieve similar success rates under B-RF and B-RDF; the same holds for C-RF and C-RDF. But the success rate is sensitive to the forgetting rate. In all forgetting regimes, the success rate goes down with increasing rates of forgetting. In the case of Forgetting C, moreover, there seems to be a critical threshold for the forgetting rate above which sender and receiver are no longer able to communicate. This can be seen from the fact that sender and receiver successfully learn to communicate for $f = 0.00001$ but perform no better than chance for $f = 0.1$. At some intermediate rate of forgetting, communication breaks down. A threshold effect is not observed for Forgetting B, as this is a less severe form of forgetting than Forgetting C; sender and receiver therefore perform much better than chance even for very high rates of forgetting.

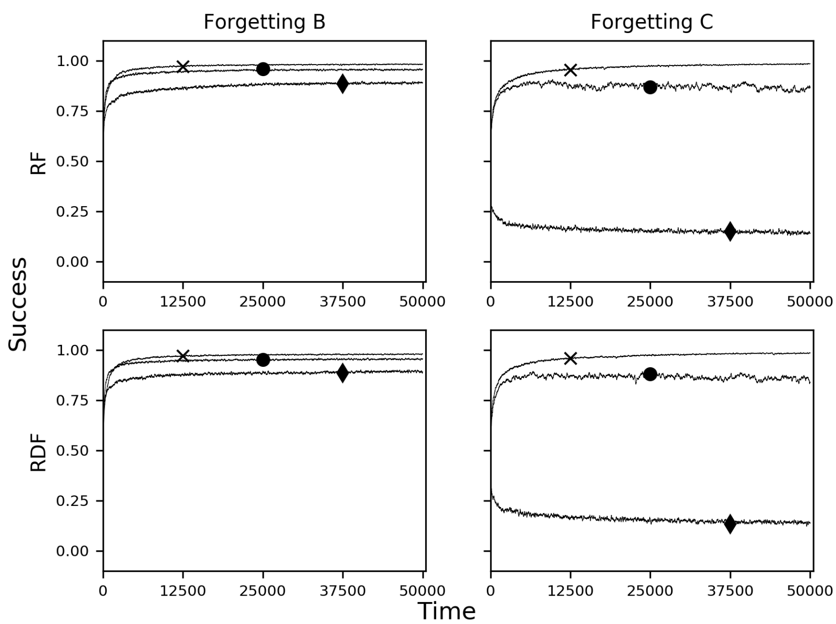

Figure 2: Success rate. Left panel: Forgetting B with forgetting rules B-RF (*top*) and B-RDF (*bottom*) and forgetting rates equal to 0.1 (*cross*), and 0.3 (*circle*), 0.9 (*diamond*). Right panel: Forgetting C with forgetting rules C-RF ((*top*)) and C-RDF (*bottom*) and forgetting rates equal to 0.00001 (*cross*), and 0.001 (*circle*), 0.1 (*diamond*). Results are average of 100 runs, each with 50,000 iterations.

These results stand in contrast to Alexander (2014). In a model in which the correct state-action pair is randomly swapped, signaling evolves and persists with reinforcement learning and discounting the past provided that the swapping rate is kept to a minimum. But in a model in which new states arrive at a constant rate, Alexander shows that the success rate converges to zero if new states are introduced at or above a certain rate. As agents encounter a constantly growing number of states, this eventually overwhelms their capacity to learn. In my model, however, the rate with which new states are discovered goes down with time. Agents thus face an easier coordination problem. This difference in modeling assumption makes it possible for communication to emerge across a wide range of parameter values.

My model also recreates striking features observed in natural languages. Since Zipf (1949), the frequency of word use has been known to be inversely proportional to word rank. This relationship is robust across languages and word categories, such as determiners, prepositions, verbs, and different classes of nouns (e.g. number words, names of chemical elements, and taboo words)—for a recent review, see Piantadosi (2014). Similar findings have been reported in some cases of animal communication, although this remains controversial (McCowan et al., 1999, 2005; Suzuki et al., 2005). In English, large-scale studies confirm that the frequency of word use approximates Zipf's law (Moreno-Sánchez et al., 2016). In particular, the frequency of word use has been found to best approximate a distribution given by $f(k) = \frac{1}{k(s-1)} - \frac{1}{(k+1)(s-1)}$, where $f(k)$ is the frequency of the k^{th} word, k is a word's rank, and s is a parameter in the interval $[0, \infty)$ that determines the concavity of the distribution (lower values produce approximately uniform and higher values produce more concave distributions).

To determine whether the present model can give rise to Zipf's inverse power law, two regimes for the reinforcement of states were considered. Under Preferential Reinforcement (PR), a state is reinforced whenever it is observed. PR assumes that once a state is observed it is more likely to be observed again, just as in Hoppe's original urn model. Under Random Reinforcement (RR), on the other hand, every time a state is observed any state is just as likely to be reinforced as the observed one. Clearly, RR is not a realistic scenario. But I include it here to contrast it with PR, show how the distribution of signal frequencies tends to follow the distribution of state frequencies in both the RR and RF regime, and thus provide a possible explanation for the Zipfian distribution of signal frequencies.

Results show that PR is remarkably successful at recreating Zipf's inverse power law, while RR is not (Figure 3). After collecting the frequency of all signals used over $50,000$ rounds in the PR regime, signal frequency was indeed inversely proportional to signal rank. In particular, the observed frequencies of signal use closely followed the distribution predicted by the equation above. Although estimates for the parameter s vary slightly across datasets, the distribution of signal frequencies observed in the model approximates well the predicted distribution when $s = 2$ for a wide range of forgetting rates, different forgetting rules (RF and RDF), and both forgetting types (Forgetting B and C). This parameter value is close to estimates between 1.9 and 2.1 reported for natural language corpora (Moreno-Sánchez et al., 2016).

As expected, the distribution of state frequencies is similarly Zipfian under

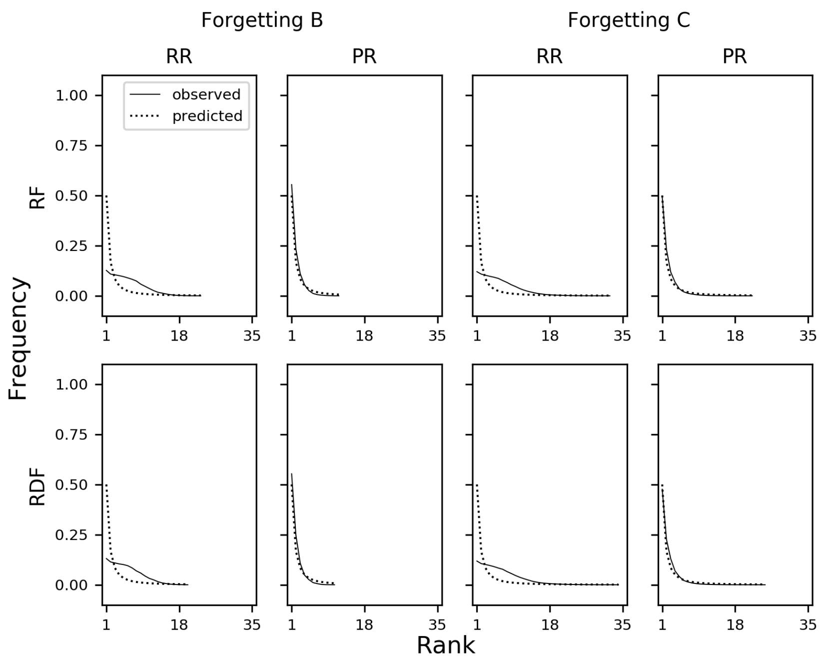

Figure 3: Observed (*solid lines*) and predicted (*dotted lines*) distribution of signal frequencies. Left and middle-left panels: Forgetting B with forgetting rules B-RF (*top*) and B-RDF (*bottom*) and forgetting rate equal to 0.1. Middle-right and right panels: Forgetting C with forgetting rules C-RF (*top*) and C-RDF (*bottom*) and forgetting rate equal to 0.00001. Results are average of 100 runs, each with 50,000 iterations.

PR but not under RR regardless of forgetting type. But sender and receiver do learn to communicate under RR with both Forgetting B and Forgetting C, even though the distribution of state frequencies does not approximate Zipf's law under this reinforcement rule (Figure 4). However, learning happens more slowly and the success rate appears to be slightly more sensitive to the forgetting rate than under PR. This is true with both Forgetting B and Forgetting C. For both types of forgetting (Forgetting B and C), the forgetting rules RF and RDF do not seem to affect the success rate.

Given that new signals can be invented and old ones forgotten, signals can also change. When a signal is lost, a new signal may become associated with the same state that the lost signal used to be associated with. This resembles a process of linguistic replacement in which language users develop neologisms and words fall into disuse. Empirical studies suggest that replacement rates follow predictable patterns across languages. In an oft-cited paper, Pagel et al. (2007) find that the replacement rate depends on the frequency of word use across

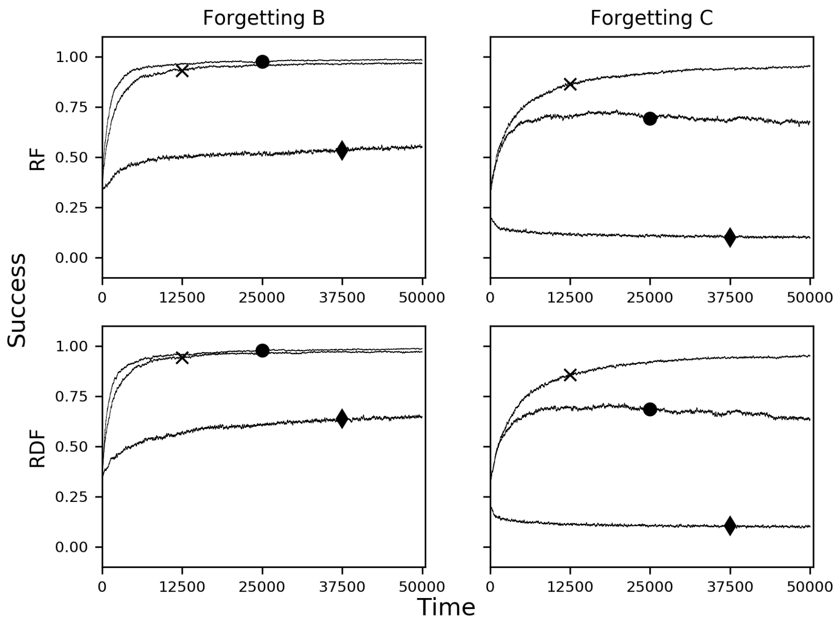

Figure 4: Success rate with Random Reinforcement. Left panel: Forgetting B with forgetting rules B-RF (*top*) and B-RDF (*bottom*) and forgetting rates equal to 0.1 (*cross*), and 0.3 (*circle*), 0.9 (*diamond*). Right panel: Forgetting C with forgetting rules C-RF (*top*) and C-RDF (*bottom*) and forgetting rates equal to 0.00001 (*cross*), and 0.001 (*circle*), 0.1 (*diamond*). Results are average of 100 runs, each with 50,000 iterations.

different languages. In particular, their study shows that words are replaced in inverse proportion to their frequency of use, as more common words evolve at slower rates. Lieberman et al. (2007) report similar results, with common English verbs changing to regular forms at a slower rate than rare ones.

A consequence of these findings is that the average age of signals can be taken as a proxy for their rate of change. That is, if frequently used signals change at a slower rate than infrequently used ones, then frequently used signals should be on average older than their less frequently used counterparts. This is in line with Zipf's (1949) original observation that word age seems to be inversely proportional to the frequency of word use.

To find out whether the present model can recreate the relationship between frequency of word use and word age observed in natural languages, the age of each signal was recorded. In the model, signal age denotes the round in which the signal was first introduced. After collecting the age of all signals after $50,000$ rounds of interaction, signal age was normalized to a scale in the unit interval and then scaled by the total number of iterations. In this scale, the largest possible value denotes the age of a signal invented in the first round of interaction (oldest) and the lowest possible value denotes the age of a signal invented in the last round of interaction (youngest).

I then conducted a regression analysis of the resulting signal age on the frequency with which each signal was used throughout the entire run. As expected, commonly used signals were on average older than infrequently used ones for both Forgetting B and Forgetting C (Figures 5 and 6). The association between signal use and signal age was found to be robust across variations in the forgetting rule, as both RF and RDF led to qualitatively similar results. But the association depends on signals actually carrying information: with Forgetting C and a forgetting rate that is too high for communication to take place, signal age does not correlate with frequency of use.

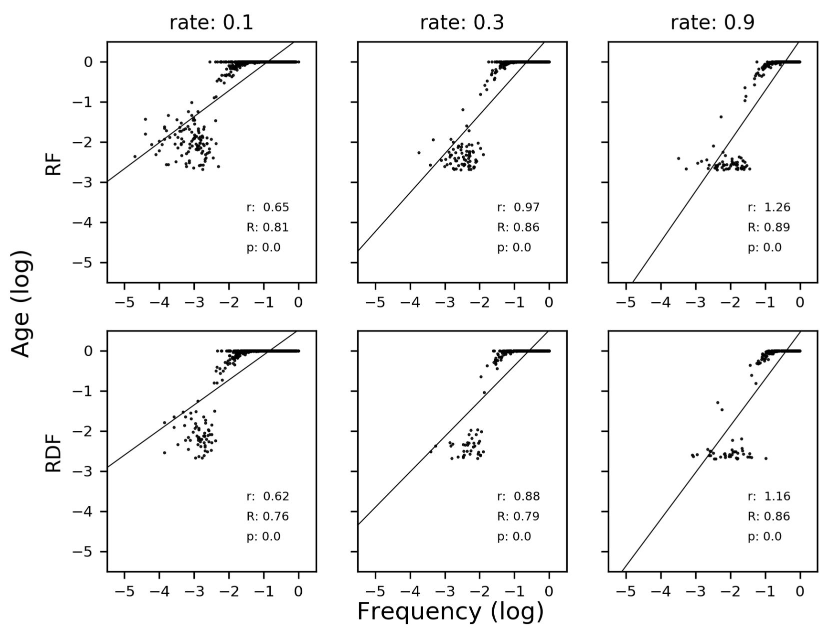

Figure 5: Log-log regression of signal age on signal frequency with B-RF (*top*) and B-RDF (*bottom*) and forgetting rates equal to 0.1 (*left*), and 0.3 (*middle*), 0.9 (*right*). Regression line is given by $log(y) = a + rlog(x)$. Values are regression coefficients (r), associated p-values (p), and adjusted R^2 (R). Results are for data from 100 runs with PR, each with 50,000 iterations.

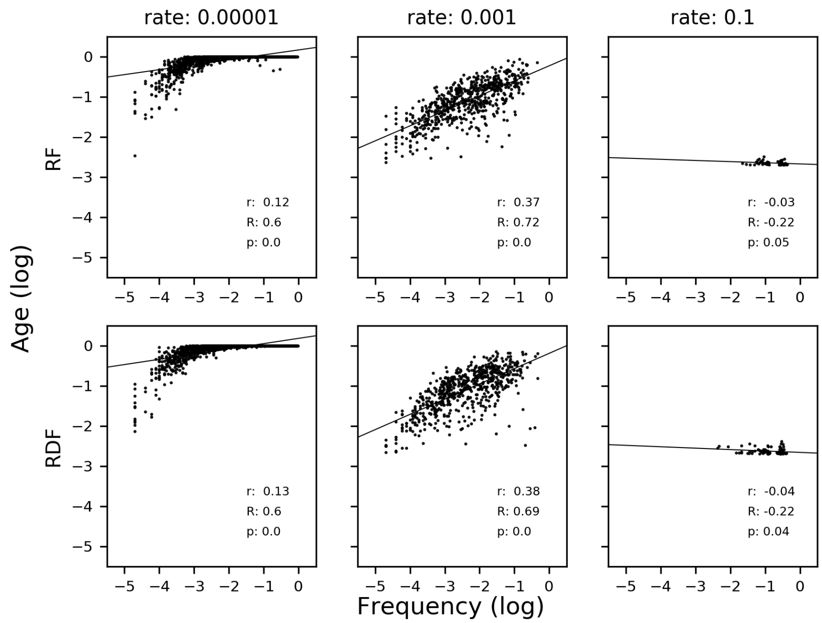

Figure 6: Log-log regression of signal age on signal frequency with C-RF (*top*) and C-RDF (*bottom*) and forgetting rates equal to 0.00001 (*left*), and 0.001 (*middle*), 0.1 (*right*). Regression line is given by $log(y) = a + rlog(x)$. Values are regression coefficients (r), associated p-values (p), and adjusted R^2 (R). Results are for data from 100 runs with PR, each with 50,000 iterations.

4 Discussion

While it is reasonable to suppose that the need for communication sometimes arises when agents inhabit a world with a dynamic state space, previous sender-receiver models tend to ignore this issue. Alexander (2014) is an exception. When the correct state-action pair is randomly swapped at a low rate, Alexander shows that reinforcement learning with discounting the past allows for signaling to evolve and persist. He also shows that the ability of agents to signal is overwhelmed when the number of states increases without bound at a constant rate. This may be taken to suggest that it is difficult for signaling to evolve with a growing number of states. However, the present model shows that communication can easily evolve even when agents face the two-fold challenge of navigating an increasingly complex environment and learning to signal. In situations like this, successful communication depends on a fine balance between the rate at which agents discover new states and their ability to maintain a sufficiently large repertoire of signals. The balance can be achieved if the probability of discovering a new state goes down with the number of known states so that agents ultimately face an easier problem than in Alexander's (2014) model.

More importantly, these results indicate that sender-receiver models can be brought to bear on observational data. In particular, the model proposed here recreates Zipf's inverse power law—the pattern between frequency of word use and word rank observed in natural languages. Additionally, the model correctly predicts the relationship observed between frequency of word use and word age. Given that sender-receiver models were first introduced to show how referential meaning could emerge by convention, it is remarkable to find that these models recreate features of human languages other than referentiality—such as frequency of word use and evolutionary tempo. Not only does this increase our confidence in the results that sender-receiver models sought to provide in their original domain of application, but it also provides a possible mechanism for how these features arise.

This is important because studies in evolutionary linguistics tend to focus on large-scale patterns that are the product of biological and cultural evolution. But this often means that the corresponding low-level processes remain poorly understood. Pagel et al. (2007), Lieberman et al. (2007), and Moreno-Sánchez et al. (2016) are cases in point: their studies describe large-scale patterns but do not shed light on the underlying mechanism. In contrast, sender-receiver models can offer insights into the processes that underpin the observed patterns. This could ultimately contribute to the synthesis that Mesoudi (2011) rightly advocates "between the microlevel, that is, those small-scale, individual-level processes that act to change the frequency of culturally transmitted traits within a single population, and the macrolevel, that is, large-scale patterns and trends at or above the level of entire societies" (p. 51).

One of the most striking large-scale patterns in linguistics is Zipf's inverse power law: the inverse correlation between frequency of word use and word rank first noted by Zipf (1949). Although ample empirical evidence supports it, there is no consensus on a psychologically compelling and theoretically motivated account of its origin (Piantadosi, 2014). For example, several models show how Zipf-like distributions can arise from purely random or stochastic processes (Li, 1992; Zanette and Montemurro, 2005). Such models successfully recreate Zipf's law but offer no reason to think that language users produce signs according to random or stochastic processes. These accounts therefore lack any psychological plausibility. Other attempts to explain Zipf's law invoke principles of optimization and efficiency in information transmission (i Cancho and Solé, 2003; Salge et al., 2015). This is in line with Zipf's (1949) original account. But these proposals fail to account for the fact that words with similar meanings occur with similar frequencies across distantly related languages (Calude and Pagel, 2011; Dehaene and Mehler, 1992). A full account of Zipf's law should therefore also appeal to the semantic properties of language. The model proposed by Manin (2008) attempts to do just that, as it derives Zipf's law from the assumption that meanings are arranged hierarchically and that there is pressure to avoid excessive synonymy. But being designed to uniquely explain Zipf's law, Manin's model is extremely narrow in scope.

In contrast, the sender-receiver model proposed here provides a possible mechanism for the emergence of Zipf's law that is psychologically plausible, has a very general scope, and takes into account semantic properties of language. The model is psychologically plausible because it explicitly represents agents following simple learning and forgetting rules to coordinate on signal-act pairs and thus learn to communicate. Other existing accounts of Zipf's law do not explicitly represent language users. The present model is also very general in that the sender-receiver framework has been used in a wide variety of contexts to account for very disparate phenomena—for an overview of sender-receiver models and some applications, see Skyrms (2010). Finally, the model takes into account semantic properties of language to explain the emergence of Zipf's law.

To see why, note that signals approximate Zipf's law only when the distribution of states is similarly Zipfian. When states are reinforced according to PR, the frequency of states and the frequency of signals follow Zipf's distribution; when states are reinforced according to RR, neither the frequency of states nor the frequency of signals is Zipfian. Given that built into the PR regime is a mechanism to generate state frequencies that follow a particular distribution, the fact that the distribution of signal frequencies is roughly Zipfian in this regime is by itself not very surprising. What is interesting, however, is that signal frequencies mirror state frequencies in both regimes. This suggests that a full explanation of Zipf's inverse power law must account not just for word frequencies, but also for features of the world that are external to language. An explanation of Zipf's inverse power law that appeals to language-external features of the world would also be consistent with the finding that words with similar meanings occur with similar frequencies across different languages: if the distribution of word frequencies is not specific to a particular language, then the distribution of word frequencies should reflect features of the world that are shared across languages.

What could these language-external features of the world be? One possibility is that the distribution of word frequencies simply follows the distribution of frequencies with which language users encounter different states of the world—see Barrett and LaCroix (2020) for a related point. In the present model, signal frequencies directly follow state frequencies. This could be taken to indicate that word frequencies are not tied down to language-specific factors because they track the frequencies with which language users encounter their referents. But the frequency of referents cannot be the only or even the main factor to determine the frequency of the words used to pick out those refererents. In natural languages, the frequency of words also seem to respond to the relevance of their referents. As an anonymous referee for this journal points out, the word "gold" occurs roughly three times more often than "dirt" in the Corpus of Contemporary American English. It is unlikely that this is because most of us encounter gold more often than dirt. This difference in word frequency likely reflects a difference in the relevance of their referents, and not simply a difference in their frequencies.

The present model, however, cannot easily distinguish between the relevance and the frequency of different referents. A natural way to think of relevance in the context of a sender-receiver model is in terms of payoff, but the payoff that sender and receiver get from communicating is the same in every state. As some states occur more often than others, sender and receiver get an overall higher payoff when they can communicate in common states as opposed to in rare ones. So this model cannot determine whether signal frequencies track state frequencies simply because of the frequency with which each state occurs or because of the total payoff that sender and receiver get in each state. Yet, my main point remains that a full account of the distribution of signal frequencies should appeal to language-external factors. All else being equal, these factors may be the raw frequency of states. When all else is not equal, they may include relevance- or payoff-weighted frequencies of states—or, in the apt formulation of the same anonymous referee, "pragmatically-motivated situations".

The present model also provides a possible account of why common words change less rapidly than rare ones. Recall the two forgetting rules used in the model. Both B-RF and C-RF ignore reinforcement level, as every signal is lost with the same probability regardless of reinforcement level. Given that reinforcement level positively correlates with frequency of use, RF does not make the rate with which signals are forgotten dependent on their frequency of use. On the other hand, B-RDF and C-RDF take reinforcement level into consideration as signals are forgotten depending on reinforcement level. As reinforcement level positively correlates with frequency of use, RDF makes the rate with which signals are forgotten dependent on their frequency. But the observed correlation between frequency of use and evolutionary tempo does not depend on the forgetting rule, as both RF and RDF give rise to the correlation. This suggests that the correlation may not be due to frequency-dependent rates of change. Contrary to what Pagel et al. (2007) claim, similar patterns in natural languages may therefore not be due to "linguistic, frequency-dependent, purifying selection" (p. 719) against lexical change. After all, both RF and RDF give rise to a similar pattern in the model. But while RF makes the rate of change dependent on the frequency of signal use, RDF does not.

An alternative explanation is that rare words change more rapidly because they are subject to stronger drift. At this point, it is worth noting that Lynch (2011) offers a similar mechanism to explain the high error rates observed in infrequently used polymerases. In periods of cellular stress, rarely active and error-prone polymerases are recruited to repair DNA lesions. As Lynch points out, this has led many to hypothesize that natural selection may constrain the efficiency of lesion-repair polymerases and thus increase rates of evolutionary change during stressful periods. But as these polymerases are only infrequently activated, Lynch proposes instead that their low frequency of use simply exposes them to stronger drift as they are less often under the action of frequency-independent selection. As a result, higher error rates and higher rates of change can persist when these enzymes are at work. This accords with the expectation that the rate of evolutionary change is likely to increase when evolutionary forces are less pronounced (McShea and Brandon, 2010), as it is the case at low frequencies.

In principle, then, it seems that both selection and drift could account for frequency-dependent rates of evolutionary change. When purifying selection is stronger on more frequent traits, it can cause frequent traits to exhibit a lower rate of change. But given that more frequent traits tend to be under weaker drift as well, their lower rate of change could also simply be due to drift. In the former case, the rate of change depends on selection that is frequency-dependent in the sense that it is stronger on more frequent traits; in the latter, there is also a frequency effect but this effect is the product of drift. Whether it is drift or selection that is responsible for frequency-dependent rates of change is therefore something that should be shown rather than assumed. Following Lynch's (2011) lead in the case of frequently and infrequently used polymerases, I have tried to show here that a plausible explanation for why infrequently used words have a higher rate of change is that they are under stronger drift. The same may be true of other rare cultural traits that display a high evolutionary tempo.

5 Conclusion

This study proposes an extension of the sender-receiver model to account for the emergence of communication in a dynamic world, validates the model against observational data, and provides a possible mechanism for Zipf's inverse power law relating the frequency of word use and word rank and for the negative correlation between the frequency of word use and rate of lexical change. In the face of other empirical studies on large-scale features of natural languages (Atkinson et al., 2008; Dunn et al., 2011), it also hopes to pave the way for future attempts to integrate modeling approaches and empirical data in philosophy and linguistics. For example, multiplayer versions of the sender-receiver model in which populations of agents give rise to daughter populations could be used to help us better understand large-scale patterns such as the "punctuational bursts" reported in Atkinson et al. (2008).

More generally, studies of this sort also help underscore the importance of integrating modeling techniques and data-based approaches in other areas of philosophy. The history and philosophy of science is a particularly promising case. In recent years, a large number of studies in this field have turned to empirical data. Take, for example, Byron (2007), Wray (2010), Machery and Cohen (2011), Overton (2013), or Weingart (2015). Common to these studies is the use of data to address questions in the history and philosophy of science.[1] At the same time, modeling techniques have also been extensively applied to issues in the history and philosophy of science—for a few illustrative examples, see Weisberg and Muldoon (2009), Zollman (2010), Alexander (2013), and Grim et al. (2019). Yet, while data-driven and model-based approaches have for the most part been pursued independently of one another, it is clear that models should ultimately face the tribunal of experience (Martini and Pinto, 2017). Ensuring that models make contact with real-world data is a straightforward way of doing so. Conversely, data-driven studies stand to benefit from the use

[1] Pence and Ramsey (2018) offer an overview of data repositories, analytical tools, and related challenges in this emerging field.

of agent-based and mathematical models that offer a deeper understanding of empirical patterns. The present study thus hopes to encourage a closer dialogue between data-driven and model-based approaches not only to the emergence of meaning and the evolution of language, but also to questions in other areas of philosophy as well.

Acknowledgement

I would like to thank Brian Skyrms, Louis Narens, Cailin O'Connor, Simon Huttegger, Jeffrey Barrett, Travis LaCroix, and other members of the Social Dynamics Seminar at UC Irvine for their comments on an earlier version of this paper. I am equally grateful for the organizers and the audience of the conference "Generalized Theory of Evolution" in Düsseldorf. I would also like to thank Hannah Read for her numerous comments on previous versions of this paper.

References

Alexander, J. M. (2013). Preferential attachment and the search for successful theories. *Philosophy of Science*, 80(5):769–782.

Alexander, J. M. (2014). Learning to signal in a dynamic world. *The British Journal for the Philosophy of Science*, 65(4):797–820.

Alexander, J. M., Skyrms, B., and Zabell, S. L. (2012). Inventing new signals. *Dynamic Games and Applications*, 2(1):129–145.

Atkinson, Q. D., Meade, A., Venditti, C., Greenhill, S. J., and Pagel, M. (2008). Languages evolve in punctuational bursts. *Science*, 319(5863):588–588.

Barrett, J. and Zollman, K. J. (2009). The role of forgetting in the evolution and learning of language. *Journal of Experimental & Theoretical Artificial Intelligence*, 21(4):293–309.

Barrett, J. A. and LaCroix, T. (2020). Epistemology and the Structure of Language. *Erkenntnis*, pages 1–15. Publisher: Springer.

Blume, A., DeJong, D. V., Kim, Y.-G., and Sprinkle, G. B. (1998). Experimental evidence on the evolution of meaning of messages in sender-receiver games. *The American Economic Review*, 88(5):1323–1340.

Blume, A., DeJong, D. V., Kim, Y.-G., and Sprinkle, G. B. (2001). Evolution of communication with partial common interest. *Games and Economic Behavior*, 37(1):79–120.

Bruner, J., O'Connor, C., Rubin, H., and Huttegger, S. M. (2014). David Lewis in the lab: experimental results on the emergence of meaning. *Synthese*, pages 1–19.

Byron, J. M. (2007). Whence philosophy of biology? *The British Journal for the Philosophy of Science*, 58(3):409–422.

Calude, A. S. and Pagel, M. (2011). How do we use language? Shared patterns in the frequency of word use across 17 world languages. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 366(1567):1101–1107.

Dehaene, S. and Mehler, J. (1992). Cross-linguistic regularities in the frequency of number words. *Cognition*, 43(1):1–29.

Dunn, M., Greenhill, S. J., Levinson, S. C., and Gray, R. D. (2011). Evolved structure of language shows lineage-specific trends in word-order universals. *Nature*, 473(7345):79.

Fisher, R. A., Corbet, A. S., and Williams, C. B. (1943). The relation between the number of species and the number of individuals in a random sample of an animal population. *The Journal of Animal Ecology*, pages 42–58.

Gotelli, N. J. and Colwell, R. K. (2001). Quantifying biodiversity: procedures and pitfalls in the measurement and comparison of species richness. *Ecology letters*, 4(4):379–391.

Grim, P., Singer, D. J., Bramson, A., Holman, B., McGeehan, S., and Berger, W. J. (2019). Diversity, ability, and expertise in epistemic communities. *Philosophy of Science*, 86(1):98–123.

Hoppe, F. M. (1984). Pólya-like urns and the Ewens' sampling formula. *Journal of Mathematical Biology*, 20(1):91–94.

Huttegger, S. M. (2017). *The Probabilistic Foundations of Rational Learning.* Cambridge University Press.

Huttegger, S. M., Skyrms, B., and Zollman, K. J. (2014). Probe and adjust in information transfer games. *Erkenntnis*, 79(4):835–853.

i Cancho, R. F. and Solé, R. V. (2003). Least effort and the origins of scaling in human language. *Proceedings of the National Academy of Sciences*, 100(3):788–791.

LaCroix, T. (2019). Evolutionary Explanations of Simple Communication: Signalling Games and Their Models. *Journal for General Philosophy of Science*, pages 1–25. Publisher: Springer.

Levy, A. (2011). Game Theory, Indirect Modeling, and the Origin of Morality. *The Journal of Philosophy*, 108(4):171–187.

Lewis, D. (1969). *Convention: A philosophical study.* Harvard University Press, Cambridge.

Li, W. (1992). Random texts exhibit Zipf's-law-like word frequency distribution. *IEEE Transactions on information theory*, 38(6):1842–1845.

Lieberman, E., Michel, J.-B., Jackson, J., Tang, T., and Nowak, M. A. (2007). Quantifying the evolutionary dynamics of language. *Nature*, 449(7163):713.

Lynch, M. (2011). The lower bound to the evolution of mutation rates. *Genome biology and evolution*, 3:1107–1118.

Machery, E. and Cohen, K. (2011). An evidence-based study of the evolutionary behavioral sciences. *The British Journal for the Philosophy of Science*, 63(1):177–226.

Manin, D. Y. (2008). Zipf's law and avoidance of excessive synonymy. *Cognitive Science*, 32(7):1075–1098.

Martini, C. and Pinto, M. F. (2017). Modeling the social organization of science. *European Journal for Philosophy of Science*, 7(2):221–238.

McCowan, B., Doyle, L. R., Jenkins, J. M., and Hanser, S. F. (2005). The appropriate use of Zipf's law in animal communication studies. *Animal Behaviour*, 69(1):F1–F7. Publisher: Academic Press.

McCowan, B., Hanser, S. F., and Doyle, L. R. (1999). Quantitative tools for comparing animal communication systems: information theory applied to bottlenose dolphin whistle repertoires. *Animal behaviour*, 57(2):409–419. Publisher: Academic Press.

McShea, D. W. and Brandon, R. N. (2010). *Biology's first law: the tendency for diversity and complexity to increase in evolutionary systems*. University of Chicago Press, Chicago, London.

Mesoudi, A. (2011). *Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences*. University of Chicago Press.

Moreno-Sanchez, I., Font-Clos, F., and Corral, A. (2016). Large-scale analysis of Zipf's law in English texts. *PloS one*, 11(1):e0147073.

Overton, J. A. (2013). "Explain" in scientific discourse. *Synthese*, pages 1–23.

Pagel, M., Atkinson, Q. D., and Meade, A. (2007). Frequency of word-use predicts rates of lexical evolution throughout Indo-European history. *Nature*, 449(7163):717–720. Publisher: Nature Publishing Group.

Pence, C. and Ramsey, G. (2018). How to do digital philosophy of science. *Philosophy of Science*, 85(5):930–941.

Piantadosi, S. T. (2014). Zipf's word frequency law in natural language: A critical review and future directions. *Psychonomic bulletin & review*, 21(5):1112–1130.

Roth, A. E. and Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, 8(1):164–212.

Rubin, H., O'Connor, C., and Bruner, J. (2019). Experimental Economics for Philosophers. In Fischer, E. and Curtis, M., editors, *Methodological Advances in Experimental Philosophy*. Bloomsbury Publishing.

Salge, C., Ay, N., Polani, D., and Prokopenko, M. (2015). Zipf's law: balancing signal usage cost and communication efficiency. *PLoS one*, 10(10):e0139475.

Seyfarth, R. M. and Cheney, D. L. (2017). The origin of meaning in animal signals. *Animal Behaviour*, 124:339–346. Publisher: Elsevier.

Silk, J. B., Kaldor, E., and Boyd, R. (2000). Cheap talk when interests conflict. *Animal behaviour*, 59(2):423–432.

Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford University Press, Oxford.

Suzuki, R., Buck, J. R., and Tyack, P. L. (2005). The use of Zipf's law in animal communication analysis. *Animal Behaviour*, 69(1):F9–F17. Publisher: Academic Press.

Wagner, E. (2009). Communication and structured correlation. *Erkenntnis*, 71(3):377–393.

Wagner, E. O. (2011). Deterministic chaos and the evolution of meaning. *The British Journal for the Philosophy of Science*, 63(3):547–575.

Weingart, S. B. (2015). Finding the History and Philosophy of Science. *Erkenntnis*, 80(1):201–213.

Weisberg, M. and Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of science*, 76(2):225–252.

Wray, K. B. (2010). Philosophy of science: What are the key journals in the field? *Erkenntnis*, 72(3):423–430.

Zabell, S. L. (1992). Predicting the unpredictable. *Synthese*, 90(2):205–232.

Zanette, D. and Montemurro, M. (2005). Dynamics of text generation with realistic Zipf's distribution. *Journal of quantitative Linguistics*, 12(1):29–40.

Zipf, G. K. (1949). *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology.* Addison-Wisley, Cambridge.

Zollman, K. J. (2005). Talking to neighbors: The evolution of regional meaning. *Philosophy of Science*, 72(1):69–85.

Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1):17.

Zollman, K. J., Bergstrom, C. T., and Huttegger, S. M. (2013). Between cheap and costly signals: the evolution of partially honest communication. *Proceedings of the Royal Society B: Biological Sciences*, 280(1750):1–8.