

Towards an ecology of vagueness

José Pedro Correia* Michael Franke

Abstract

A vexing puzzle about vagueness, rationality, and evolution runs, in crude abbreviation, as follows: vague language use is demonstrably suboptimal if the goal is efficient, precise and cooperative information transmission; hence rational deliberation or evolutionary selection should, under this assumed goal, eradicate vagueness from language use. Since vagueness is pervasive and entrenched in all human languages, something has to give. In this paper, we focus on this problem in the context of signaling games. We provide an overview of a number of proposed ways in which vagueness may come into the picture in formal models of rational and evolutionary signaling. Most argue that vague signal use is simply the best we can get, given certain factors. Despite the plausibility of the proposals, we argue that a deeper understanding of the benefits of vagueness needs a more ecological perspective, namely one that goes beyond the local optimization of signaling strategies in a homogeneous population. As an example of one possible way to expand upon our current models, we propose two variants of a novel multi-population dynamic of imprecise imitation where, under certain conditions, populations with vague language use dominate over populations with precise language use.

1 Vagueness and rationality

The classical philosophical problem of vagueness is most starkly embodied by the sorites paradox. The original formulation is attributed to Eubulides, an ancient Megarian philosopher (Sorensen, 2009), and uses the example of a heap of sand: if no removal of one grain of sand can make a heap into a non-heap, one can repeatedly remove all but one grain of sand from something that is clearly a heap and be forced to acknowledge that the remaining single grain of sand is still a heap; otherwise, it seems, one would have to accept that there is a determinate number of grains that forms a heap, and anything under it is not a heap. Neither choice is, however, intuitively satisfying. The paradox is interesting because it can be made general and re-applied to many other words besides ‘heap’. Predicates for which one can find a suitable instance of the general formulation of the sorites paradox are called *vague*. Paradigmatic examples besides ‘heap’ include ‘tall’, ‘red’, ‘bald’, ‘tadpole’, and ‘child’ (Keefe and P. Smith, 1997). How widespread is the problem? It is easy to find more examples of predicates based on more finely grained properties—as ‘tall’ is intuitively based on height—for which constructing a sorites paradox would be easy. Mereological nihilists argue that instances of the sorites paradox can be designed for any material object that can be decomposed into small enough parts. If one subscribes to the scientific picture of matter as composed of molecules and atoms, this applies to tables and chairs, cats and mats, and any other ordinary thing (Unger, 1979). Bertrand Russell famously argued (1923) that all words, including “the words of pure logic,” are vague when used by human beings.

If one thinks of language as governed by logical rules, and of rationality as including the ability to follow those rules, the sorites paradox seems *prima facie* to demonstrate that vagueness and rationality are incompatible. But one can think of language in different terms. One possibility is to think of signs as tools agents use to coordinate actions. An example of a way to formally study language along those lines is the game-theoretic framework of signaling games (Lewis, 1969/2011). Within such an approach, rationality can be seen as the ability to choose the use of signs that is optimal to achieve some form of coordination. The existence of vagueness in such models

*Holder of a PhD studentship by the *Fundação para a Ciência e a Tecnologia* (SFRH/BD/100437/2014).

would not be at odds with rationality as long as vague languages turn out to be optimal for the purposes at hand. However, as Barton Lipman (2009) argues in detail, this is typically not the case. The problem can be put very succinctly as follows. In standard game-theoretic models of communication, vague signal use yields a lower expected utility than crisp use. Therefore, given that the dynamics (be it natural selection, cultural evolution, or rational choice) maximize utility, vagueness should be weeded out by these forces, giving rise to only precise languages. Thus it would be irrational to stick to vague language use when one could (theoretically) switch to a better system. But vagueness is pervasive in natural language, and there is no reason to believe it is going away. The problem seems to be a theoretically serious one, but it does not obviously undermine our everyday linguistic practices. Most of the time we seem to communicate just fine. Therefore, the issue must lie with the conceptions we have of the forces or mechanisms that underlie those practices. Lipman concludes that “we cannot explain the prevalence of vague terms in natural language without a model of bounded rationality which is significantly different from anything in the existing literature” (2009, p. 1).

Here we survey proposals for addressing vagueness in the context of signaling games, but before we start we want to establish some vocabulary to better frame the upcoming discussion. Rationality is an elusive notion that is frequently debated in areas like philosophy, psychology, and economics. Discussions around it usually touch on different aspects of the concept without always clearly demarcating them; this is what we want to do before going further. The aforementioned logical picture of language and meaning is focused on dichotomies like true and false, meaningful and meaningless, correct and incorrect. A sentence can be true or false only if it is meaningful, and it is meaningful if it is constructed according to correct rules. Rationality is intimately connected with the ability to follow certain procedures, not only of sentence production, but ultimately of sentence combination and reasoning (think of what is required for making a logically valid deductive inference). This is an example of what we will call a *procedural* account of rationality, one which focuses more on the means rather than the ends. One can also do the opposite and focus on the consequences instead. *Instrumental* rationality, as it is typically called, characterizes an agent’s choice as rational if it maximizes the possibility of achieving a desired goal, regardless of the means. The notion is linked to David Hume (1738/2005), epitomized in the following assertion: “Reason is, and ought to only be the slave of the passions.” Instrumental rationality is close to a notion of rational choice that is used in economics and game theory:¹ agents are rational if and only if they make decisions that maximize their expected utility. A more in-depth discussion of the opposition in the context of theoretical economics can be found, for example, in the work of Herbert A. Simon (*e.g.* 1986).²

When developing models where rationality is relevant, be it constructing a logical system or setting up a formal game, the assumed epistemic relation between agents and their environment can come to bear on considerations of rationality. The verdict over how rational certain choices are can vary depending on how accurate and complete an agent’s knowledge is of its environment, the goals to be achieved, the choices or rules available, the relation between those choices and the objectives, and so forth. In this respect, we will call an agent *omniscient* if it is in possession of the same information as the modeler, whereas of an agent with less than that we will say that it only has *limited* awareness of the relevant aspects of the model. Models working within the logical picture typically do not make a distinction between modeler and agent, and thus lack room to express these epistemic gaps. By abstracting away from language users, these models also typically do not represent potential interactions between them, let alone allow for repeated interaction and language change. In other types of models, however, a further aspect of rational choice needs to be considered, namely the ability of agents to make accurate predictions about how other agents behave. We can say that an agent is more or less *strategic* depending on the extent to which she is able to anticipate the actions and beliefs of other agents, and to predict medium/long term gains from repeated play. Lack of perfect knowledge of the situation or lack

¹In fact, it has been argued (Vanderschraaf, 1998) that Hume’s whole account of convention is very closely in line with modern game theory.

²Simon uses the term substantive instead of instrumental rationality, but the characterization is basically the same (Simon, 1986, pp. 210-212).

of ability to choose strategically can be caused by many possible factors. These include, among other things, limitations in handling information (receiving, storing, retrieving, transmitting) and limited computational resources to solve complex problems. We can talk about *bounded* rationality to characterize the choices of such agents. While the above definition of instrumental rationality is usually understood as requiring a single choice to maximize the expected utility in a single concrete decision situation, one might also be interested in more general choice mechanisms (*e.g.* Fawcett, Hamblin, and Giraldeaub, 2013; Galeazzi and Franke, 2017; Hagen et al., 2012; Zollman and Smead, 2010). A choice mechanism is a general way of behaving for an agent involved in a variety of decision situations. When considering only a single situation, rationality can only be *local*. If we take into account the possibility of the agent’s choice or choice mechanism hinging on multiple situations, we can also talk about *global* rationality. This can be important because, hypothetically, there could be choice mechanisms that are sub-optimal at a local level (for each situation) but are actually perfectly rational at a global level.

Note that we consider that, in theory, most combinations of these aspects are possible. Although we have been pinning procedural rationality to the logical picture of language, this is only with the most traditional logical systems in mind. We are not denying that advances in dynamic, epistemic, fuzzy, paraconsistent, and other types of logic could potentially enable one to capture procedural rationality with different characteristics. Game-theoretical models of language can, on the other hand, also combine local instrumental rationality with omniscient highly strategic agents. Our objective here is not to survey all the possibilities. We want to focus on signaling games as a framework for the study of language use and meaning. The vocabulary just introduced will, we hope, help inform the discussion that follows. We proceed by introducing the framework of signaling games in section 2. In section 3 we look into explanations of vagueness in a particular kind of signaling game. Section 4 tries to generalize the considerations of these proposals to argue for an approach to vagueness anchored in a more global notion of rationality. We propose and analyse the results of a novel multi-population model of imprecise imitation in section 5, and summarize our conclusions in section 6.

2 Signaling games

Signaling games were first introduced as models of communication by David Lewis (1969/2011). In order to support the idea that linguistic conventions can arise without any prior conventional activity, Lewis considers situations where agents’ choices involve sending and receiving signals or messages.³ One can think of two players with different roles. The first player, the sender, has knowledge about which of a number of possible states of affairs obtains and, depending on this information, chooses a signal to send. The second player, the receiver, has knowledge about which signal the sender chose and, based on this information, chooses one of several possible responses. A preference relation exists between responses and states of affairs, and a payoff is attributed to each player based on the choices of both. Note that Lewis assumes that no player has any preference regarding the particular signal that is used for a given state, provided that it enables advantageous coordination with responses. Formally, in order to describe the setup all we need is to specify a set of possible states of affairs T , a probability measure P such that $P(t)$ is the probability or frequency with which $t \in T$ occurs, a set of available signals or messages M , a set of responses or actions A , and a pair of utility functions $U_{S,R} : T \times A \rightarrow \mathbb{R}$, one for the sender and one for the receiver, each of which yields a payoff value for each possible pairing of state and action. These so-called signaling problems can be seen as particular cases of coordination problems if we consider the players’ choices to be of contingency plans or strategies. A sender strategy is a specification of a choice of message for each possible state of affairs. It thus describes the sender’s behavior conditional on the state of affairs that obtains. A receiver strategy analogously specifies a choice of action for each possible message. Formally, what the sender chooses is a function $\sigma : T \rightarrow M$ and the receiver a function $\rho : M \rightarrow A$. The expected utility EU of a strategy can be calculated by using the utility function and aggregating payoffs for all pairings of states of

³These terms will be used interchangeably.

affairs and actions, weighted by the probability of each state. Concretely, the expected utility of σ given ρ is $\text{EU}_S(\sigma \mid \rho) = \sum_{t \in T} P(t) U_S(t, \rho(\sigma(t)))$, and the expected utility of ρ given σ is $\text{EU}_R(\rho \mid \sigma) = \sum_{t \in T} P(t) U_R(t, \rho(\sigma(t)))$. As an example, consider a game with $T = \{t_1, t_2\}$, $M = \{m_1, m_2\}$, $A = \{a_1, a_2\}$, $P(t_1) = P(t_2) = 0.5$ and the following utility matrix:

	a_1	a_2
t_1	1, 1	0, 0
t_2	0, 0	1, 1

Consider sender strategy $\sigma = \{t_1 \mapsto m_2, t_2 \mapsto m_1\}$ and receiver strategy $\rho = \{m_1 \mapsto a_2, m_2 \mapsto a_1\}$. These would have an expected utility of 1 for both sender and receiver, since when t_1 obtains with probability 0.5 the sender will use m_2 , and to this message the receiver will respond with a_1 , which achieves a payoff of 1, and when t_2 obtains with probability 0.5 the sender will use m_1 , and to this message the receiver will respond with a_2 , which also achieves a payoff of 1. They also represent one of the two stable conventions in this game, the other being the pair of strategies $\sigma = \{t_1 \mapsto m_1, t_2 \mapsto m_2\}$ and $\rho = \{m_1 \mapsto a_1, m_2 \mapsto a_2\}$. Conventions of this kind in a signaling problem are what Lewis calls *signaling systems*. An example of complete miscoordination would be $\sigma = \{t_1 \mapsto m_1, t_2 \mapsto m_2\}$ and $\rho = \{m_1 \mapsto a_2, m_2 \mapsto a_1\}$. Partial coordination is achieved, for example, by $\sigma = \{t_1 \mapsto m_1, t_2 \mapsto m_1\}$ and $\rho = \{m_1 \mapsto a_1, m_2 \mapsto a_2\}$.

Lewis' account of the stability of conventions rests on what could be considered strong demands. Namely, there needs to be a state of affairs that indicates to everyone involved that a certain regularity will hold, as well as "mutual ascription of some common inductive standards and background information, strategic rationality, mutual ascription of strategic rationality, and so on" (1969/2011, pp. 56–57). Agents are thus envisioned as omniscient and highly strategic. These requirements can seem excessive, and even more so if we consider how simple signaling systems are when compared to human languages. The models were introduced to help explain how language could have gotten off the ground as a conventional system without any sort of prior agreement. However, if one considers the circumstances of the origins of human language, it seems implausible that the agents that started making use of primordial signaling systems which (hypothetically) evolved into languages possessed such advanced rationality. Furthermore, communication through simple message exchange is something that almost all animals do: monkeys use calls, birds use singing, bees use dances, ants use pheromone trails, and so on. A plausible account of the origin of language should first explain how signaling systems could get started, without requiring high standards of rationality from the agents involved.

In order to address this problem, Brian Skyrms (1996) proposes studying signaling problems in evolutionary terms. Rather than imagining, as Lewis does, rational agents making conscious decisions in possession of knowledge of the game and expectations of the behavior of other agents, one can imagine a simpler scenario inspired by biological evolution: there is a population of agents with biologically hardwired behaviors for engaging in interactions characteristic of a signaling problem; utility does not represent preference, but rather fitness for survival and reproduction; the make-up of the population evolves based on the relative fitness of the strategies represented in the population. Such a setup attempts to capture the main features of natural selection: in a diverse population, agents with more successful strategies thrive, while agents with less fit strategies die off. Although the inspiration for this scenario is biological evolution, similar things could be said about how ideas spread in a population of agents who can adopt or abandon them depending on how successful they prove to be (*e.g.* Benz, Jäger, and van Rooij, 2006; Pagel, 2009; Thompson, Kirby, and K. Smith, 2016), *i.e.* we can interpret these notions in terms of cultural evolution (Boyd and Richerson, 1985; Dawkins, 1978). The principles can be captured in formal models that abstract away from details of single interactions and behavior of individual agents, for example in the replicator dynamics (Taylor and Jonker, 1978). The only things relevant to this equation are the relative proportions of strategies in a given population and the utility function. Using it, one can compute which strategies evolve under which conditions.

Skyrms' evolutionary game theory approach to signaling games not only gives more plausible grounds to support Lewis' discussion of convention, it also accomplishes an important conceptual change: it moves most of the theory and mathematical formalism to the descriptive side of the

investigation. Utility represents how the modeler views the signaling problem and understands the relative advantages or disadvantages of different possible strategy combinations. Dynamics describe how strategies can evolve when driven by mechanisms of utility maximization. The shift in perspective allows interpretations that accommodate limited non-strategic agents. While the general framework manages to abstract quite some details away from the formalization, it nevertheless leaves room for them, especially when it comes to the dynamics. We have already mentioned the replicator equation that can be seen as representing biological or cultural evolution, but one can also use dynamics inspired by learning mechanisms (*e.g.* Roth and Erev, 1995), or even ones assuming a high degree of knowledge of the game and other players (*e.g.* Gilboa and Matsui, 1991; Mühlenbernd, 2011; Spike et al., 2017). This range of options goes hand in hand with a range of pictures of rationality, from nothing more than survival of the fittest in a biologically-inspired setting, to a certain degree of instrumental but limited and non-strategic rationality in the case of learning dynamics, to higher levels of rationality and even recursive strategic reasoning about the co-players’ beliefs and choices. Each of these can be utilized depending on the problem that one is interested in characterizing. Thus, although Skyrms shows that high requirements of rationality are not necessary for signaling conventions to evolve, the framework does leave room for the study of linguistic interactions between highly strategic agents.

The characterization of signaling problems in terms of evolutionary game theory allows us to explain why certain equilibria come to be and how. A core notion in this context is that of an *evolutionary stable state* (Maynard Smith, 1982): an equilibrium situation that a population tends to under standard evolutionary pressures, and to which it returns if slightly disturbed. With these tools, one can better understand why signaling systems are stable even without any strong assumptions of rationality. One can also map out which initial conditions drive the system towards which equilibria and which do not. In a simple case like the example discussed above, an evolutionary process of the kind described always drives the population into a state where one signaling system takes over completely. More complex signaling problems may have different evolutionary outcomes, sometimes unexpected ones. Skyrms (2010) gives an overview of different topics studied using signaling games, including expansions of the framework itself (for example, considering other dynamics beyond the replicator equation), exploration of other factors that impact the evolution of signaling (for example, how agents are interconnected), or variations on the signaling problem and its basic assumptions (for example, loosening the alignment of interests in order to provide accounts of deceptive signal use). Other uses of signaling games include discussions of categorization (*e.g.* Jäger and van Rooij, 2007), compositionality (*e.g.* Barrett, 2009), incommensurability (*e.g.* Barrett, 2010), to name a few. More recent overviews are given by Huttegger (2014), Huttegger et al. (2014), and Franke and Wagner (2014). In the following section, we discuss how a particular type of signaling game has been used to address the problem of vagueness.

3 Vagueness in sim-max games

The sorites paradox requires us to assume a relation between the vague terms and a more precise underlying dimension (height for tallness, number of hairs for baldness, number of grains of sand for “heapness”, and so on). Not only does this property need to be much more fine-grained than the vague term, but it also needs to have some structure: there is at least an order between the elements in it (thinking of height in centimeters, $180 > 179 > \dots > 120$), and usually even a degree of how far apart these elements are from each other. In terms of signaling games, one can model this using a state space constituted of values of the underlying dimension, and a message space constituted by the terms in question. Because of the difference in granularity, we will typically be interested in cases where the state space is much larger than the message space. We can model the structure of the state space by defining a distance or similarity function between every value, effectively making it a metric space. Another important ingredient of the sorites paradox is the acknowledgment of a certain degree of tolerance with respect to whether a certain term applies or not. This tolerance decreases with distance in state space: assuming a 180cm person is tall,




Figure 1: Example strategies for a state space with 50 states. Each line corresponds to a message. For the sender, it plots for each state the probability that the message is used. For the receiver, it plots for each state the probability that the response is that state, given the message.

one would easily tolerate the use of the term for a person measuring 179cm, less so for someone who is 170cm, and much less so for 160cm. This can be modeled using a utility function that is continuous rather than discrete and that monotonously decreases with distance, *i.e.* success is not a matter of black and white, right or wrong, but a matter of degree, of how close the receiver got to the optimal response to the sender’s perceived state.

The simplest type of game to study in this scenario is one where the state space and the action space are the same. We can imagine this as a game of guessing states of affairs: the sender has knowledge of a particular state, sends a message to the receiver, who in turn has to guess it; their payoff, as discussed above, is proportional to how close the guess got to the original state. These games, called similarity-maximization or sim-max games for short, were first introduced by Gerhard Jäger and Robert van Rooij (Jäger, 2007; Jäger and van Rooij, 2007) and further studied by Jäger, Metzger, and Riedel (2011). What these authors find about this setup is that the evolutionary stable states are what they call Voronoi languages. Roughly, these are situations where the sender uses messages in a way that can be seen as partitioning the state space into convex regions, and the receiver responds with the central element of those regions.

In an abstract setup, using 50 states uniformly distributed over the unit interval and two possible messages, such an optimal language looks like what we see in Figure 1a: at a specific point, the probability that the sender uses one message decays sharply from 1 to 0, and increases sharply from 0 to 1 for the other message; the response of the receiver for each message is a degenerate distribution over the state space which assigns all probability mass to a single state. These strategies give mutually optimal behavior for a case where the prior probability is the same for each state and utility is a linear or quadratic function of the distance between the actual state and the receiver’s guess. In general, at which point the sender switches the use of messages and which guesses of the receiver are optimal critically hinges on the priors over states and the utility function. Still, confirming Lipman’s argument, there is no vagueness in such optimal languages.

The sender/receiver strategy pairs that we are looking for look more like Figure 1b, where the sender’s probability of choosing a certain message gradually changes, and the receiver strategy assigns a positive probability to more than one state for each message. What characterizes vagueness in these models is thus a smooth and monotonous transition between parts of the state space where a message is clearly used and where it is clearly not. This means that, for some states in the middle of the state space, there is uncertainty as to which message will be used by the sender, whereas for states in the extremes this is as good as certain. For the receiver part, there is uncertainty as to which state will be picked for each message.

The interpretation of this uncertainty will be different depending on the interpretation of the model. If we see it as an explanatory model of how two agents play the game, we can see it as randomization. If we interpret the model as descriptive, it simply represents expected behavior in a manner agnostic to the underlying mechanisms. A third option is to see probabilities as

capturing relative numbers in a population of agents. For example, if the sender strategy assigns a probability of 0.4 to the event of message m being sent for state t , this would mean that 40% of the population uses that message for that state. This latter option leaves open the possible interpretation that each agent commands a crisp language and vagueness is only a population-level effect. However, given the level of abstraction of the description so far, none of this is necessarily implied by the model. Our preferred stance is to see the model as descriptive, but we do not wish to focus on that debate here. The question we want to address is: what additional modifications to sim-max games are sufficient for optimal languages to be more like Figure 1b, rather than like Figure 1a?

Franke, Jäger, and van Rooij (2011) make two suggestions as to how vagueness in a signaling game can be boundedly rational, *i.e.* how vagueness could arise as a consequence of cost-saving limitations in the cognitive capacities of instrumentally rational agents. The first proposal is called limited memory fictitious play (LMF) and models agents playing a sim-max game where their ability to recall past interactions with others is limited to a certain number. For a given interaction, each agent uses her limited memory of the other agent’s past behavior to estimate the other’s strategy, and plays an instrumentally rational best response to that strategy. In order to study the evolution of strategies in repeated interaction, the authors model several individual agents in actual play. What they observe is the emergence of vague signaling at the level of the population, *i.e.* population averages of individual strategies exhibit the hallmarks of a vague language. This is because each agent recalls a different history of previous play and so holds slightly different beliefs about language use. However, each agent still commands a crisp language, which is an inadequate feature of the model if the intention is to capture how vagueness presents itself in human languages.

In order to overcome this limitation, Franke, Jäger, and van Rooij make another proposal using the notion of a quantal response equilibrium (QRE). The idea, inspired by experimental psychology, is to model the choice of best response as stochastic rather than deterministic.⁴ A prominent explanation for such soft-max or quantal choice behavior is that agents make small mistakes in the calculation of their expected utilities (Train, 2009). They still choose the option with the highest expected utility, but each assessment of the expected utilities is noise perturbed. This, in turn, may actually be boundedly rational since the calculation of expected utilities relies on assessing stochastic uncertainty, which in turn may be costly to calculate precisely. Choice based on a few samples from a distribution can be optimal if taking more samples or other means of better approximating probabilistic beliefs is resource costly (*e.g.* Sanborn and Chater, 2016; Vul et al., 2014). The degree to which agents tremble in the calculation of expected utilities and therefore deviate from the instrumentally rational behavior can be characterized by a parameter. Franke, Jäger, and van Rooij find that for low values of this parameter, only babbling equilibria are possible, where sender and receiver simply randomize, respectively, message and interpretation choice uniformly. Above a certain value of the parameter, other equilibria of the kind described in the beginning of this section arise, where agents communicate successfully, though not perfectly, using fuzzy strategies similar to those depicted in Figure 1b. However, it is not clear whether soft-max choices capture the right stochastic trembles in decision making as they would arise under natural sources of uncertainty about the context (see Franke and Correia, 2017).

Cailin O’Connor (2014) proposes a way in which vagueness could be expected to evolve as a side-effect of a particular type of learning process. She studies sim-max games driven not by rational choice dynamics, but by generalized reinforcement learning (GRL), a variant of Herrnstein reinforcement learning (HRL) (Roth and Erev, 1995). In HRL, agents learn to play a signaling game by strengthening particular choices (of messages for the sender, of responses for the receiver) proportionally to how successful those choices prove to be in an interaction. O’Connor’s proposal is to model generalization as the propagation of reinforcement to nearby states, where “nearby” is defined in terms of distance in state space. For example, if a sender was successful in using message m for state t , she will not only positively reinforce that choice of message for t , but also

⁴Probabilistic choice rules are also the source of vagueness in recent accounts by Lassiter and Goodman (2017) and Qing and Franke (2014).

for states similar to t . This is done to a degree that is proportional to the similarity between t and other states. The dynamics gives rise to vague signaling of the kind we are looking for.

Although there is a close relationship between reinforcement learning and population-level dynamics (Beggs, 2005; Börgers and Sarin, 1997), O’Connor’s GRL is, on the face of it, an account of learning between individual players. Also, we need further justification for linking generalization in reinforcement so closely to the underlying payoff function of the sim-max game. Why should agents evolve to generalize in exactly the right way? O’Connor suggests that, despite a language with vagueness having lower expected utility than a language without it, the learning mechanism that induces vagueness does have evolutionary advantages: it achieves higher payoffs in a shorter period of time. From a global point of view, learning speed can be an advantage. Imagine an initial population of agents with random strategies, some using GRL and others using classical HRL. Although the latter type of agent can hypothetically develop a precise and more efficient signaling system, agents using GRL could coordinate on vague signaling strategies with high (though not optimal) expected utility sooner than agents using HRL. In such a scenario, they could drive the other agents to extinction before the latter had time to achieve coordination and reap the benefits of a more precise signaling system.

A similar finding is made by Franke and Correia (2017) when studying a variant of the replicator dynamics in which individual agents do not have the ability to generalize but simply make perceptual mistakes. In a scenario where agents learn by imitation, if one assumes that they do not have perfect perception, there will always be the possibility for senders to confuse states and thus learn associations with messages that are different than the ones observed, and for receivers to similarly mix up responses. Furthermore, it seems reasonable to assume that this confusability is proportional to state similarity, *i.e.* that the more similar two states are, the more likely it is that they will be mistaken for each other. Incorporating these considerations into a derivation of the replicator dynamics based on imitation processes, Franke and Correia develop a variant of the dynamic that also induces vague signal use of the kind we expect here. The consequence is very similar to that of the GRL model discussed above, in that the way the behavior for a given state is updated takes into account the behavior of similar states, proportionally to their similarity. Given the known relation between reinforcement learning and the replicator dynamics (Beggs, 2005), it is actually quite plausible that the two are tightly related (although this would need to be formally demonstrated). The account is, furthermore, interpretable at a lower level of rationality.

The motivation underlying this model of vague signaling is still one of inevitability. A vague strategy is not claimed to have higher expected utility than a crisp one. However, the authors observe an effect similar to that pointed out by O’Connor: signaling converges faster and more often in scenarios where there is some degree of state confusability. Furthermore, they observe one additional potentially beneficial property. Running several rounds of simulation for each parameter set, they measure for each group of results how close resulting strategies are to each other, and how they would fare playing against one another. The results show that the within group distance between strategies becomes smaller with growing confusability, and the within group expected utility is actually higher for strategies evolving under a certain degree of state confusion. Thus, some amount of uncertainty seems to promote more homogeneous populations of signalers that are better at achieving cooperation within a group. What is left to show, as it was with O’Connor’s GRL approach, is whether the potential payoff advantages that were observed in simulations actually suffice to promote vague language use in an encompassing model of multi-level selection. The following section motivates the need for taking a more ecological approach to the evolution of vagueness, before section 5 gives a concrete model.

4 The ecology of vagueness

Arguments of the kind presented by Lipman (2009), that vague signal use is suboptimal when compared to crisp use, work under various assumptions. Part of the picture formed by these assumptions is a highly idealized conception of the agents involved and of the context in which they develop and use signals. These idealizations probably originate, via game theory, from the

conception of rationality of traditional theoretical economics. Herbert A. Simon describes this picture as follows:

Traditional economic theory postulates an “economic man,” who, in the course of being “economic” is also “rational.” This man is assumed to have knowledge of the relevant aspects of his environment which, if not absolutely complete, is at least impressively clear and voluminous. He is assumed also to have a well-organized and stable system of preferences, and a skill in computation that enables him to calculate, for the alternative courses of action that are available to him, which of these will permit him to reach the highest attainable point on his preference scale. (Simon, 1955, p. 99)

Both Simon and Lipman call for this picture to be revised, and this is what the proposals surveyed here all do. In order to account for vagueness in natural language in the context of these models, they peel away from this idealized picture and bring some of these assumptions down to earth. In the process, they point to ways in which we, as language learners and language users, are finite beings finding ways to cope with a highly complex and dynamic environment:

1. Our existence is temporally finite; language does not have an infinite amount of time to evolve, nor can it take an infinite time to be learned. The faster a language can start being useful, the better;
2. Language learning through experience has to rely on a limited number of observations. Not only is the state space typically much larger than one can survey in sufficient time, it is even potentially infinite and constantly changing;
3. A corollary of the former is that there will always be heterogeneity in a population of language learners, at the very least in their prior experience, since each agent will have relied on a different set of observations. Furthermore, this information is not directly or fully accessible to others.

All of these observations support the weakening of modeling assumptions. The research surveyed here shows us some examples of assumptions which, when weakened, make vague signal use a natural outcome of certain evolutionary dynamics. But it gives us even more. It suggests ways in which the mechanisms that lead to vagueness can have positive effects that are extremely important in the context of the points just enumerated. We learned from O’Connor and Franke and Correia that vague languages are quicker to converge and adapt, which is valuable given the finite and dynamic character of our experience (point 1). O’Connor also showed how generalization, an invaluable feature of any procedure for learning from a limited number of observations (point 2), leads to vagueness. We also learned from Franke and Correia that state confusability, a mechanism that leads to vague signal use, can have a homogenizing effect on vocabularies, potentially compensating for the heterogeneity of agents’ experiences (point 3).

What do these observations tell us about rationality? GRL (O’Connor, 2014) and the work of Franke and Correia (2017) both assume a picture of agents with a basic level of instrumental rationality, possibly limited awareness of the game, and a lack of strategic capabilities, adapting their behavior with only short-term gains in sight. These approaches introduce constraints on agent behavior or information processing that prevent the evolution of crisp signal use. But a crisp language would still have a higher expected utility than the evolved strategies. Agents in those models seem to be only as rational as the modelers allow them to be. Despite the plausibility of the mechanisms proposed (limited memory, imprecise calculation of expected utilities, generalization, state confusability), the results of these models feel somewhat bittersweet because of the hypothetical possibility of an ideal strategy, seemingly barred from the agents in an artificial manner. Couldn’t a more rational agent evolve and drive the system into crisp signal use? Aren’t we, human beings, that kind of agent?

Perhaps a deeper understanding of vagueness and the reasons for its pervasiveness in natural language are to be found only when we broaden the scope of the models employed. All the models discussed so far explore evolutionary dynamics for one homogeneous population playing

one game. Different types of agents and different game setups are considered, but each of these different possible scenarios is tested separately. We see at least two ways in which one could embrace a more ecological perspective. The first is to think about meaning and vagueness from a more Wittgensteinian perspective. One can see each signaling game as embodying a particular language game. In Wittgenstein’s picture of language, however, we do not play only one language game over the course of our existence; there is a plurality of them and which one an agent is engaged in at a particular moment is never clearly identified, neither are the exact benefits one might gain by choosing a certain behavior over another. These are furthermore not fixed in time; old language games fall out of fashion or stop being useful, and new ones emerge all the time (see Wittgenstein, 1953/2009, and in particular §23).

We can look for rationality at several levels in this pluralistic picture. First, as before, there is the actual behavior of a single agent in each actualized language game. As mentioned above in connection the soft-max choice function used by Franke, Jäger, and van Rooij (2011), behavior that strictly maximizes expected utility under uncertainty may be resource heavy, so it might be compatible with local strategic rationality that agents’ production choices are stochastic. Second, if we look at behavior across many game types and contexts, there is also the level of an agent’s internal theory of how words and phrases are likely to be used (or even normatively: how messages should be used), conditional on a given context. Notice that a single agent’s rational beliefs about linguistic practices or linguistic meaning may well have to reflect the actual stochasticity: under natural assumptions about information loss, the best belief for prediction matches the actual distribution in the real world (*e.g.* Vehtari and Ojanen, 2012). In sum, both at the level of behavior and at the level of beliefs about use or meaning, we should expect to find vagueness. Still, despite the natural vagueness, there does not seem to be anything fundamentally missing or conceptually incoherent in a naturalistic, rationality- or optimality-driven explanation of what each agent is doing or what each agent believes about language, use and meaning.

Another way to go beyond locality is to work with more heterogeneous population models. The mechanisms that lead to vague signal use, as O’Connor (2014) and Franke and Correia (2017) stress, have the aforementioned important advantages of faster speed of convergence, higher flexibility, and homogenization. The argument goes that these side-effects, by temporarily enabling a higher expected utility, could allow a population using some generalization (or affected by some imprecision) to take over. However, despite its intuitiveness, the argument is based on comparing isolated runs of different dynamics. The models do not allow the hypothesis to be tested, because they do not accommodate different populations evolving together. In the remainder of the paper we propose a way to do this based on the model of Franke and Correia (2017). We introduce two variants of a multi-population model of the imprecise imitation dynamics, where populations characterized by different imprecision values interact and evolve together. Using this model, we can better see under which conditions the hypothesized advantages of some imprecision can lead to the evolution of vagueness.

5 Two multi-population models of imprecise imitation

We build upon the model of Franke and Correia (2017), adding to the imprecise imitation dynamics support for multiple populations with different imprecision values evolving together. The resulting evolutionary dynamic has two layers: first, signaling strategies change by imprecise imitation following the model of Franke and Correia (2017); second, there is evolutionary selection at the level of the imprecision values themselves. More concretely, imagine agents who are born with varying perceptual abilities, subsequently learn a signaling strategy by imitation of other agents (either within or across populations), and depending on the success of the strategies they develop, are more or less likely to survive and reproduce. In such a setup, evolutionary processes will not only change the distribution of strategies within a population, where a population is identified by a shared level of imprecision in imitation, they will also promote those levels of imprecision which result in the development of more successful signal use.

This is closely related to the ideas behind theories of kin selection (Hamilton, 1964) and

multi-level (or group) selection (Wilson, 1975). These theories build upon the hypothesis that selection acts not only to directly favor genes that result in behaviors that benefit individuals, but also to indirectly favor genes that lead to behaviors that benefit either genotypically (kin) or socially (group) related individuals (see Okasha, 2006, for more details). In our model there is a similar structure: there is a process of selection of behavior within each population, and levels of imprecision are selected across populations based on the strategies they give rise to. We believe there is an important difference with group and kin selection models in that the two selection processes in our model act on different entities: the inner shaping signaling behavior, and the outer selecting levels of imprecision. In any case, we intend our model to be descriptive rather than causal or explanatory. That means that we do not want to commit to seeing populations either as kin or as social groups, and use them as merely descriptive abstractions that allow us to capture the hypothetical impact of indirect selection processes.⁵

In the following, we lay down the formal details of our model. Let's start by defining A to be the set of imprecision values considered. For each value $\alpha \in A$, the proportion of its population is given by $P(\alpha)$, such that $\sum_{\alpha \in A} P(\alpha) = 1$. Each population has its own sender and receiver strategies, represented as σ^α and ρ^α . The probability that a given agent with imprecision α (or of type α) observes t_o when the actual state is t_a is given by $P_o^\alpha(t_o|t_a)$. If the same agent intends to realize t_i , the probability that she actually realizes t_r instead is given by $P_r^\alpha(t_r|t_i)$. Following Franke and Correia (2017, p. 26), we can then define the following values.

Probability that t_a is actual if t_o is observed by an agent of type α :

$$P_o^\alpha(t_a|t_o) \propto P_a(t_a)P_o^\alpha(t_o|t_a)$$

Probability that a random sender of type α produces m when the actual state is t_a :

$$P_\sigma^\alpha(m|t_a) = \sum_{t_o} P_o^\alpha(t_o|t_a)\sigma^\alpha(m|t_o)$$

Probability that the actual state is t_a if a random sender of type α produced m :

$$P_\sigma^\alpha(t_a|m) \propto P_a(t_a)P_\sigma^\alpha(m|t_a)$$

Probability that t_r is realized by a random receiver of type α in response to message m :

$$P_\rho^\alpha(t_r|m) = \sum_{t_i} P_r^\alpha(t_r|t_i)\rho^\alpha(t_i|m)$$

These formulations are merely parameterized versions of the single-population model. They encapsulate calculations that one can use to compute expected utilities and strategy update steps for each type. The latter, however, depend on the types of interaction that we imagine occurring between populations. In the following sections, we consider two different possibilities.

5.1 Tight population interaction

In a multi-population model with tight interaction between populations, each agent plays with, observes, and potentially imitates any other agent, regardless of their type. This has an impact on the expected utilities of sender and receiver strategies of each type, and on the update steps for those strategies. Let's start with the expected utilities. For a sender of type α , the expected utility of its strategy σ^α against all other receiver strategies ρ^* , is given by:

$$EU_\sigma^\alpha(m, t_o, \rho^*) = \sum_{t_a} P_o^\alpha(t_a|t_o) \sum_{\alpha' \in A} P(\alpha') \sum_{t_r} P_\rho^{\alpha'}(t_r|m)U(t_a, t_r)$$

⁵We make this note because of the heated debate between the two theories. See Kohn (2008), and Kramer and Meunier (2016) for more details.

and, for a receiver of type α , the expected utility of its strategy ρ^α against all other sender strategies σ^* , is given by:

$$\text{EU}_\rho^\alpha(t_i, m, \sigma^*) = \sum_{\alpha' \in A} P(\alpha') \sum_{t_a} P_{\tilde{\sigma}}^{\alpha'}(t_a | m) \sum_{t_r} P_r^\alpha(t_r | t_i) U(t_a, t_r)$$

Expected utilities thus take into account the existence of strategies of other types, and weigh the relevance of each type α' according to its relative proportion $P(\alpha')$.

Another important value to calculate has to do with the types that agents observe and imitate. In a model with tight interaction, we imagine this occurring across populations. Therefore, we can define the probability that a sender of type α observes a randomly sampled agent play message m for observed state t_o as:

$$P_o^\alpha(m | t_o) = \sum_{t_a} P_o^\alpha(t_a | t_o) \sum_{\alpha' \in A} P(\alpha') P_{\sigma'}^{\alpha'}(m | t_a)$$

and the probability that a receiver of type α observes a randomly sampled agent choose interpretation t_o given message m as:

$$P_o^\alpha(t_o | m) = \sum_{t_r} P_o^\alpha(t_o | t_r) \sum_{\alpha' \in A} P(\alpha') P_{\rho'}^{\alpha'}(t_r | m)$$

Again, these calculations incorporate the probabilities that the imitating agent might observe the behavior of an agent of another type α' , weighed by its relative proportion. Finally, the update step for a sender strategy of type α at time instant $i + 1$ is given by:

$$\check{\sigma}_{i+1}^\alpha(m | t) \propto P_o^\alpha(m | t) \text{EU}_{\sigma_i}^\alpha(m, t, \rho_i^*)$$

and similarly for a receiver strategy of type α by:

$$\check{\rho}_{i+1}^\alpha(t | m) \propto P_o^\alpha(t | m) \text{EU}_{\rho_i}^\alpha(t, m, \sigma_i^*)$$

We here use $\check{\sigma}$ and $\check{\rho}$ since there is still an additional adjustment to these values to be calculated before we get the final strategies σ and ρ .

These formulations cover the evolution of the particular strategies of each type. We can think of this as the level of cultural evolution: agents are born with a certain level of imprecision and adopt strategies based on the behavior of others. Alongside this process, we can imagine another level of selection, where agents die and new agents are born. More successful agents have a higher likelihood of surviving and reproducing, giving rise to more agents with their level of imprecision. Levels of imprecision are thus subject to an evolutionary dynamic that is indirectly influenced by the cultural dynamic. Importantly, only the level of imprecision is passed on to new generations under this dynamic, not the actual strategies developed by the agents at the cultural level.

We model this process by changing the proportion of each type $P(\alpha)$ according to the replicator dynamic. A population of type α consists of agents employing both a sender and a receiver strategy, thus the overall fitness of the population must include the expected utilities of both. We could imagine this process happening at a different speed to the cultural process, in which case we could have different time scales. For the sake of simplicity, we choose to have them both happen at each time step, but calculate the changes occurring at this level of selection after the calculation for the other level. We define the proportion of type α at time step $i + 1$ as:

$$P_{i+1}(\alpha) \propto P_i(\alpha) (\text{EU}_{\sigma_{i+1}}^\alpha(m, t, \rho_i^*) + \text{EU}_{\rho_{i+1}}^\alpha(t, m, \sigma_i^*))$$

In order to additionally account for the fact that strategies are not passed on to new generations, we mix the evolved strategies of a certain type σ^α and ρ^α with new random strategies $\tilde{\sigma}^\alpha$ and $\tilde{\rho}^\alpha$ (generated at each time step). The idea is that, at each time step, a certain percentage of each population will consist of “newborn” agents, *i.e.* agents that haven’t yet had time to evolve their strategies. We define a parameter γ that quantifies this percentage, or as we can also call it the

birth rate, which we consider to be the same for every population. This mixing finally defines the strategies for time step $i + 1$ and can be described in the following formulas:

$$\begin{aligned}\sigma_{i+1}^\alpha(m|t) &= (1 - \gamma)\check{\sigma}_{i+1}^\alpha(m|t) + \gamma\tilde{\sigma}_{i+1}^\alpha(m|t) \\ \rho_{i+1}^\alpha(t|m) &= (1 - \gamma)\check{\rho}_{i+1}^\alpha(t|m) + \gamma\tilde{\rho}_{i+1}^\alpha(t|m)\end{aligned}$$


Simulation results We performed 25 simulation runs of this model for each of three population scenarios: only one population with $\alpha = 0$ (for reference), two populations with $\alpha \in \{0, 0.05\}$, and three populations with $\alpha \in \{0, 0.05, 0.1\}$. For each scenario, starting proportions $P(\alpha)$ were equal for each value of α . Given the observations by Franke and Correia (2017) that state space size and tolerance parameter β do not result in important qualitative difference, we fixed these values at $n_S = 30$ and $\beta = 0.1$. We also used a fixed uniform distribution for the priors and a message space with 2 messages. Each type started with its own randomly generated strategy. Regarding the duration of the simulations, due to the mixing in of new individuals into the population (birth rate was fixed at $\gamma = 0.05$), the convergence criteria is no longer applicable because each strategy is randomly perturbed at each time step. Therefore, all simulation runs were stopped after 200 iterations.

The first thing to observe from the simulation results is that the population with no imprecision ($\alpha = 0$) dominated the other populations in every run. In Figure 2 we plot the evolution of the proportion of each population in the two-population and three-population scenarios for all trials. As the plot shows, the population with no imprecision steadily increased its proportion against the others in every trial. In the three-population scenario, the proportion of type $\alpha = 0.05$ sees a slight increase in the beginning of the simulation, but inexorably starts a downward trend. These observations go against the expectation of Franke and Correia (2017) that faster convergence to a convex strategy by populations with a certain level of imprecision could give them a temporary advantage to take over and eliminate other types. The reason for this is interesting in itself. What happens is that, because of the tight interaction between populations, the strategies of each type evolve in close tandem with each other. One of the consequences of this is that the population with no imprecision reaches convexity faster than it would on its own because of the interaction with the populations with imprecision. We can see this effect by looking into the percentage of trials with convex sender strategies at a given iteration, for each scenario, and comparing the three scenarios: populations with no imprecision evolving alone, two populations ($\alpha \in \{0, 0.05\}$), and three populations ($\alpha \in \{0, 0.05, 0.1\}$). This is plotted in Figure 3a. What we see is more trials reaching convexity earlier for the multi-population scenarios when compared with the single-population scenario. This effect precludes the hypothesized temporary advantage of imprecision manifesting itself, but it can be seen as a positive influence of the population(s) with imprecision on the population with no imprecision.

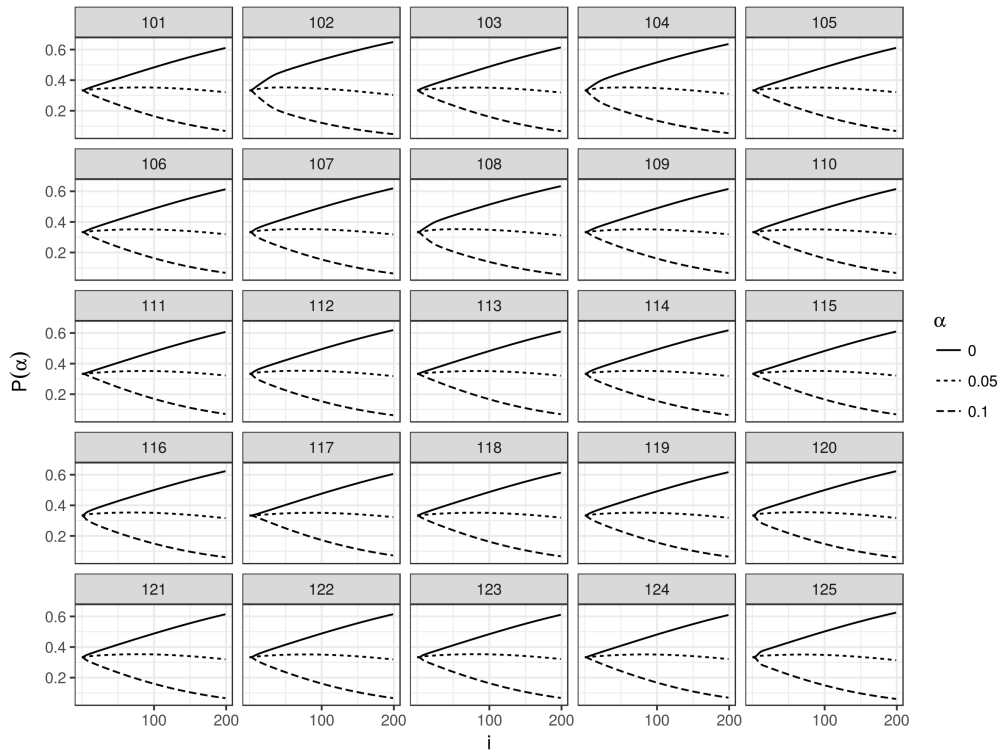
The flip side of this tight connection between populations is that strategies evolved by populations with no imprecision are also more vague (in the sense defined in section 3). This is visible by looking at mean entropy values, namely sender strategy entropy, plotted in Figure 3b, and receiver strategy entropy, plotted in Figure 3c. The values for the population with no imprecision are clearly higher in the scenarios where it evolves together with populations with imprecision than in the scenario where it evolves on its own. Given the trends in population proportions, one expects this to eventually be eliminated when the population with no imprecision finally takes over the others, but it is interesting to observe that while populations with vague strategies persist, the population with no imprecision takes much longer to evolve a fully crisp strategy.

5.2 Loose population interaction

We can also imagine a scenario where populations interact more loosely. Namely, the dynamic we want to model here is one where agents of a certain type α imitate and learn only from other agents of the same type, but nevertheless use their signaling strategies with agents from all types. In order to capture this, we need to make changes to some calculations. Any formula that is not




(a) Two-population scenario.



(b) Three-population scenario.

Figure 2: Evolution of population proportions through time for each simulation trial of the tight interaction model. Numbers on top of each plot identify each trial.



(a) Percentage of convex sender strategies for $\alpha = 0$. (b) Mean entropy of sender strategies for $\alpha = 0$. (c) Mean entropy of receiver strategies for $\alpha = 0$.

Figure 3: Development of some metrics through time for the $\alpha = 0$ population in each of three scenarios: evolving alone ('0'), with an $\alpha = 0.05$ population ('0/0.05'), and with an additional $\alpha = 0.1$ population ('0/0.05/0.1').

redefined in this section should be assumed to remain the same. First, if agents learn only within their population, the imitation dynamic needs to consider only the expected utility against agents of that population. We thus define the following expected utilities for a sender of type α :

$$\text{EU}_\sigma^\alpha(m, t_o, \rho^\alpha) = \sum_{t_a} P_\sigma^\alpha(t_a|t_o) \sum_{t_r} P_\rho^\alpha(t_r|m) U(t_a, t_r)$$

and for a receiver of type α :

$$\text{EU}_\rho^\alpha(t_i, m, \sigma^\alpha) = \sum_{t_a} P_\sigma^\alpha(t_a|m) \sum_{t_r} P_r^\alpha(t_r|t_i) U(t_a, t_r)$$

The main difference from the previous model is that these expected utilities are not calculated against all populations (σ^* and ρ^*) but only against the agent's own type (σ^α and ρ^α). This also implies that population proportions do not play a role at this level of selection.

Regarding the imitation process, we can redefine the probability that a sender of type α observes a randomly sampled agent play message m for observed state t_o as:

$$P_\sigma^\alpha(m|t_o) = \sum_{t_a} P_\sigma^\alpha(t_a|t_o) P_\sigma^\alpha(m|t_a)$$

and the probability that a receiver of type α observes a randomly sampled agent choose interpretation t_o given message m as:

$$P_\rho^\alpha(t_o|m) = \sum_{t_r} P_\rho^\alpha(t_o|t_r) P_\rho^\alpha(t_r|m)$$

Again, the main difference is that agents make observations within their own population, so to model a randomly sampled agent one needs only to take into account agents of the same type. Population proportions again do not play a role in these calculations. Based on these formulas, we can define the update step for a sender strategy of type α at time instant $i + 1$ as:

$$\check{\sigma}_{i+1}^\alpha(m|t) \propto P_\sigma^\alpha(m|t) \text{EU}_{\sigma_i}^\alpha(m, t, \rho_i^\alpha)$$

and similarly for a receiver strategy of type α as:

$$\check{\rho}_{i+1}^\alpha(t|m) \propto P_\rho^\alpha(t|m) \text{EU}_{\rho_i}^\alpha(t, m, \sigma_i^\alpha)$$

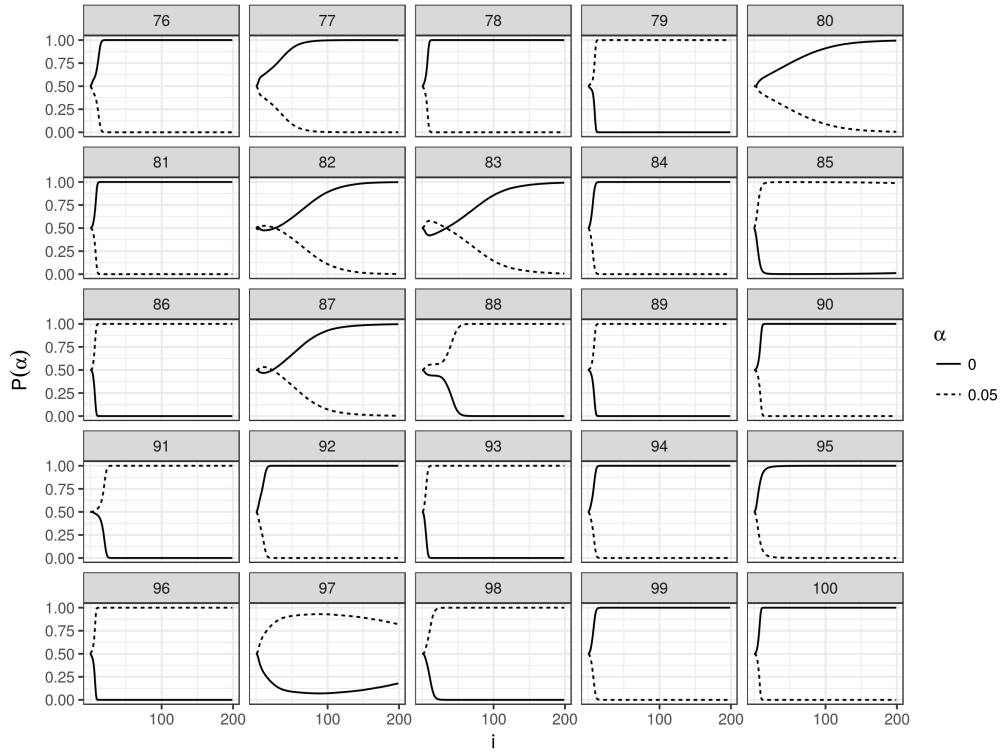
Note that, because imitation and learning occur only within populations, these formulations are essentially the same as for the single-population model of Franke and Correia (2017), only parameterized by type α .

The selection process between different populations still follows the same motivation as before: agents of a certain type that evolve successful strategies (with respect to other types) will be more likely to survive and reproduce, benefiting the proportion of their population. The formulation of the dynamic for the proportion of each type $P(\alpha)$ thus stays the same. In order to avoid confusion, we want to stress that this means that these calculations rely on the definitions of expected utility across populations (*i.e.* $\text{EU}_\sigma^\alpha(m, t_o, \rho^*)$ and $\text{EU}_\rho^\alpha(t_i, m, \sigma^*)$) and not the newly introduced $\text{EU}_\sigma^\alpha(m, t_o, \rho^\alpha)$ and $\text{EU}_\rho^\alpha(t_i, m, \sigma^\alpha)$.


Simulation results We ran the same number of simulation trials under the same conditions as for the model with tight population interaction. In Figure 4 we plot the evolution of population proportions for all trials of the multi-population scenarios. The first thing to observe is that proportions evolve faster than in the model with tight interaction. Whereas in the latter no given population ever reached much more than 70% proportion, in this model we see that many trials resulted in one population fully dominating the others. In the two-population scenario, some population reached at least 99% in 24 out of 25 trials. In the three-population scenario, this happened in 18 out of 25 trials. More interestingly for our investigation, some trials actually resulted in the population with $\alpha = 0$ being dominated. For the two-population scenario, $\alpha = 0.05$ fully reached 100% in 8 trials (79, 86, 88, 89, 91, 93, 96, 98), a point from which it is technically impossible for the other population to recover. In most cases the dominating population gains its ground from the start, but in 5 cases we see a temporary advantage of $\alpha = 0.05$ that is then lost to the other population. In one interesting trial (85), $\alpha = 0.05$ reached 99.87% only to then steadily start losing ground to $\alpha = 0$. For the three-population scenario, $\alpha = 0$ was reduced to 0% in 4 trials (152, 157, 159, and 175). In all of those cases, $\alpha = 0.05$ clearly has the upper hand over $\alpha = 0.1$, despite a temporary advantage of the latter in some trials.

Because of the loose interaction between populations, different types can now evolve separately. One consequence of this is that populations with a certain level of imprecision again reach convexity faster than those without imprecision. In Figure 5 we plot, for each trial, the first iteration when each type reached convexity. What we see is that, even though $\alpha = 0$ usually reaches convexity later, this is not always the case. This has certainly to do with the initial conditions of each trial, since the randomly generated strategies can simply by chance be more favourable to reaching convexity. More importantly, we also see that reaching convexity sooner is not a sufficient condition for achieving population dominance. There were many trials where another type reached convexity much sooner than $\alpha = 0$ but its population was completely dominated nevertheless (*e.g.* 76, 84, 92, and 99 in the two-population scenario; 156, 165, 168, and 172 in the three-population scenario). It also does not seem to be fully necessary, given a few examples where $\alpha = 0$ reached convexity early on and another type ended up dominating (89 in the two-population scenario and 152 in the three-population scenario).

Another consequence of the populations evolving separately in this model is that they do not necessarily evolve towards the same equilibrium. In a sim-max game such as the one set up here, there are only two stable equilibria. These are the two Voronoi languages of the kind shown in Figure 1: one where the first message is used for the first half of the state space, and the other where the second message is used for this region. In the multi-population model with loose interaction, each population evolves independently towards one of these two equilibria. An important factor determining which language strategies converge to is the random initial population configuration. The populations are, however, not fully independent, since the process of selection of the level of precision relies on the expected utility of one population playing against itself and the others. And this is important because strategies in one equilibrium get the lowest payoff possible against strategies in the other equilibrium. Now, if two populations evolve towards different equilibria, whatever advantage one population has playing against itself could trigger a runaway effect by causing an increase in its proportion, which in turn will increase the population's relative expected utility, potentially increasing their proportion further in the next round, and so forth. In this case, one would expect that faster convergence towards convexity would be especially important for a population's success.




(a) Two-population scenario.




(b) Three-population scenario.

Figure 4: Evolution of population proportions through time for each simulation trial of the loose interaction model. Numbers on top of each plot identify each trial.




(a) Two populations scenario.




(b) Three populations scenario.

Figure 5: First iteration with convex sender strategies for each simulation trial of the loose interaction model.



(a) Sender expected utility.



(b) Receiver expected utility.

Figure 6: Expected utilities for trial 88 of the loose interaction model. Vertical lines demarcate, for each type, the iteration where convexity was achieved.

In the two-population scenario, of the 8 trials where $\alpha = 0$ was reduced to 0%, all but two (88, 98) ended with the two populations each close to different equilibria. In the three-population scenario, the population with $\alpha = 0$ evolved towards the same equilibrium as the other two in 1 of the trials where it was eliminated. In the other 3 trials, in 2 it evolved to a different equilibrium than the other two, and in 1 it evolved towards the same equilibrium as $\alpha = 0.05$ (but different to the one $\alpha = 0.1$ evolved towards). Again, we do not find a clear case for this being a decisive factor in the success of populations with some imprecision. The same goes for the impact of an initial advantage in expected utility creating the runaway effect we just mentioned. Despite none of these three factors (including reaching convexity sooner) sentencing the demise of $\alpha = 0$ with certainty, they do seem to conspire together to bring it about. In most trials where the type was eliminated, in both the two- and three-population scenarios, $\alpha = 0$ ended up evolving towards a different equilibrium than the other types, and either had an initial disadvantage in expected utility, or reached convexity much later. These factors can thus be seen as indicators at best, and the story of how a population with imprecision ends up dominating one without it seems to be more complicated than expected.

This is not surprising, since we are facing a complex dynamical system with various interacting components (sender and receiver strategies and multiple populations). Just as an example of this, in Figure 6 we plot the expected utilities of sender and receiver strategies for both populations

of trial 88. The curves up to where $\alpha = 0.05$ achieved convexity seem to suggest that the sender strategy of that type was evolving towards the same equilibrium as the receiver strategy of the other. The moment where the population reaches convexity marks the point where the sender strategy of $\alpha = 0.05$ aligns with its type's receiver strategy and this coincides with the moment when the population seems to gain real traction over the other type (see again the plot for trial 88 in Figure 5). In this particular case, reaching convexity seems to have made an important difference.

6 Conclusions

Vagueness presents a challenge to both procedural and instrumental pictures of rationality alike. In the context of game-theoretical models of language, this takes the form of a question about evolution: how can vagueness persist if it characterizes demonstrably less efficient communication? Most existing proposals in the literature attempt to explain this by considering agents with some degree of bounded rationality. Despite the many ways in which our rationality is inevitably limited, the argument for the pervasiveness of vagueness in natural language would be much stronger if one could also find associated advantages. In this paper, we argued that finding those might require us to go beyond a local notion of rationality, as two of the proposals reviewed here (Franke and Correia, 2017; O'Connor, 2014) suggest. We advocate moving towards an ecological approach, studying vagueness in more heterogeneous ecosystems. Language is part of a very complex system (Beckner et al., 2009) that involves many components interacting in often unpredictable ways. An ecology of vagueness would involve studying models where different populations can evolve and interact with each other, where different language games can be played between individuals, where the environment is uncertain and changing, or anything else that more closely approximates the real context of language evolution.

In light of this picture, we have proposed two variants of a concrete multi-population signaling model to test the hypothesis (Franke and Correia, 2017) that certain features of imprecise imitation, like promoting faster convergence and regularity, could prove beneficial in contrast with full precision. Analysing these models did not provide us with a clear-cut answer, and revealed that the story is much more nuanced than initially expected. In a variant where populations with and without imprecision interact tightly, although precision always has the upper hand, vagueness seems to take a long time to be weeded out. When we let populations interact more loosely, we see a more complex pattern of outcomes. These include scenarios where imprecise imitation dominates over full precision, showing that strategies with vagueness can actually, under certain circumstances, be more successful. Bringing several populations together in a more complex ecosystem thus allowed us to not only spell out and test the original intuition, but also learn about unforeseen effects. These models thus serve as an example of how moving to a more global perspective on rationality can allow us to achieve a more detailed awareness of the complex interactions that might be involved in sustaining vagueness in natural language. They can be seen, we believe, as a first step towards an ecology of vagueness.

References

- Barrett, Jeffrey A. (Aug. 2009). "The Evolution of Coding in Signaling Games". In: *Theory and Decision* 67.2, pp. 223–237.
- (Jan. 2010). "Faithful Description and the Incommensurability of Evolved Languages". In: *Philosophical Studies* 147.1, pp. 123–137.
- Beckner, Clay et al. (Dec. 2009). "Language Is a Complex Adaptive System: Position Paper". In: *Language Learning* 59, pp. 1–26.
- Beggs, Alan (2005). "On the convergence of reinforcement learning". In: *Journal of Economic Theory* 122.1, pp. 1–36.

- Benz, Anton, Gerhard Jäger, and Robert van Rooij, eds. (2006). *Game Theory and Pragmatics*. Hampshire: Palgrave MacMillan.
- Börger, Tilman and Rajiv Sarin (1997). “Learning Through Reinforcement and Replicator Dynamics”. In: *Journal of Economic Theory* 77.1, pp. 1–14.
- Boyd, Robert and Peter J. Richerson (1985). *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.
- Dawkins, Richard (Mar. 1978). *The Selfish Gene*. New Edition. London: Flamingo.
- Fawcett, Tim W., Steven Hamblin, and Luc-Alain Giraldeau (2013). “Exposing the behavioral gambit: the evolution of learning and decision rules”. In: *Behavioral Ecology* 24.1, pp. 2–11.
- Franke, Michael and José Pedro Correia (2017). “Vagueness and Imprecise Imitation in Signalling Games”. In: *The British Journal for the Philosophy of Science*.
- Franke, Michael, Gerhard Jäger, and Robert van Rooij (2011). “Vagueness, Signaling and Bounded Rationality”. In: *New Frontiers in Artificial Intelligence*. Berlin, Heidelberg: Springer, pp. 45–59.
- Franke, Michael and Elliott Wagner (2014). “Game Theory and the Evolution of Meaning”. In: *Language and Linguistics Compass* 8/9, pp. 359–372.
- Galeazzi, Paolo and Michael Franke (2017). “Smart Representations: Rationality and Evolution in a Richer Environment”. In: *Philosophy of Science* 84.3, pp. 544–573.
- Gilboa, Itzhak and Akihiko Matsui (1991). “Social Stability and Equilibrium”. In: *Econometrica: Journal of the Econometric Society*, pp. 859–867.
- Hagen, Edward H. et al. (2012). “Decision Making: What Can Evolution Do for Us?” In: *Evolution and the Mechanisms of Decision Making*. Ed. by Peter Hammerstein and Jeffrey R. Stevens. Cambridge, MA: MIT Press, pp. 97–126.
- Hamilton, W. D. (July 1964). “The genetical evolution of social behaviour. I”. In: *Journal of Theoretical Biology* 7.1, pp. 1–16.
- Hume, David and Michael P Levine (1738/2005). *A Treatise of Human Nature*. New York: Barnes & Noble.
- Huttegger, Simon (Jan. 2014). “How Much Rationality Do We Need to Explain Conventions?” In: *Philosophy Compass* 9.1, pp. 11–21.
- Huttegger, Simon et al. (July 2014). “Some Dynamics of Signaling Games”. In: *Proceedings of the National Academy of Sciences of the United States of America* 111.Suppl 3, pp. 10873–10880.
- Jäger, Gerhard (2007). “The Evolution of Convex Categories”. In: *Linguistics and Philosophy* 30.5, pp. 551–564.
- Jäger, Gerhard, Lars P. Metzger, and Frank Riedel (Nov. 2011). “Voronoi Languages: Equilibria in Cheap-Talk Games with High-Dimensional Types and Few Signals”. In: *Games and Economic Behavior* 73.2, pp. 517–537.
- Jäger, Gerhard and Robert van Rooij (Nov. 2007). “Language Structure: Psychological and Social Constraints”. In: *Synthese* 159.1, pp. 99–130.
- Keefe, Rosanna and Peter Smith (1997). “Introduction: theories of vagueness”. In: *Vagueness: A Reader*. Ed. by Rosanna Keefe and Peter Smith. MIT Press, pp. 1–57.
- Kohn, Marek (Nov. 2008). “Darwin 200: The needs of the many”. In: *Nature News* 456.7220, pp. 296–299.
- Kramer, Jos and Joël Meunier (Apr. 2016). “Kin and multilevel selection in social evolution: a never-ending controversy?” In: *F1000Research* 5.776.
- Lassiter, Daniel and Noah D. Goodman (2017). “Adjectival vagueness in a Bayesian model of interpretation”. In: *Synthese* 194.10, pp. 3801–3836.
- Lewis, David (1969/2011). *Convention: A Philosophical Study*. Oxford: Blackwell.
- Lipman, Barton L. (2009). “Why Is Language Vague?”
- Maynard Smith, John (1982). *Evolution and the Theory of Games*. Cambridge; New York: Cambridge University Press.
- Mühlenbernd, Roland (2011). “Learning with Neighbors: Emergence of Convention in a Society of Learning Agents”. In: *Synthese* 183, pp. 87–109.
- O’Connor, Cailin (2014). “The Evolution of Vagueness”. In: *Erkenntnis* 79.4, pp. 707–727.

- O'Connor, Cailin (2015). "Evolving to Generalize: Trading Precision for Speed". In: *The British Journal for the Philosophy of Science*.
- Okasha, Samir (2006). *Evolution and the Levels of Selection*. Oxford: Oxford University Press.
- Pagel, Mark (2009). "Human Language as a Culturally Transmitted Replicator". In: *Nature Reviews Genetics* 10, pp. 405–415.
- Qing, Ciyang and Michael Franke (2014). "Gradable Adjectives, Vagueness, and Optimal Language Use: A Speaker-Oriented Model". In: *Proceedings of SALT 44*. Ed. by Jessi Grieser et al. elanguage.net, pp. 23–41.
- Roth, Alvin E. and Ido Erev (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term". In: *Games and economic behavior* 8.1, pp. 164–212.
- Russell, Bertrand (June 1923). "Vagueness". In: *Australasian Journal of Psychology and Philosophy* 1.2, pp. 84–92.
- Sanborn, Adam N. and Nick Chater (2016). "Bayesian Brains without Probabilities". In: *Trends in Cognitive Sciences* 20.12, pp. 883–893.
- Simon, Herbert A. (1955). "A Behavioral Model of Rational Choice". In: *The Quarterly Journal of Economics* 69.1, pp. 99–118.
- (1986). "Rationality in Psychology and Economics". In: *Journal of Business*, S209–S224.
- Skyrms, Brian (June 1996). *Evolution of the Social Contract*. Cambridge University Press.
- (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press.
- Sorensen, Roy (Mar. 2009). "Sorites Arguments". In: *A Companion to Metaphysics*. Ed. by Jaegwon Kim, Ernest Sosa, and Gary S. Rosenkrantz. Second Edition. New York: John Wiley & Sons, pp. 565–566.
- Spike, Matthew et al. (2017). "Minimal requirements for the emergence of learned signaling". In: *Cognitive Science* 41.3, pp. 623–658.
- Taylor, Peter D. and Leo B. Jonker (1978). "Evolutionary Stable Strategies and Game Dynamics". In: *Mathematical Bioscience* 40.1–2, pp. 145–156.
- Thompson, Bill, Simon Kirby, and Kenny Smith (2016). "Culture Shapes the Evolution of Cognition". In: *Proceedings of the National Academy of Sciences of the United States of America* 113.16, pp. 4530–4535.
- Train, Kenneth E. (2009). *Discrete Choice Methods with Simulation*. Cambridge, MA: Cambridge University Press.
- Unger, Peter (1979). "There Are No Ordinary Things". In: *Synthese* 41.2, pp. 117–154.
- Vanderschraaf, Peter (1998). "The Informal Game Theory in Hume's Account of Convention". In: *Economics and Philosophy* 14.02, pp. 215–247.
- Vehtari, Aki and Janne Ojanen (2012). "A survey of Bayesian predictive methods for model assessment, selection and comparison". In: *Statistics Surveys* 6, pp. 142–228.
- Vul, Edward et al. (2014). "One and Done? Optimal Decisions From Very Few Samples". In: *Cognitive Science* 38.4, pp. 599–637.
- Wilson, D. S. (Jan. 1975). "A theory of group selection". In: *Proceedings of the National Academy of Sciences of the United States of America* 72.1, pp. 143–146.
- Wittgenstein, Ludwig (1953/2009). *Philosophical Investigations*. Ed. by P. M. S. Hacker and Joachim Schulte. Revised 4th edition. Chichester, U.K.: Wiley-Blackwell.
- Zollman, Kevin J. S. and Rory Smead (2010). "Plasticity and Language: An Example of the Baldwin Effect?" In: *Philosophical Studies* 147.1, pp. 7–21.