Vagueness and imprecise imitation in signalling games

Michael Franke & José Pedro Correia

Abstract

Signalling games are popular models for studying the evolution of meaning, but typical approaches do not incorporate vagueness as a feature of successful signalling. Complementing recent like-minded models, we describe an aggregate population-level dynamic that describes a process of imitation of successful behaviour under imprecise perception and realisation of similar stimuli. Applying this new dynamic to a generalisation of Lewis' signalling games, we show that stochastic imprecision leads to vague, yet by-and-large efficient signal use, and, moreover, that it unifies evolutionary outcomes and helps avoid sub-optimal categorisation. The upshot of this is that we see 'as-if'-generalisation at an aggregate level, without agents actually generalising.

1 Introduction

Many concepts and expressions are vague. A vague category knows clear cases that fall under it, clear cases that do not, and also so-called borderline cases. Borderline cases do not clearly belong to the category, nor do they clearly not belong, and there may be differences between borderline cases in terms of how well they represent the category in question. Vagueness does not seem to dramatically affect the success of everyday communication, but it is troublesome for some of the most prominent theories of language and meaning. This is especially so for the logico-positivist tradition of Frege, Russell, and early Wittgenstein, which is challenged by the paradoxes vagueness gives rise to.

There are other intriguing aspects about vagueness. One perplexing issue is how vagueness could arise and be maintained in the first place. This is an apparent puzzle for functionalist accounts that maintain that concepts and linguistic meanings evolved driven towards efficiency. Lipman (2009) argues that, in common interest signalling situations, the existence of unclear borderline cases entails inefficiency of categorisation or communication, or at least no advantage. The challenge is then to explain how vagueness can persist both: (i) under evolutionary pressure to be optimally discriminative, and (ii) without undermining the possibility of evolving, learning and communicating with a meaningful language. A number of authors have consequently tried to explain why vagueness evolved as something that is itself useful (e.g. de Jaegher 2003; van Deemter 2009; Blume and Board 2014). Others have argued that vagueness is a natural byproduct of limitations in information processing (e.g. Franke et al. 2011) or of generalisation in low-level learning strategies (e.g. O'Connor 2014a). This paper contributes to the latter line of thought.

We believe that vagueness in language and thought may have many reasons, not just one. We focus here on one *a priori* plausible reason for why vagueness is natural and pervasive. The idea is that vagueness is, at least in part, due to imprecision in the perception of similar stimuli and imprecision in the realisation of similar responses. On this view, vagueness in language may be seen as a necessary sub-optimality due to limitations in another domain of cognition, first and foremost perception. It could then be speculated that the whole system, perception and language together, may be an almost optimal adaptation in a larger frame of reference, for example if we take into account the metabolic costs for increased perceptual accuracy. We will not engage in such speculation here. Rather we will explore the consequences of perceptual limitations on processes of meaning evolution in a suitable formal framework. In other words, in order to address Lipman's challenge seriously, and not just hand-wave it away by appeal to the naturalness of vagueness, one needs to spell out how exactly could confusability of stimuli come into play in a process of meaning evolution and how it could lead to vague but by-and-large informative signal meaning. The formal model that this paper introduces does exactly that. But it also does more. We find that confusability of stimuli can regularise and systematise evolving meaning. This suggests a possibly advantageous side-effect that a natural cause of vagueness may have and that might compensate some of the disadvantages that vagueness may have for communicative efficiency (c.f. O'Connor 2014a).

The next section introduces the background against which the work presented here can be appreciated. Section 3 introduces a generalisation of the replicator dynamic which is derived from the idea that agents imitate other agents' behaviour while possibly confusing similar states.[1] This imprecise imitation dynamic is explored in Section 4. Section 5 reflects and compares our approach to that of others. Appendix A provides formal detail.

2 Background

2.1 Sim-max games & conceptual spaces

Signalling games, as introduced by Lewis (1969), have a sender and a receiver. The sender knows the true state of the world, but the receiver does not. The sender can select a signal, or message, to reveal to the receiver, who then chooses an act. In Lewis' games, if the receiver chooses the act that corresponds to the actual state, the play is a success, otherwise a failure. Certain regular combinations of sender signalling and receiver reaction make messages meaningful, in the sense that their use is correlated systematically to certain states or acts. To investigate the conditions under which such meaning-generating behaviour can evolve is a topic that we are only beginning to fully understand (e.g. Wärneryd 1993; Blume et al. 1993; Huttegger 2007a; Pawlowitsch 2008; Barrett 2009; Huttegger et al. 2010; Skyrms 2010).

Similarity-maximising (short: sim-max) games are a variation of Lewis' games where the receiver's actions are equated with the state space (one can think of the actions as choosing states) and different states are allowed to be more or less similar to one another. While Lewis' games treated communicative success as a matter of black and white, sim-max games allow for shades of grey: the more similar the receiver's interpretation is to the actual state, the better. Signalling games with utility-relevant similarities in the state space are fairly standard in economics (e.g. Spence 1973; Crawford and Sobel 1982), but have received particular attention in a more philosophical context for reasons that will become clear presently.

Formally, a sim-max game consists of a set of states T, a set of messages M typically with much fewer messages than states, a probability distribution $P \in \Delta(T)$ such that $P(t)$ gives the probability that state t occurs, a similarity metric[2] on states $\mathrm{Sim} : T \times T \to \mathbb{R}$ such that $\mathrm{Sim}(t_1, t_2)$ is the (physical) similarity between t_1 and t_2, and a utility function $\mathrm{U} : T \times T \to \mathbb{R}$ such that $\mathrm{U}(t_1, t_2)$ is the payoff for sender and receiver for a play with actual state t_1 and receiver interpretation t_2. We identify the receiver's acts with the states of the world, so that the game is one of guessing the actual state, so to speak. For the modelling purposes of this paper, we make the simplifying assumption that T is a set of points in Euclidean space, whose closeness to each other tracks physical similarity. Perceived similarity, where it is necessary, would be a monotonic function of physical similarity. Likewise, the utility function should be a monotonically increasing function of physical similarity.[3] Non-probabilistic sender behaviour can

[1]Much previous work has investigated the interplay of adaptive dynamics, including imitation-based update protocols, on the one hand, and noise or mutation, on the other hand (e.g. Foster and Young 1990; Fudenberg and Harris 1992; Kandori et al. 1993; Young 1993; Fudenberg and Imhof 2006). While this line of research often looks at finite populations and the extreme long-term behavior of the system under generic randomness, the focus here is on a quite particular source of stochastic noise and its quite particular role in the evolution of meanings through signalling.

[2]A metric is a function of distance between any two points in a given space. It should satisfy certain axioms that ensure behaviour that one would intuitively expect from the words 'similarity' and 'distance' alone, but these details do not matter for the purposes of this paper. The assumption that state similarity forms a metric is conceptually loaded, but we follow the literature here and conceive of it as a first and pragmatic simplification, possibly to be dispensed with later.

[3]Section 4.1 motivates particular choices of similarity and utility functions that we will explore in more detail.

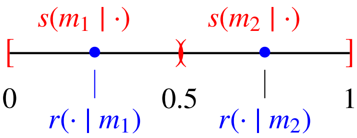

Figure 1: Example of a Voronoi language on $T = [0, 1]$. The pure sender strategy s uses one signal for the lower half, and another for the upper half of the unit interval. The pure receiver strategy r selects the central elements in the respective intervals. See Jäger et al. (2011) for further details.

be represented by a pure strategy $s \in M^T$, which deterministically defines which message would be used for each state. Similarly, a pure receiver strategy is a function $r \in T^M$.

Jäger et al. (2011) showed that the evolutionarily stable states of sim-max games with infinitely many states in n-dimensional Euclidean space $T \subseteq \mathbb{R}^n$ and a quadratic loss function for utilities $\mathrm{U}(t_1, t_2) = -(t_1 - t_2)^2$ are remarkably systematic: the evolutionarily stable states are demonstrably so-called Voronoi languages. Roughly put, a Voronoi language is a pair of sender and receiver strategies, such that the sender strategy partitions the state space into convex categories, while the receiver's interpretations are the central spots in each category. A subset X of $\mathbb{R}^n$ is convex if, informally put, all points in X are connected via a straight line that lies entirely in X; X has no gaps or dents. For example, if T is the unit interval and all states are equiprobable, a Voronoi language with two messages could have the sender use one message exclusively for all points in the lower half of the unit interval and another for all points in the upper half; the receiver's interpretations of messages are the central points, .25 and .75, in the respective intervals. See Figure 1 for an illustration.

It is intuitive to think that linguistic and conceptual categories are orderly in such a manner. For example, if you consider two people 'tall', one with a height of 2m and another with a height of 2.2m, it would be difficult to defend not considering a person with a height of 2.1m 'tall' as well. Another example where this intuition is additionally supported by empirical data is that of colour categorisation. The World Color Survey project (Cook et al. 2005; Kay et al. 2009) collected colour naming data for 110 unwritten languages of 45 language families. For the great majority of these languages a pattern can be observed: basic colour terms are by and large convex (Regier et al. 2007; Jäger 2010). It is premature to argue that these observations can be extended to all cases of categorisation. However, for the cases where it does apply, the result of Jäger et al. (2011) is interesting because it demonstrates that signalling can impose this kind of orderly categories on a metric space without that being the ulterior purpose of it all.[4] Finally, these considerations are in line with a prominent school of thought in comparative linguistics which also assumes that more abstract conceptual domains (e.g., spatial-topological relations, temporal reference, or the meanings expressible by indefinite pronouns) are preferably carved up by the languages of the world in such a way that meanings are connected regions on a so-called 'semantic map' (e.g. Croft 2003; Haspelmath 2003; Levinson et al. 2003).

There are at least two potential ways of interpreting the signalling setup. Sender and receiver can be distinct entities, whose purpose is to communicate effectively about the actual state. In that case, evolving Voronoi languages would explain why linguistic categories are well-behaved and orderly in the way they appear to be. More abstractly, sender and receiver can also be thought of as distinct modules in a single system, where the first module must discretise the information it is fed by selecting a small sample of, suggestively, category labels. These are passed to a second module that tries to decode the original information. In this case, evolving Voronoi languages would explain why conceptual categories are well-behaved and orderly in the way that they appear to be. Seen in this light, sim-max games may provide a foundation to approaches in cognitive semantics that rely on the notion of conceptual spaces. Gärdenfors (2000, pp. 70–77), for example, has prominently argued that natural categories are convex regions in conceptual space. If the conceptual space has a suitable metric, convex categories can be derived from a set of prototypes. The category corresponding to prototype p is the set of points

[4] For the concrete case of colour categorisation, see also Jäger and van Rooij (2007) and Correia et al. (2016)

that are more similar to p than to any other. In this way, Gärdenfors argues, an efficient categorisation system can be obtained: storing the prototypes lets us recover the categories without having to store each category's extension. However, what is left unexplained is where the prototypes come from, and why we would not see just any distribution of prototypes as an equally efficient classification system. This is where sim-max games can contribute a principled approach to deriving, in an independent way, not only convex categories but also prototypical exemplars belonging to them. These ideas, and more, are developed further by Jäger (2007), Jäger and van Rooij (2007), Jäger et al. (2011), O'Connor (2014b), among others.

2.2 Vagueness in sim-max games and conceptual spaces

This outline of an approach to categorisation using sim-max games leaves some problems unaddressed. One of them is that, usually, natural categories for continuously variable stimuli neither have clear boundaries, nor do they have unique, point-valued prototypes. We would like to account for the possibility of such vagueness, in particular: (i) clear positive examples of a vague category should show a gradient transition to clear negative examples; (ii) prototypes should likewise be gradient regions, peaking at the centre of the vague category they represent.

Douven et al. (2011) show that Gärdenfors's conceptual spaces approach can be extended to account for the existence of borderline cases. From the assumption that prototypes are extended, yet convex regions in conceptual space, a construction algorithm is available that yields 'collated Voronoi diagrams' with thick boundaries representing borderline regions. Decock and Douven (2012) show further how it is possible to arrive at a gradient transition between categories, by weighing in the distance of different borderline cases to various prototypical regions. This accounts for the first of the two desiderata mentioned above, but still assumes that crisp prototype regions must be given.

Alternative approaches are taken by, for example, Franke et al. (2011) and O'Connor (2014a), who show different ways how the above desiderata can be met by evolving strategies in sim-max games. To illustrate what vague signalling would look like, let us briefly consider a sim-max game with six equiprobable states and two messages and what its equilibria would be like (see Jäger 2007 and O'Connor 2014a for a more thorough discussion of equilibria of sim-max games). We assume that utilities are linearly decreasing with decreasing similarity. Figure 2 shows three pairs of sender and receiver strategies. States are arranged according to their similarity: the closer they are to each other, the more similar they are. The pair in Figure 2a is not an equilibrium, because the sender's non-convex use of signals is suboptimal given the receiver's behaviour. Namely, if m_1 is interpreted as t_2 and m_2 as t_4, then the sender would get a higher payoff from sending m_1 in t_3 than from sending m_2, because (by assumption) t_3 is more similar to t_2 than t_5 is. In contrast, Figure 2b shows a maximally efficient equilibrium. This is a partial pooling equilibrium in the sense that the speaker uses the same message for several states. Partial pooling equilibria can be less inefficient than other non-pooling equilibria, if there are enough messages. The strategy pair of Figure 2b is maximally efficient for the two-message case. So, while partial pooling may entail inefficiency in some sense, and while partial pooling can hamper the evolution of maximally efficient signal use (Huttegger 2007b; Pawlowitsch 2008; Huttegger et al. 2010), this is orthogonal to our concerns about vagueness. A regular and natural yet vague signal use would look like the pair in Figure 2c. This is not an equilibrium, but it gets close, so to speak. It shows smooth transitions across similar states at the boundaries of categories and across acts around the most prototypical instances (as indicated by decreasing thickness of arrows in Figure 2c).

2.3 Vagueness, functional pressure & transmission biases

The approach we take here is similar in spirit to that of Franke et al. (2011) and O'Connor (2014a), but different in relevant detail. A more in-depth comparison is deferred until Section 5. Let us first motivate our approach here, and spell it out in more detail in the following section.

Our conceptual starting point is the widely shared conviction that language is shaped by at least two forces, which may, on occasion, pull in opposite direction. On the one hand, there is functional pressure

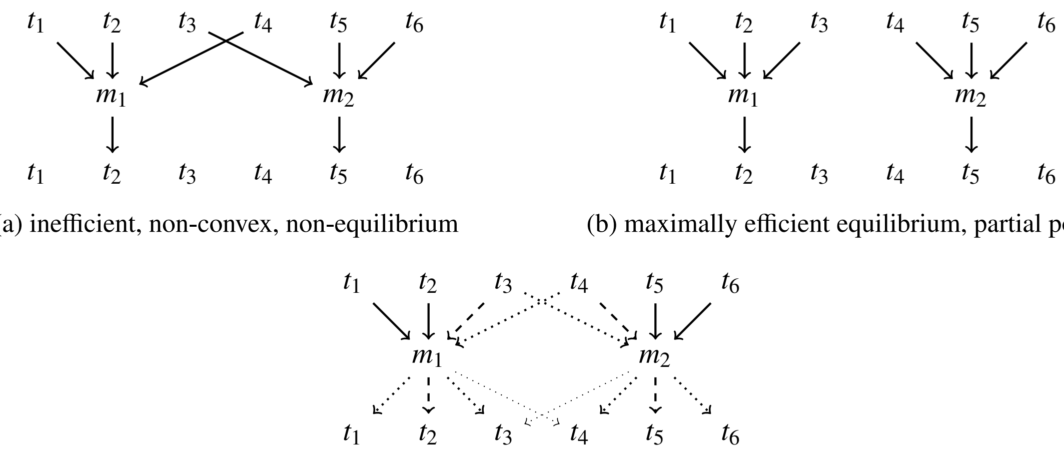

(a) inefficient, non-convex, non-equilibrium

(b) maximally efficient equilibrium, partial pooling

(c) almost efficient, probabilistic, intuitively vague

Figure 2: Examples of strategy pairs in sim-max games. Sender strategies map states onto messages (top two rows); receiver strategies map messages onto states (bottom two rows). In Figure 2, the thickness of arrows indicates the probability of a choice in a probabilistic strategy.

towards efficient communication. On the other hand, there is systematic error, noise or imprecision in the transmission of linguistic behaviour, knowledge or traits. As an example of the latter, consider a child learning syntactic rules from a parent generation. The child must infer these unobservable rules from observable speech. Inductive biases may influence which syntactic rules are likely to be inferred from (finite) parental input. Over the course of many generations, the effects of such biases can, in a manner of speaking, 'accumulate' and lead to surprising results, such as the evolution of compositional form-meaning mappings or the use of regular recursive syntactic structure. This can happen when the effects of transmission biases are isolated, as in iterated learning models (Kirby and Hurford 2002; Smith et al. 2003; Griffiths and Kalish 2007; Kirby et al. 2014), or when they interact with functional pressure towards efficient communication, for example as formalised in the replicator mutator dynamic (e.g. Nowak et al. 2000; Nowak et al. 2001).

The emphasis of previous models that studied the effects of transmission infidelity on the evolution of language has been on inductive biases and the systematicity, compression and regularisation that they can introduce. Here, we would like to show that transmission noise of a different kind can lead to regularisation as well and also give rise to vague meaning. We find that shared perceptual biases that perturb the transmission of successful signalling behaviour can regularise, facilitate and accelerate the evolution of meaning conventions, albeit at the cost of vagueness. Concretely, we formalise the expected change in the behaviour of a population of agents that try to imitate other agents' signalling behaviour. We assume, however, that both observation of others' behaviour, as well as realisation of behaviour, are both systematically perturbed by noise. The resulting population-level dynamic generalises the replicator dynamic (Taylor and Jonker 1978) in its interpretation as a cultural evolutionary dynamic based on imitation (e.g. Helbing 1996; Schlag 1998). The inclusion of confusability of stimuli does not undermine the possibility of evolving communicative signalling behaviour. Instead, it leads to the evolution of vague meanings. It moreover accelerates the emergence of communicative signalling because it unifies and regularises evolutionary outcomes, making it appear as if agents were applying inductive biases or generalising over partial observations, when this is actually the sole effect of confusion of similar stimuli.

$$\xrightarrow{\text{N}} t_1 \in T \xrightarrow{\text{S}} m \in M \xrightarrow{\text{R}} t_2 \in T$$

$$P(t_1) \qquad \sigma(m \mid t_1) \qquad \rho(t_2 \mid m)$$

Figure 3: A round of play in a sim-max game with behavioural strategies: Nature (N) chooses a state t_1 with probability $P(t_1)$; a random sender (S) selects a message m with probability $\sigma(m \mid t_1)$; a random receiver (R) selects a state t_2 with probability $\rho(t_2 \mid m)$. Payoff for both sender and receiver is given by $\mathrm{U}(t_1, t_2)$.

3 Imprecise imitation

Signalling agents can adapt their dispositions to act, given some feedback about their past success, in multiple ways. Usually, we would assume that changes in behaviour should, at least on average, lean towards increasing chances of communicative success. Such behavioural adaptations can be described at different levels of abstraction. At the level of individual agents, we can picture a more or less idealised process of how each agent adapts dispositions for future actions based on various pieces of information available to the agent. More abstractly, at the level of a population of agents, we can describe how average behavioural dispositions will evolve. The population-level perspective abstracts over small stochastic fluctuations and zooms in on the general tendency or direction of evolution that ensues from behaviour at the agent-level.

To better understand the interaction of confusability of stimuli and selective pressure towards successful communication, we look at a population-level dynamic that describes the most likely evolutionary path of a population of signalling agents who imitate other agents' behaviour but are liable to confuse states for one another. In the special limiting case where state confusability vanishes, the process is just the well-known replicator dynamic (Taylor and Jonker 1978), which should therefore be briefly reviewed first.

3.1 Replicator dynamic in behavioural strategies

Fix a sim-max game with finite states T and messages M. As usual, we assume that the receiver chooses states in T in response to messages. Let $P(\cdot) \in \Delta(T)$ be the prior distribution over states and $\mathrm{U} : T \times T \to \mathbb{R}$ the utility function shared by senders and receivers in the population. A behavioural strategy is a function that maps an agent's choice points to a probability distribution over available choices.[5] The sender's behavioural strategies are functions $\sigma \in \Delta(M)^T$, thus mapping each state $t \in T$ to a probability of each message $m \in M$ being sent in t; the receiver's are functions $\rho \in \Delta(T)^M$, thus mapping each message $m \in M$ to a probability of each interpretation $t \in T$ being chosen in response to m. Although behavioural strategies are probabilistic, evolutionary modelling usually imagines that every individual agent has a non-probabilistic strategy. Behavioural strategies then capture average population behaviour. Assuming a virtually infinite population, the number $\sigma(m \mid t)$, for instance, is then the probability that a randomly sampled sender would send message m if the actual state was t. Similarly, $\rho(t \mid m)$ is then the probability with which a randomly sampled receiver interprets m as t. A play of a single evolutionary game with behavioural strategies is illustrated in Figure 3.

[5] Our focus is on behavioural strategies, not mixed strategies, i.e., probability distributions over functions from each choice point to an act, such as $s \in \Delta(M^T)$. Dynamics on behavioural strategies assume that agents can adapt their behaviour locally, i.e., independently at each choice point. Our focus on behavioural strategies greatly reduces the complexity of the dynamic and simplifies numerical simulations. But it also seems the more plausible choice for imitation-based update protocols of the kind we consider here: agents only observe how, on some occasion, some other agent behaved in one particular situation, not how that agent would behave in all relevant choice situations; they imitate the use of a single word, so to speak, not a whole lexicon. (See Cressman 2003 for more on the difference between dynamics on mixed or behavioural strategies.) It is in this respect that the cultural evolutionary dynamic introduced here differs most visibly from the replicator mutator dynamic (e.g. Nowak et al. 2000; Nowak et al. 2001), which operates on mixed strategies and is motivated by assumptions of (asexual) biological inheritance with transmission infidelity.

The expected utility of choices at each choice point is:[6]

$$\mathrm{EU}(m,t,\rho) = \sum_{t' \in T} \rho(t' \mid m)\, U(t,t')$$

$$\mathrm{EU}(t',m,\sigma) = \sum_{t \in T} P(t \mid m)\, U(t,t')\,, \quad \text{where} \qquad P(t \mid m) \propto P(t)\,\sigma(m \mid t)\,.$$

Expected utilities are sums, over all possible concrete outcomes, of the utilities of these outcomes, weighted by how likely they occur. For instance, the expected utility for the sender of choosing message m, given state t and a receiver strategy ρ is the sum, for each interpretation t', of the probability that some random receiver will choose t' given m times the actual utility of pair t and t'.

The discrete-time replicator dynamic tracks changes in frequency of choices in the population as proportional to their expected utilities (Hofbauer and Sigmund 1998):[7]

$$\sigma'(m \mid t) \propto \sigma(m \mid t)\,\mathrm{EU}(m,t,\rho)\,, \qquad \rho'(t \mid m) \propto \rho(t \mid m)\,\mathrm{EU}(t,m,\sigma)\,. \tag{3.1}$$

For a given choice point, say a state t, the probability of seeing m played by an average agent in the population after the update is proportional to the probability of seeing it before the update, which is $\sigma(m \mid t)$, times the expected utility of m at state t. Intuitively put, frequencies of choices change by a gradient of current frequency and a measure of how good they are.

The replicator dynamic is an abstract population-level dynamic that describes the mean expected change of behavioural dispositions in a population of signallers. There are several ways of deriving the replicator dynamic from agent-level processes of behavioural adaptation. We focus here on one of the simplest: imitation of success (see Sandholm 2010). The intuitive idea is the following: every now and then a random agent gets a chance to alter his behaviour for one of his choice points (say this is a sender who gets to 'reconsider' his choice of message for state t); the revising agent then observes what some random agent does at t, say m, and will henceforth play m in t with a probability given by the expected utility of m for state t. See Appendix A.1 for a derivation of the standard replicator dynamic from this update scheme.

3.2 Noise perturbed conditional imitation

Imitation of success, as described above, presupposes that agents make no mistakes when observing states, or choosing interpretations. This may not always be an appropriate assumption, especially when some states can be perceptually similar and therefore likely to be confused for one another. Human performance in these situations has been studied by a number of authors in experimental psychology. In a stimulus identification experiment (Luce 1963), subjects are presented in each trial with a stimulus to be identified out of a fixed set. The more similar the stimuli are to each other, the larger the number of errors subjects make. For example, in Robert Nosofsky's experiments (1986), stimuli consisted of 16 semicircles with a radial line from the centre to the rim, varying in size (radius of 0.478, 0.500, 0.522, or 0.544 cm) and angle of the line (50°, 53°, 56°, or 59°). In over 9000 trials, the two subjects identified the stimulus correctly only approximately 44% and 35% of the time. Similar experiments have been conducted with other identification tasks, for example, for frequency and intensity of tones, taste, hue of colours, and magnitude of lines and areas (see Donkin et al. 2015). The empirical data confirms not only the pervasiveness of variation in the subjects' ability to correctly identify stimuli, but also the relation between similarity and likelihood of stimulus confusion, and further suggests that the phenomenon might extend to all types of perception.

[6] The notation $\propto$, for 'proportional to', that is used here and hereafter means that the right-hand side might still need to be normalised so as to have probabilities sum to one. In general, writing $P(x) \propto f(x)$ for any function $f : X \to \mathbb{R}$ is shorthand for $P(x) = \frac{f(x)}{\sum_{x' \in X} f(x')}$.

[7] The formulation given in Equation (3.1) is adequate only for cases like the one we will be looking at in Section 4, where utilities are always non-negative and expected utilities are always positive.

$$
\begin{array}{cccccccccc}
\xrightarrow{\text{N}} & t_{\text{a}} \in T & \longrightarrow & t_{\text{o}} \in T & \xrightarrow{\text{S}} & m \in M & \xrightarrow{\text{R}} & t_{\text{i}} \in T & \longrightarrow & t_{\text{r}} \in T \\
& P_a(t_{\text{a}}) & & P_o(t_{\text{o}} \mid t_{\text{a}}) & & \sigma(m \mid t_{\text{o}}) & & \rho(t_{\text{i}} \mid m) & & P_r(t_{\text{r}} \mid t_{\text{i}})
\end{array}
$$

Figure 4: A round of play in a sim-max game with probabilistic confusability of states: Nature (N) chooses an actual state with probability $P_a(t_{\text{a}})$; a randomly sampled sender (S) observes t_{o} with probability $P_o(t_{\text{o}} \mid t_{\text{a}})$ and subsequently selects m with probability $\sigma(m \mid t_{\text{o}})$; a randomly sampled receiver (R) intends to realise interpretation t_{i} with probability $\rho(t_{\text{i}} \mid m)$ but actually realises interpretation t_{r} with probability $P_r(t_{\text{r}} \mid t_{\text{i}})$. Payoff for both sender and receiver is given by $\text{U}(t_{\text{a}}, t_{\text{r}})$.

In the context of imitation of success in a game where states have a degree of similarity between them, the possibility of agents mistaking one state for another is thus something that should be taken into account. Imitation dynamics could be affected by at least two sources of probabilistic noise:

1. *observation noise:* whenever a state t_{a} actually occurs, the probability that an agent observes it as t_{o} is $P_o(t_{\text{o}} \mid t_{\text{a}})$;
2. *realisation noise:* whenever an agent intends to realise interpretation t_{i}, the probability that t_{r} is realised is $P_r(t_{\text{r}} \mid t_{\text{i}})$.

A round of play of a sim-max game with these two sources of confusability of states is pictured in Figure 4. Note that, and this is crucial, senders respond to observations not actual states, receiver strategies determine intentions not realised states, but payoff is calculated based on actual and realised states. Behavioural strategies thus encode what agents actually do from their subjective point of view; or, put differently, what they would do in a noise-free world. The noise-free situation depicted in Figure 3 is then the special case where $P_o(t_{\text{o}} \mid t_{\text{a}}) = 1$ iff $t_{\text{a}} = t_{\text{o}}$ and $P_r(t_{\text{r}} \mid t_{\text{i}}) = 1$ iff $t_{\text{i}} = t_{\text{r}}$.

The presence of observation and realisation noise also affects imitation of successes. Here, we focus on the main ideas, formal detail is provided in Appendix A.2. Take a sender who gets to revise behaviour for state t. What agents can plausibly revise by imitation is their pure strategy, which maps perceived states onto messages. When an agent gets to revise his strategy for the perceived state t, this need not necessarily be the actual state. Moreover, when that agent observes t, another agent may perceive yet a different state. Given what that latter agent perceives, his actual (pure) strategy will determine what he plays. In sum, to describe how likely the potential imitator observes a message choice m, we are interested in the conditional probability $P_o(m \mid t)$ that some other random agent selects m when the first agent perceives t. This $P_o(m \mid t)$ is derived from the prior probability of states, the current sender population behaviour σ and the given observation noise (see Appendix A.2). Eventually, the imitating agent adopts m as his choice for t with a probability given by the expected utility of sending m when perceiving state t. Expected utility should, of course, take the probabilistic confusability of states into account as well.

Similar considerations apply to the receiver side. If an agent gets a chance to change his intended interpretation of m, we need to look at the conditional probability $P_o(t \mid m)$ of observing another agent realise interpretation t given that the first agent (and therefore the second as well) perceived m. This depends on observation and realisation noise, as well as on the current receiver population behaviour ρ.

Appendix A.2 shows how imprecise imitation of this sort leads to mean changes in population frequencies of choices that can be covered by the following discrete-time formulation:

$$\sigma'(m \mid t) \propto P_o(m \mid t)\, \text{EU}(m, t, \rho)\,, \qquad\qquad \rho'(t \mid m) \propto P_o(t \mid m)\, \text{EU}(t, m, \sigma)\,. \tag{3.2}$$

This looks very much like the discrete-time formulation of the standard replicator dynamic in (3.1), but there are, of course, the aforementioned differences. Firstly, expected utility here takes stochastic confusability of states into account. Secondly, where the standard replicator dynamic had probabilities $\sigma(m \mid t)$ and $\rho(t \mid m)$, we now have $P_o(m \mid t)$ and $P_o(t \mid m)$ respectively. If there is no observation or realisation noise, $P_o(m \mid t)$ reduces to $\sigma(m \mid t)$ and $P_o(t \mid m)$ reduces to $\rho(t \mid m)$. The imprecise imitation dynamic in (3.2) conservatively extends the classic case in (3.1).

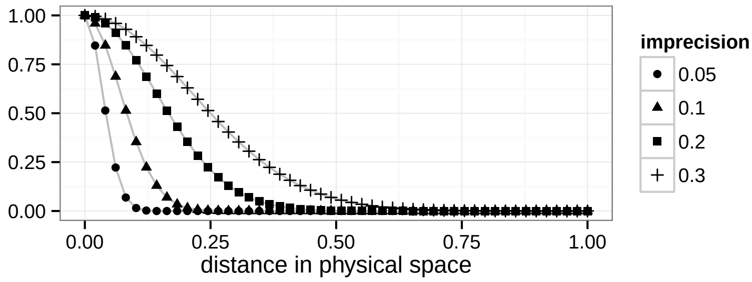

Figure 5: Examples of Nosofsky-similarity for different values of imprecision.

4 Exploring imprecise imitation

How does a tendency to confuse similar states interact with selective pressure towards more efficient signalling strategies under the imprecise conditional imitation dynamic? If confusion probabilities are moderate and regular, in that they track similarity of states, we might expect a regularising effect also on evolving signalling strategies. Indeed, we hypothesise that confusion of states can give rise to population-level aggregate behaviour that looks as if information about what is good for one state percolates also to similar states. In other words, imprecise imitation may make signalling behaviour on average look as if agents generalise across similar states, even if no agent actually generalises. To explore whether confusion of states can have this effect, we turn to numerical simulation.

4.1 Setting the stage

To obtain concrete results, we must fix how to represent states, similarity of states, the conditional confusion probabilities P_o and P_r and the utilities of our sim-max games. Confusion probabilities between states should be a function of perceptual similarity: the more similar two states are, the more likely they could be mistaken for each other.

Let the state space consist of $n_s \geq 2$ states that are equally spaced across the unit interval, including 0 and 1. All states occur, for simplicity, with the same probability (i.e., P_a is uniform). The distance $\left| t_\mathrm{i} - t_\mathrm{j} \right|$ is the objective, physical similarity between two states t_i and t_j. Distance in physical space feeds into a perceptual similarity function, as described by Nosofsky (1986):

$$\mathrm{Sim}(t_\mathrm{i}, t_\mathrm{j}\, ;\, \alpha) = \begin{cases} 1 & \text{if } \alpha = 0 \text{ and } t_\mathrm{i} = t_\mathrm{j} \\ 0 & \text{if } \alpha = 0 \text{ and } t_\mathrm{i} \neq t_\mathrm{j} \\ \exp\left(-\frac{\left| t_\mathrm{i} - t_\mathrm{j} \right|^2}{\alpha^2}\right) & \text{otherwise,} \end{cases}$$

where $\alpha \geq 0$ is an imprecision parameter. When $\alpha = 0$ agents perfectly discriminate between states; when $\alpha \to \infty$ agents cannot discriminate states at all. Figure 5 gives an impression of Nosofsky-similarity for different parameter values. Other formalisations of perceptual similarity are possible, including ones that allow for different discriminability in different areas of the state space, but we stick with Nosofsky's similarity function for the time being, because it is mathematically simple, and an established notion in mathematical psychology.

We further assume that the probability of confusing any two states t_i and t_j is proportional to their perceived similarity and, to keep matters simple, that observation noise P_o is simply the same as realisational confusability P_r and that both are governed by the same imprecision parameter α:

$$P_o(t_\mathrm{o} \mid t_\mathrm{a}) \propto \mathrm{Sim}(t_\mathrm{o}, t_\mathrm{a}\, ;\, \alpha)\,, \qquad\qquad P_r(t_\mathrm{r} \mid t_\mathrm{i}) \propto \mathrm{Sim}(t_\mathrm{r}, t_\mathrm{i}\, ;\, \alpha)\,.$$

For $\alpha = 0$, we obtain trivial confusion probabilities: everything is reduced to perfect imitation and the replicator dynamic. For $\alpha > 0$, any state can be confused for any other state with some positive probability. For the uninteresting case of $\alpha \to \infty$, confusion is maximal and every state can be perceived or realised as any other state with probability $^1/_{|T|}$.

As for utility, we define it in terms of similarity and introduce another free parameter, $\beta \geq 0$. The intention is for this parameter to model the amount of tolerable pragmatic slack, which should be allowed to vary separately from perceptual imprecision:

$$\mathrm{U}(t_\mathrm{i}, t_\mathrm{j}\,;\beta) = \mathrm{Sim}(t_\mathrm{i}, t_\mathrm{j}\,;\beta)\,.$$

With $\beta = 0$, we regain the case of Lewis' games where only a perfect match of actual state t_i with receiver interpretation t_j leads to a positive payoff. The higher β the more acceptable a wider environment of interpretations t_j around t_i is (see Figure 5). This choice of utility function is governed partly by convenience, but also because we believe it has the right general properties for a communicative payoff function. Unlike utilities that are, say, linearly or quadratically decreasing in physical distance (c.f. Jäger et al. 2011; Franke et al. 2011), utilities that are exponentially decreasing in negative quadratic distance can model situations where a small amount of imprecision in communication is tolerable, whereas similarly small differences in intolerably far away interpretations matter very little, with a smooth transition between these regimes (c.f. O'Connor 2014a).

In order to simplify the analysis, we focus on games with two messages, i.e., we fix $|M| = 2$. In sum, to structure thinking about the behaviour of our imprecise imitation dynamic, the system is governed by three parameters: the number $n_s = |T|$ of states in the state space, imprecision α and tolerance β.

4.2 Simulation set-up

We ran 50 trials of the discrete-time dynamic in (3.2), starting with randomly sampled sender and receiver strategies, for each triplet of independent parameter values: $n_s \in \{6, 10, 50, 90\}$, $\alpha \in \{0, 0.05, 0.1, 0.2, 0.3\}$, $\beta \in \{0.05, 0.1, 0.2, 0.3\}$. Each trial ran for a maximum of 200 update steps. A trial was considered converged, and thus stopped before the maximum of 200 rounds, if the total amount of change between strategies before and after an update step was smaller than a suitably chosen threshold. It is not guaranteed that strategies at halting time had converged to the eventual attracting state, whether they ran for 200 rounds or not. Our notion of convergence is therefore only a categorical measure for reaching a certain (well-considered, but eventually arbitrary) degree of stability. In other words, our notion of 'convergence' is a measure of relative speed: is it true that the system reached a state in which evolutionary adaptations had slowed down almost to a halt before 200 update steps? This is motivated by practical concerns regarding length of simulation time, but also theoretically justifiable, because we hypothesise that confusability of states leads to regularisation of evolving strategies, which would show exactly in an increased speed of evolutionary trajectories towards well-behaved and regular signalling behaviour.

Representative examples for resulting strategy pairs are given in Figure 6. Figure 6a shows a strategy pair at stopping time with 90 states, tolerance $\beta = 0.1$ and imprecision $\alpha = 0$. Zero imprecision means that the trial was effectively an application of the standard replicator dynamic. Noteworthily, the given sender strategy approximates a pure sender strategy that crisply partitions the state space into non-convex sets. The irregular shape of the receiver strategy suggests that the pictured strategy pair has not yet reached a stable state. Indeed, the trial was stopped when reaching the maximum of 200 rounds. In contrast, the outcome of a trial with identical parameters, except with imprecision $\alpha = 0.05$, which is shown in Figure 6b, had converged (in our technical sense) after 99 rounds. The sender strategy shows a smooth blending from one 'category' to the other, and the receiver's interpretations are rather extended curves, peaking at a central point in the relevant 'categories.'

These examples already show two interesting things. Firstly, inclusion of imprecision can lead to seemingly well-behaved, yet vague strategies in the sense that we are after (see again Section 2.2). The sender strategy in Figure 6b identifies clear positive and clear negative cases for each signal, with a smooth transition in-between. The receiver's interpretations of signals can be seen as smoothed-out prototype regions. Secondly, (sender) strategies can approach non-convex pure strategies under the replicator dynamic and linger there for vast amounts of time, possibly indefinitely. We see this in our

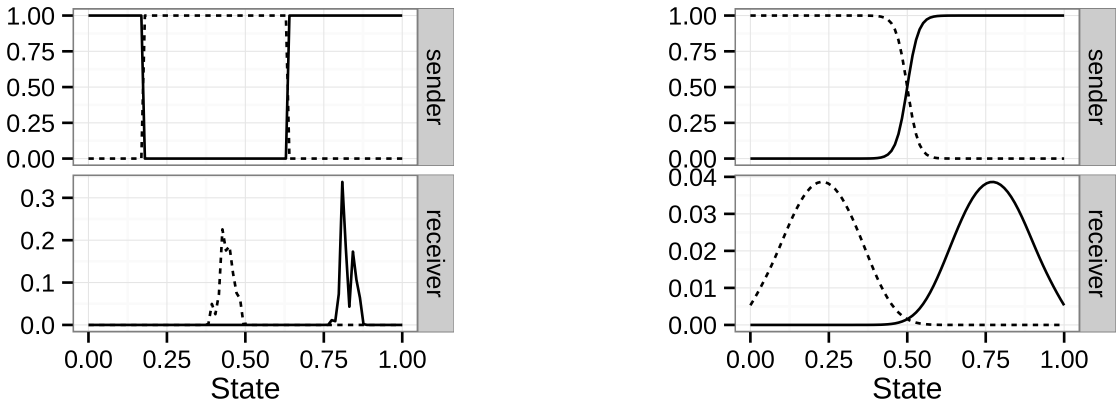

(a) $n_s = 90, \beta = 0.1, \alpha = 0$

(b) $n_s = 90, \beta = 0.1, \alpha = 0.05$

Figure 6: Example strategies at stopping time. Each line corresponds to a message and plots, for each state, the probability that the message is used.

limited-time simulations (e.g. Figure 6a), but this also holds, for some types of utility functions, for the limiting case. This was first observed by Elliott Wagner, as mentioned by O'Connor (2014b). A full analysis of the dynamics of sim-max games is beyond the scope of this paper, but we will see shortly that diffusion from confusability of states clearly prevents evolutionary paths that meander for a long time in the vicinity of non-convex strategies.

4.3 Measures of interest

To further explore our simulation results, we calculated metrics that aim to numerically capture how vague, generally well-structured, and communicatively efficient the recorded strategy pairs were. Entropy captures the amount of systematicity or regularity in signal use. Convexity captures whether a behavioural strategy would project onto a convex pure strategy. Expected utility measures the communicative efficiency of evolved strategy pairs.

Entropy. This classic information-theoretic notion captures the amount of uncertainty in a probability distribution. Roughly put, entropy of a signalling strategy captures inverse distance from a pure strategy. The usual definition of entropy applies directly to mixed strategies (see Footnote 5), but provably equivalent metrics for behavioural strategies are ready to hand:

$$E(\sigma) = -\sum_{t \in T} \sum_{m \in M} \sigma(m \mid t) \cdot \log(\sigma(m \mid t)), \qquad E(\rho) = -\sum_{m \in M} \sum_{t \in T} \rho(t \mid m) \cdot \log(\rho(t \mid m)).$$

Values obtained by these definitions have lower bound of zero and an upper bound of, respectively, $\log(|M^T|) = |T| \cdot \log(|M|)$ and $\log(|T^M|) = |M| \cdot \log(|T|)$. We work with values rescaled to lie in $[0; 1]$ for cross-comparability. The sender strategies in Figures 6a and 6b have entropy $1.19e^{-5}$ and 0.08, respectively. The receiver strategies have respective entropies 0.43 and 0.81. In general, we expected that vague languages would have higher entropy than crisp ones and that increasing imprecision would lead to increased entropy, all else being equal.

Convexity. At least for sender strategies, which develop faster than receiver strategies, it also makes sense to define a categorical measure of convexity that compensates for potential vagueness. To determine whether a sender strategy σ is convex despite possibly being vague, we look at the derived pure

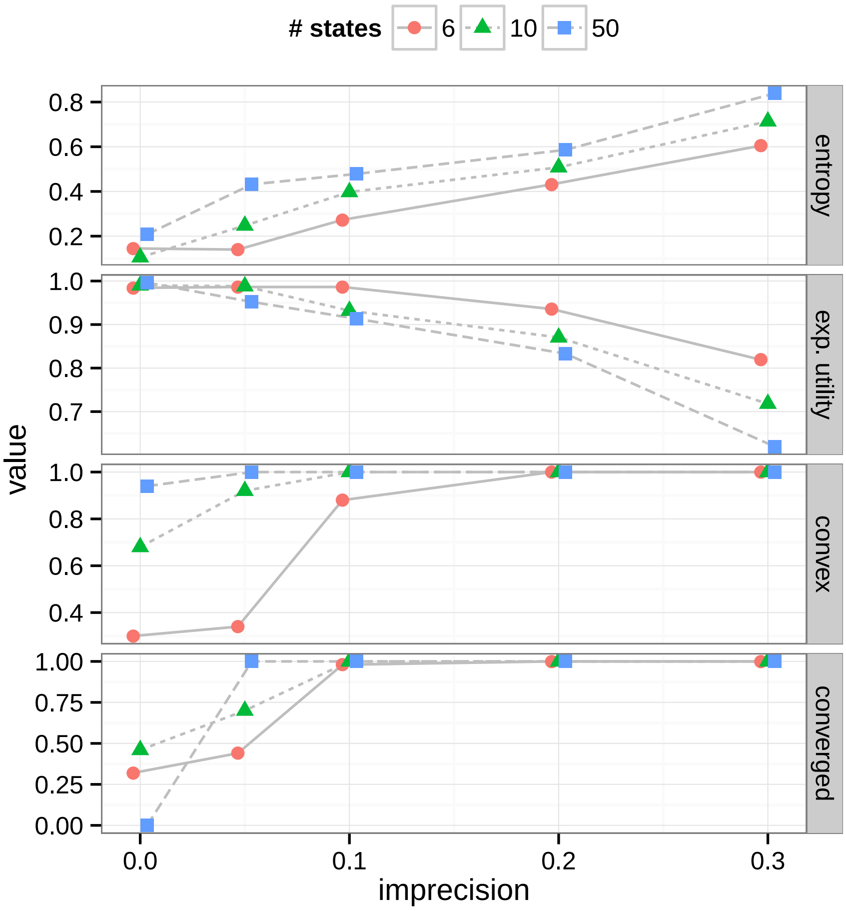

Figure 7: Means of gradient and proportions of categorical measures for $\beta = 0.1$, $n_s \in \{6, 10, 50\}$, and $\alpha \in \{0, 0.05, 0.1, 0.2, 0.3\}$. The plot shows the average of the entropies for the sender and receiver strategy.

strategy s for which $s(t) = \arg\max_{m' \in M} \sigma(t, m')$. If that s is convex, we also count σ as convex. The sender strategy in Figure 6a is not convex, while the one in Figure 6b is. If confusion of states can regularise signalling strategies, like we hypothesised, we should see more convexity with increasing imprecision all else equal.

Expected utility. We also recorded the expected utility of a strategy pair:

$$\mathrm{EU}(\sigma, \rho\,; \beta) = \sum_{t \in T} \sum_{m \in M} \sum_{t' \in T} P(t) \cdot \sigma(t, m) \cdot \rho(m, t') \cdot \mathrm{U}(t, t'\,; \beta)\,.$$

To make direct comparisons across different parameter settings, we normalise expected utility by the maximal amount of expected utility obtainable in the relevant game. The strategy pair in Figure 6a has a normalised expected utility of 0.99, the pair in Figure 6b has 0.95. Generally, vagueness and imprecision can be expected to decrease expected utility (c.f. Lipman 2009). The crucial question is whether communicative success drops unacceptably fast with moderate levels of vagueness and imprecision.

4.4 Results

Figure 7 shows plots summarising a selected part of our findings. For perspicuity, we only plot results for one level of tolerance $\beta = 0.1$, and leave out the case of $n_s = 90$. Still, every qualitative trend

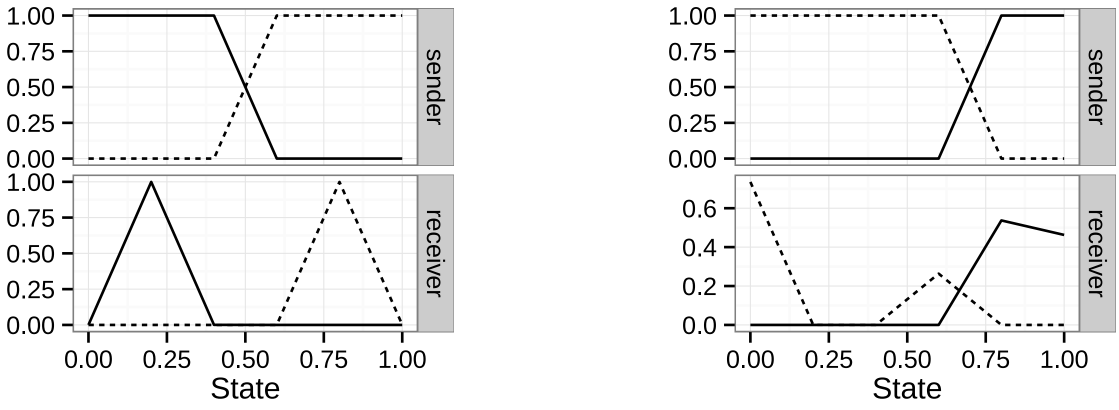

(a) $n_s = 6, \beta = 0.2, \alpha = 0$ (b) $n_s = 6, \beta = 0.2, \alpha = 0$

Figure 8: More example strategies at stopping time of our simulations.

mentioned in the following applies to the whole set of results.

As expected, increasing imprecision leads to higher entropy and lower expected utility. Importantly, however, imprecision does not necessarily lead to disastrous decline of communicative success. What is more, in line with our hypothesis that mere imprecision in imitation behaviour can lead to behaviour that looks as if agents generalise across similar stimuli, higher imprecision led to a higher number of outcomes with convex sender strategies. It also led to higher rates of convergence. In fact, sufficient imprecision always ensured convergence and convexity. It appears that perceptual imprecision leads to more vagueness, slightly less communicative efficiency, but more regular, well-behaved languages in shorter time.

Beyond promoting convexity and convergence, diffusion also has another interesting regularising effect on the evolution of signalling. There is very little variation in the recorded metrics for evolved strategies, at least for higher values of imprecision. On closer inspection, it turns out that variability in low-imprecision conditions is not only due to non-convergence or non-convexity. Figure 8 gives two more examples of strategy pairs at stopping time. Both are obtained for the same triple of parameters, both converged before the maximum number of rounds, and both have convex sender strategies. However, they are not equally efficient. In fact, the pair in Figure 8a has a normalised expected utility of 0.99 while the one in Figure 8b only has 0.89.

Interestingly, this type of variability in evolutionary outcomes can be weeded out by imprecision. To investigate this, we calculated the average distance between evolved sender strategies within each group of trials that had identical parameter values. We determined the distance between strategies σ and σ' as the average Hellinger distance between distributions $\sigma(t)$ and $\sigma'(t)$ at each choice point t:

$$\mathrm{HD}(\sigma, \sigma') = \frac{1}{|T| \cdot \sqrt{2}} \cdot \sum_{t \in T} \sqrt{\sum_{m \in M} \left(\sqrt{\sigma(t,m)} - \sqrt{\sigma'(t,m)} \right)^2} .$$

To compensate for the arbitrariness of message use, we set the distance between strategies σ and σ' to be the maximum of $\mathrm{HD}(\sigma, \sigma')$ and $\mathrm{HD}(\sigma^*, \sigma')$ where σ^* is σ with reversed message indices. An example of the 'within group distance,' i.e., the average distances between all sender strategies obtained for the same parameter values, is plotted in Figure 9a for $\beta = 0.1$ and $n_s = 10$. Despite some quantitative differences, the general trend is the same for all other parameter settings that we tested: with increasing imprecision, the resulting sender strategies were much more alike (modulo swapping of messages). This means that perceptual imprecision can speed up and unify evolutionary outcomes. It can amplify the emergence of sender strategies that are not only convex, but also regular in that they induce a vague category

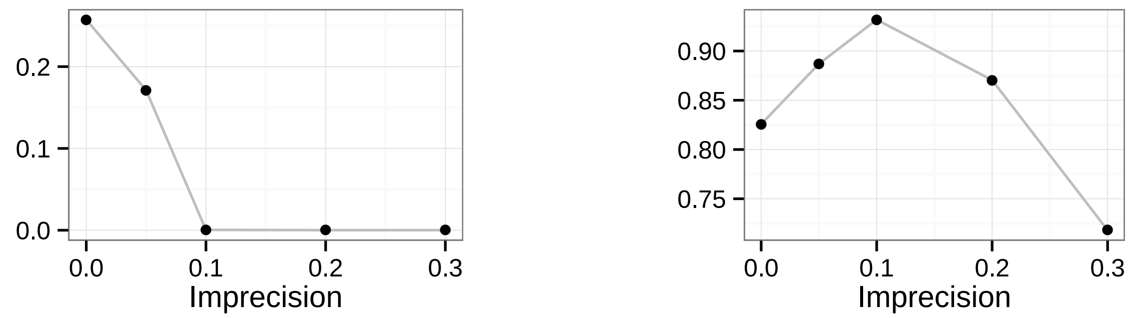

(a) Within group distance

(b) Within group expected utility

Figure 9: Within group measures for all runs with $\beta = 0.1$ and $n_s = 10$.

split exactly in the middle of the unit interval. This is then reflected in the 'within group expected utility', defined as the average expected utility that each evolved language scored when playing against an arbitrary other language obtained for the same parameter values. Figure 9b gives a representative example.

5 Discussion

Our imprecise imitation dynamic leads to by-and-large successful signalling behaviour, even in the presence of noise, that shows the hallmarks of vagueness as desired. It also gives rise to population-level behaviour that looks as if agents are generalising across similar stimuli. Here, we would like to reflect briefly on some further conceptually relevant points and compare our approach to related work.

5.1 Levels of vagueness

The imprecise imitation dynamic was introduced in Section 3 as tracing changes in the overall distribution of pure strategies in a population: $\sigma(m \mid t)$ was said to represent the probability that a randomly sampled sender would have a strategy that responds with m to t (likewise for the receiver). This is in line with the standard interpretation of the replicator equation, but we should consider its philosophical implications. Based on this picture, vagueness in signal use would seem to be characterised as a strictly population-level phenomenon since it arises in a signalling system from the inability of individual agents to fully align their (non-vague) strategies because of imprecision. This is, we believe, a plausible mechanism that can already explain the existence of vagueness in a language even if we assume that each agent commands a non-vague idiolect.

We would not, however, want to commit to the idea that vagueness does not exist at the level of individual agents. True, our derivation of the imprecise imitation dynamic assumed that agents carry and revise pure strategies. But that was an assumption of convenience, not of conviction. Moreover, even if individual agents command a non-vague pure strategy, the realisation of that pure strategy, according to our model, is bound to be vague: the same agent could signal differently in repeated exposure to the same state because of the non-deterministic nature of observation noise. We have used the term 'observation noise' here, but this could equally well be interpreted as an inseparable component of an agent's 'signalling faculty.' In this sense, then, the model might be compatible with a picture of agents who have internalised a vague signalling strategy. It would need to be seen, however, how revision of non-deterministic individual-level behaviour must be spelled out rigorously and whether the resulting population-level dynamic would be equivalent to our present proposal in all relevant respects.

5.2 Evolutionary benefits of imprecision

The inclusion of confusability of similar states has noteworthy effects on the evolving meaning of signals. It transpired from our results that imprecision can have further accelerating and, surprisingly, unifying effects on meaning evolution. The unifying property of perceptual imprecision could be considered an evolutionarily beneficial side-effect. A certain degree of imprecision can lead to higher 'within group expected utility', defined as the average expected utility that each evolved language scored when playing against an arbitrary other language obtained for the same parameter values. Figure 9b gives a representative example. The observation repeats for other parameter values: while imprecision might decrease the communicative efficiency of individual languages, it increases the conceptual coherence and communicative success between independently evolving strategies. It is as if mere confusability of states imposes a regularity constraint on evolving categories.

The phenomenon could potentially be more than just a side effect. Given the benefit of a certain amount of imprecision we observe when comparing within group expected utility, it would be interesting to study whether, under certain conditions, this group-level advantage could trump the individual-level disadvantage of a vague language, and thus actually select for a certain amount of imprecision. This could be achieved by letting the imprecision parameter be an evolving part in the dynamics as well. The idea is for now mostly speculative, but we consider it an interesting avenue for future research. There is additional motivation to consider its potential if we see it in the eyes of multilevel selection theory (Wilson and Sober 1994; O'Gorman et al. 2008).

5.3 Related work

O'Connor (2014a) makes a proposal related to ours based on a version of reinforcement learning for sim-max games, in which successful play leads to reinforcement of choice options also for states similar to the ones that actually occurred. This not only leads to vague signalling of the appropriate kind, but also speeds up learning in such a way that, especially for sim-max games with higher numbers of states, higher levels of communicative success are reached in shorter learning periods. Our results complement and extend O'Connor's. The most important differences are that (i) we obtain similar regularising effects also for cases with low numbers of states and (ii) we do not assume that agents have any kind of generalising capacity in and of themselves, even if that is only implicit in O'Connor's generalised reinforcement learning. State confusability has an effect on aggregate signalling behaviour that can be described as generalisation without generalisers: the dynamics of imprecise imitation look as if 'conclusions' about what works for one state are 'carried over' to similar states. This, however, is merely an epiphenomenon in the sense that no single agent genuinely generalises over stimuli or reasons about what a more systematic signalling strategy would be.

Franke et al. (2011) suggested a number of ways in which information-processing limitations could lead to vague strategies. The model that is most clearly related to the present approach uses the notion of a quantal response, also known as a soft-max response function (e.g. Luce 1959; McFadden 1976; Goeree et al. 2008). The main difference between this and our present approach is in where stochastic noise is assumed to reside. In case of a quantal response dynamic, it resides in the computation of expected utilities; in case of imprecise imitation, it resides in perception and realisation of similar states. There are cases, then, where evolving signalling behaviour, as predicted by these two approaches, is quite different. Intuitively speaking, for a case with two messages, the further we venture away from a prototypical interpretation of either message, the less discriminative a signalling strategy would be when the source of 'trembles' is the computation of expected utilities: to wit, since both 'tall' and 'short' are almost equally bad descriptions for a giant, quantal response dynamics predict that senders would be almost indifferent. Sender behaviour that evolves under confusion of states does not have this puzzling property, because a giant would not likely be confused for a dwarf.

6 Conclusion

We set out to meet a technical challenge posed by Lipman's problem (2009): is there a conceptually sound and mathematically coherent formal model that shows how vague language can evolve under selective pressure for efficient communication if agents tend to confuse similar stimuli? To address this, we derived a generalisation of the replicator dynamic from an agent-level process of imprecise imitation. The resulting population-level dynamic produced signalling behaviour that is at the same time regular and by-and-large communicatively efficient, while also showing the crucial marks of vagueness. In a sense, the model derives vagueness as a by-product of an arguably natural limitation on the discriminatory power of signalling agents. Although inability to sharply discriminate similar stimuli may lead to vagueness and bring about a (slight) decrease in communicative efficiency, there may also be an advantage, not of vagueness itself, but of its cause. Systematicity in the confusability of states (which may be a natural by-product of the perceptual system) supports 'as-if'-generalisation at the population-level without having to assume that agents themselves have any generalisation power. In this way, the presented model extends research into the effects of transmission biases on processes of meaning evolution. While most previous models have focused on inductive biases of language learners and the regularisation that these may effect (e.g. Nowak et al. 2000; Nowak et al. 2001; Kirby and Hurford 2002; Smith et al. 2003; Griffiths and Kalish 2007; Kirby et al. 2014), we have shown here that shared perceptual biases, of which such an effect was not necessarily expected, can also regularise, facilitate and accelerate the evolution of meaning conventions.

Appendix A Imprecise conditional imitation

The goal of this section is to provide technical details for Section 3. We first show, in Section A.1, how to derive the standard replicator dynamic from noise-free imitation of success. Then, in Section A.2, we derive the imprecise imitation dynamic by a parallel chain of arguments.

A.1 Deriving the replicator dynamic from imitation of success

The main idea behind imitation of success is that agents imitate the behaviour of other agents at some choice point with a probability that is proportional to the expected utility of the latter agents' choice. Since the sim-max games that we are looking at here have positive utilities upper bound by 1, we can identify the switching probability with the expected utility.

Let's consider sender strategies, as the receiver case is parallel. Call (misleadingly!) the agent who gets a chance to change behaviour 'learner' and the possibly to-be-imitated agent 'teacher.' A random learner is drawn from the population and given a chance to change behaviour at choice point t. The probability that our learner plays m is $\sigma(m \mid t)$. The learner observes what a randomly sampled teacher does at t. That would be m' with probability $\sigma(m' \mid t)$. The learner then starts using m' instead of m with probability $\mathrm{EU}(m', t, \rho)$. (Of course, m' and m could be the same; the learner could even be the teacher as well, by random sampling.)

If agents get repeated update chances for their choice points, the expected change of frequency of m-choices at t becomes:

$$\dot{\sigma}(m \mid t) = P(m' \to m, t) - P(m \to m', t)\,, \tag{A.1}$$

where $P(m' \to m, t)$ is the 'inflow' probability that agents switch from any m' to m and $P(m \to m', t)$ is the 'outflow' probability that agents switch from m to any m'. Since we are dealing with expectations in a huge population, these can be spelled out as:

$$P(m' \to m, t) = \sum_{m'} \underbrace{\sigma(m' \mid t)}_{\text{learner plays } m'} \cdot \underbrace{\sigma(m \mid t)}_{\text{teacher plays } m} \cdot \underbrace{\mathrm{EU}(m, t, \rho)}_{\text{EU teacher choice}}$$

$$P(m \to m', t) = \sum_{m'} \underbrace{\sigma(m \mid t)}_{\text{learner plays } m} \cdot \underbrace{\sigma(m' \mid t)}_{\text{teacher plays } m'} \cdot \underbrace{\mathrm{EU}(m', t, \rho)}_{\text{EU teacher choice}}$$

From this, we can simplify the expression of expected change in Equation (A.1) to:

$$\begin{aligned}\dot{\sigma}(m \mid t) &= \sigma(m \mid t) \cdot \mathrm{EU}(m, t, \rho) - \sigma(m \mid t) \cdot \sum_{m'} \sigma(m' \mid t) \cdot \mathrm{EU}(m', t, \rho) \\ &= \underbrace{\sigma(m \mid t)}_{\text{frequency of } m \text{ at } t} \Big(\underbrace{\mathrm{EU}(m, t, \rho)}_{\text{EU of } m \text{ at } t} - \underbrace{\sum_{m'} \sigma(m' \mid t) \cdot \mathrm{EU}(m', t, \rho)}_{\text{average EU at choice point } t} \Big).\end{aligned}$$

This latter formulation is the continuous-time version of the replicator dynamic. We obtain a discrete-time formulation from it by assuming that discrete update steps are infinitesimally small, so that:

$$\begin{aligned}\dot{\sigma}(m \mid t) &= \sigma'(m \mid t) - \sigma(m \mid t) \\ &= \frac{\sigma(m \mid t)\,\mathrm{EU}(m, t, \rho)}{\sum_{m'} \sigma(m' \mid t)\,\mathrm{EU}(m, t, \rho)} - \sigma(m \mid t) \\ &= \frac{\sigma(m \mid t)\,\mathrm{EU}(m, t, \rho) - \sigma(m \mid t) \sum_{m'} \sigma(m' \mid t)\,\mathrm{EU}(m, t, \rho)}{\sum_{m'} \sigma(m' \mid t)\,\mathrm{EU}(m, t, \rho)}.\end{aligned}$$

By dropping the denominator, which is constant for all m for fixed t, we obtain the above continuous-time formulation.

A.2 Imitation of success with imprecision

The above derivation of the replicator dynamic assumes that agents can discriminate choices and choice points perfectly. Let's dispense with that assumption. With an eye toward sim-max games, we will assume that states, but not messages, may be confused for one another.[8] Confusability of states will affect how agents behave, how they perceive the behaviour of others, and the expected utilities of behavioural dispositions.

To keep matters simple, let us assume that agents carry pure dispositions to act. Noise can affect the realisation of these strategies. As a sender, every agent maps states to messages: these are subjectively perceived states, and no longer necessarily also the actually occurring states. As a receiver, every agent maps messages to state interpretations: these are intended interpretations that need not always be faithfully realised. This means that behavioural strategies σ and ρ represent the average proportions of actual behavioural dispositions in the population, the realisation and observation of which can be distorted by agents' confusion of similar states.

If t_a is the actual state, let $P_o(t_o \mid t_a)$ be the probability that a given agent observes state t_o. Similarly, if a given receiver intends to select interpretation t_i, let $P_r(t_r \mid t_i)$ be the probability with which state t_r is realised. A single round of play of a sim-max game is then governed by five pieces of stochastic information, where previously there were only three (see Figures 3 and 4).

Expected utilities of choices at choice points should likewise take into account that actual states need not be observed states, and intended interpretations need not be realised interpretations. First, note that

$$P_{\bar{o}}(t_a \mid t_o) \propto P_a(t_a) P_o(t_o \mid t_a)$$

[8] It is relatively straightforward to also incorporate confusability of messages, but this is irrelevant to our present purposes.

is the probability that t_a is actual if t_o is observed by an agent. The probability that a random sender produces m when the actual state is t_a is:

$$P_\sigma(m \mid t_\mathrm{a}) = \sum_{t_\mathrm{o}} P_o(t_\mathrm{o} \mid t_\mathrm{a})\sigma(m \mid t_\mathrm{o})\,.$$

The probability that the actual state is t_a if a random sender produced m is:

$$P_{\overline{\sigma}}(t_\mathrm{a} \mid m) \propto P_a(t_\mathrm{a})P_\sigma(m \mid t_\mathrm{a})\,.$$

The probability that t_r is realised by a random receiver in response to message m is:

$$P_\rho(t_\mathrm{r} \mid m) = \sum_{t_\mathrm{i}} P_r(t_\mathrm{r} \mid t_\mathrm{i})\rho(t_\mathrm{i} \mid m)\,.$$

This lets us capture the expected utilities for observed states (sender) and intended interpretations (receiver) by taking into consideration what the likely actual states and realised interpretations will be:

$$\begin{aligned}\mathrm{EU}(m, t_\mathrm{o}, \rho) &= \sum_{t_\mathrm{a}} P_{\overline{o}}(t_\mathrm{a} \mid t_\mathrm{o}) \sum_{t_\mathrm{r}} P_\rho(t_\mathrm{r} \mid m)\mathrm{U}(t_\mathrm{a}, t_\mathrm{r})\,,\\ \mathrm{EU}(t_\mathrm{i}, m, \sigma) &= \sum_{t_\mathrm{a}} P_{\overline{\sigma}}(t_\mathrm{a} \mid m) \sum_{t_\mathrm{r}} P_r(t_\mathrm{r} \mid t_\mathrm{i})\mathrm{U}(t_\mathrm{a}, t_\mathrm{r})\,.\end{aligned}$$

If the conditional probabilities P_o and P_r are trivial, i.e., assign probability 0 to the confusability of non-identical states, above definitions reduce to the previous definitions of expected utilities. This also legitimates the overload of notation.

Presence of potential imprecision in the form of non-trivial P_o and P_r will also affect the dynamic that ensues from imitation of successes. Since imprecision works slightly differently on senders and receivers (the former confuse choice points, the latter confuse choices), we need to look separately at each case.

As before, suppose that senders receive a chance to change their behaviour independently for a given choice point. In the present case, this would be a chance to change how to respond to a subjectively perceived state t_o, which need not be the actual one. We must then consult the probability $P_o(m \mid t_\mathrm{o})$ that, given that the learner observed t_o, he will simultaneously observe a randomly sampled teacher play m. This is (with $P_{\overline{o}}$ and P_σ as defined above):

$$P_o(m \mid t_\mathrm{o}) = \sum_{t_\mathrm{a}} P_{\overline{o}}(t_\mathrm{a} \mid t_\mathrm{o})\, P_\sigma(m \mid t_\mathrm{a})\,.$$

The 'inflow' and 'outflow' probabilities $P(m' \to m, t)$ and $P(m' \to m, t)$ that a randomly sampled learner switches from any m' to m or from m to any m' in subjectively perceived state t_o are therefore:

$$\begin{aligned}P(m' \to m, t_\mathrm{o}) &= \sum_{m'} \underbrace{\sigma(m' \mid t_\mathrm{o})}_{\text{learner plays } m' \text{ at } t_\mathrm{o}} \cdot \underbrace{P_o(m \mid t_\mathrm{o})}_{\text{observe teacher play } m} \cdot \underbrace{\mathrm{EU}(m, t_\mathrm{o}, \rho)}_{\text{EU teacher choice in learner's view}}\\ P(m \to m', t_\mathrm{o}) &= \sum_{m'} \underbrace{\sigma(m \mid t_\mathrm{o})}_{\text{learner plays } m \text{ at } t_\mathrm{o}} \cdot \underbrace{P_o(m' \mid t_\mathrm{o})}_{\text{observe teacher play } m'} \cdot \underbrace{\mathrm{EU}(m', t_\mathrm{o}, \rho)}_{\text{EU teacher choice in learner's view}}\end{aligned}$$

The mean change to the proportion of m choices at state t are then:

$$\begin{aligned}\dot\sigma(m \mid t_\mathrm{o}) &= P(m' \to m, t_\mathrm{o}) - P(m \to m', t_\mathrm{o})\\ &= \sum_{m'} \sigma(m' \mid t_\mathrm{o})\, P_o(m \mid t_\mathrm{o})\, \mathrm{EU}(m, t_\mathrm{o}, \rho) - \sum_{m'} \sigma(m \mid t_\mathrm{o})\, P_o(m' \mid t_\mathrm{o})\, \mathrm{EU}(m', t_\mathrm{o}, \rho)\\ &= P_o(m \mid t_\mathrm{o})\, \mathrm{EU}(m, t_\mathrm{o}, \rho) - \sigma(m \mid t_\mathrm{o}) \sum_{m'} P_o(m' \mid t_\mathrm{o})\, \mathrm{EU}(m', t_\mathrm{o}, \rho)\,. \qquad \text{(A.2)}\end{aligned}$$

The case of the receiver is mostly analogous. Presented with an update opportunity for choice point m, a learner will observe a random teacher choose interpretation t_{o} with probability:[9]

$$P_o(t_{\mathrm{o}} \mid m) = \sum_{t_{\mathrm{r}}} P_o(t_{\mathrm{o}} \mid t_{\mathrm{r}})\, P_\rho(t_{\mathrm{r}} \mid m)\,.$$

Parallel to the sender case, this gives rise to:

$$\begin{aligned}\dot\rho(t \mid m) &= \sum_{t'} \rho(t' \mid m)\, P_o(t \mid m)\, \mathrm{EU}(t, m, \sigma) - \sum_{t'} \rho(t \mid m)\, P_o(t' \mid m)\, \mathrm{EU}(t', m, \sigma) \\ &= P_o(t \mid m)\, \mathrm{EU}(t, m, \sigma) - \rho(t \mid m) \sum_{t'} P_o(t' \mid m)\, \mathrm{EU}(t', m, \sigma) \end{aligned} \tag{A.3}$$

The continuous-time formulations in Equations (A.2) and (A.3) have elegant and practical discrete-time solutions in:

$$\sigma'(m \mid t) \propto P_o(m \mid t)\, \mathrm{EU}(m, t, \rho)\,, \qquad \rho'(t \mid m) \propto P_o(t \mid m)\, \mathrm{EU}(t, m, \sigma)\,,$$

which is the discrete-time formulation of the imprecise imitation dynamic given in Equation (3.2). To see how the discrete-time formulation gives rise to the continuous-time formulations above, let's assume that update steps are infinitesimally small, so that, for the sender case:

$$\begin{aligned}\dot\sigma(m \mid t) &= \sigma'(m \mid t) - \sigma(m \mid t) \\ &= \frac{P_o(m \mid t)\, \mathrm{EU}(m, t, \rho) - \sigma(m \mid t) \sum_{m'} P_o(m' \mid t)\, \mathrm{EU}(m', t, \rho)}{\sum_{m'} P_o(m' \mid t)\, \mathrm{EU}(m', t, \rho)}\,.\end{aligned}$$

As before, we drop the denominator, which is constant for all m for fixed t, and obtain the above continuous-time formulation.

Acknowledgements

We would like to thank Gerhard Jäger and several anonymous reviewers for their insightful comments and suggestions on earlier versions of this paper. Funding for this research was provided by the Dutch Research Council (NWO-VENI grant 275-80-004 to Michael Franke), the Institutional Strategy of the University of Tübingen (German Research Foundation ZUK 63 to Michael Franke), the Priority Program XPrag.de (German Research Foundation Schwerpunktprogramm 1727 to Michael Franke), and the Fundação para a Ciência e a Tecnologia (PhD studentship SFRH/BD/100437/2014 to José Pedro Correia).

References

Barrett, J. A. (2009). The evolution of coding in signaling games. *Theory and Decision*, 67:223–237.

Blume, A. and Board, O. (2014). Intentional vagueness. *Erkenntnis*, 79(4):855–899.

Blume, A., Kim, Y.-G., and Sobel, J. (1993). Evolutionary stability in games of communication. *Games and Economic Behavior*, 5(5):547–575.

Cook, R. S., Kay, P., and Regier, T. (2005). The world color survey database: History and use. In Cohen, H. and Lefebvre, C., editors, *Handbook of Categorization in Cognitive Science*, pages 223–242. Elsevier, Dordrecht.

[9] We assume that the noise $P_o(t_{\mathrm{o}} \mid t_{\mathrm{a}})$ that applies to the sender's observations also applies to the receiver's observations when trying to imitate other agents' strategies. We could easily make the model more complex and introduce another measure of noise for state confusability during receiver attempts to imitate. We refrain from it here, because we see no immediate theoretical gain.

Correia, J. P., Ocelák, R., and Mašek, J. (2016). Towards more realistic modeling of linguistic color categorization. manuscript.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50:1431–1451.

Cressman, R. (2003). *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge, MA.

Croft, W. (2003). *Typology and Universals*. Cambridge University Press, Cambridge, UK.

Decock, L. and Douven, I. (2012). What is graded membership? *Noûs*, pages 1–30.

Donkin, C., Rae, B., Heathcote, A., and Brown, S. D. (2015). Why is accurately labeling simple magnitudes so hard? a past, present, and future look at simple perceptual judgment. In Busemeyer, J. R., Wang, Z., Townsend, J. T., and Eidels, A., editors, *The Oxford Handbook of Computational and Mathematical Psychology*, pages 121–141. Oxford University Press.

Douven, I., Decock, L., Dietz, R., and Egré, P. (2011). Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic*, 42(1):137–160.

Foster, D. and Young, P. (1990). Stochastic evolutionary game dynamics. *Theoretial Population Biology*, 38(2):219–232.

Franke, M., Jäger, G., and van Rooij, R. (2011). Vagueness, signaling & bounded rationality. In Onoda, T., Bekki, D., and McCready, E., editors, *JSAI-isAI 2010*, pages 45–59. Springer.

Fudenberg, D. and Harris, C. (1992). Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57(2):420–441.

Fudenberg, D. and Imhof, L. A. (2006). Imitation processes with small mutations. *Journal of Economic Theory*, 131(1):251–262.

Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. MIT Press.

Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2008). Quantal response equilibrium. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, Basingstoke.

Griffiths, T. L. and Kalish, M. L. (2007). Language evolution by iterated learning with bayesian agents. *Cognitive Science*, 31:441–480.

Haspelmath, M. (2003). The geometry of grammatical meaning: Semantic maps and cross-linguistic comparison. In Tomassello, M., editor, *The new psychology of language, Vol. II*, pages 211–242. Lawrence Erlbaum, Mahwah, New Jersey.

Helbing, D. (1996). A stochastic behavioral model and a 'microscopic' foundation of evolutionary game theory. *Theory and Decision*, 40(2):149–179.

Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, Massachusetts.

Huttegger, S. M. (2007a). Evolution and the explanation of meaning. *Philosophy of Science*, 74:1–27.

Huttegger, S. M. (2007b). Evolutionary explanations of indicatives and imperatives. *Erkenntnis*, 66:409–436.

Huttegger, S. M., Skyrms, B., Smead, R., and Zollman, K. J. S. (2010). Evolutionary dynamics of lewis signaling games: Signaling systems vs. partial pooling. *Synthese*, 172(1):177–191.

de Jaegher, K. (2003). A game-theoretic rationale for vagueness. *Linguistics and Philosophy*, 26(5):637–659.

Jäger, G. (2007). The evolution of convex categories. *Linguistics and Philosophy*, 30(5):551–564.

Jäger, G. (2010). Natural color categories are convex sets. In Aloni, M., Bastiaanse, H., Jager, T., and Schulz, K., editors, *Logic, Language and Meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18, 2009, Revised Selected Papers*, pages 11–20. Springer Berlin Heidelberg, Berlin, Heidelberg.

Jäger, G., Metzger, L. P., and Riedel, F. (2011). Voronoi languages: Equilibria in cheap-talk games with high-dimensional types and few signals. *Games and Economic Behavior*, 73(2):517–537.

Jäger, G. and van Rooij, R. (2007). Language structure: Psychological and social constraints. *Synthese*, 159(1):99–130.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Kay, P., Berlin, B., Maffi, L., Merrifield, W. R., and Cook, R. (2009). *The World Color Survey*. Center for the Study of Language and Information, Stanford.

Kirby, S., Griffith, T., and Smith, K. (2014). Iterated learning and the evolution of language. *Current Opinion in Neurobiology*, 28:108–114.

Kirby, S. and Hurford, J. R. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In Cangelosi, A. and Parisi, D., editors, *Simulating the Evolution of Language*, pages 121–148. Springer.

Levinson, S., Meira, S., and The Language and Cognition Group (2003). 'natural concepts' in the spatial topological domain - adpositional meanings in crosslinguistic perspective: An exercise in semantic typology. *Language*, 79(3):485–516.

Lewis, D. (1969). *Convention. A Philosophical Study*. Harvard University Press, Cambridge, MA.

Lipman, B. L. (2009). Why is language vague? Manuscript, Boston University.

Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley, New York.

Luce, R. D. (1963). Detection and Recognition. In *Handbook of Mathematical Psychology*, pages 1–103. John Wiley & Sons.

McFadden, D. (1976). Quantal choice analysis: A survey. *Annals of Economic and Social Measurement*, 5(4):363–390.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115(1):39–57.

Nowak, M. A., Komarova, N. L., and Niyogi, P. (2001). Evolution of universal grammar. *Science*, 291(5501):114–118.

Nowak, M. A., Plotkin, J. B., and Jansen, V. A. A. (2000). The evolution of syntactic communication. *Nature*, 404(30).

O'Connor, C. (2014a). The evolution of vagueness. *Erkenntnis*, 79(4):707–727.

O'Connor, C. (2014b). Evolving perceptual categories. *Philosophy of Science*, 81(5):110–121.

O'Gorman, R., Sheldon, K. M., and Wilson, D. S. (2008). For the good of the group? Exploring group-level evolutionary adaptations using multilevel selection theory. *Group Dynamics: Theory, Research, and Practice*, 12(1):17–26.

Pawlowitsch, C. (2008). Why evolution does not always lead to an optimal signaling system. *Games and Economic Behavior*, 63(1):203–226.

Regier, T., Kay, P., and Khetarpal, N. (2007). Color naming reflects optimal partitions of color space. *Proceedings of the National Academy of Sciences*, 104:1436–1441.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, MA.

Schlag, K. H. (1998). Why imitate, and if so, how? *Journal of Economic Theory*, 78(1):130–156.

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

Smith, K., Kirby, S., and Brighton, H. (2003). Iterated learning: A framework for the emergence of language. *Artificial Life*, 9:371–386.

Spence, A. M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87:355–374.

Taylor, P. D. and Jonker, L. B. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Bioscience*, 40(1–2):145–156.

van Deemter, K. (2009). Utility and language generation: The case of vagueness. *Journal of Philosophical Logic*, 38(6):607–632.

Wärneryd, K. (1993). Cheap talk, coordination, and evolutionary stability. *Games and Economic Behavior*, 5(4):532–546.

Wilson, D. S. and Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and Brain Sciences*, 17:585–608.

Young, P. (1993). The evolution of conventions. *Econometrica*, 61(1):57–84.