An adaptationist criterion for signal meaning

Michael Franke

Institute for Logic, Language and Computation
Amsterdam, The Netherlands

Abstract

To think that use constitutes meaning is an appealing idea, but also notoriously difficult to make precise. After reflecting on the notion of use-constituted meaning that signaling theory, in the vein of Lewis (1969) and Skyrms (2010), gives us, I suggest a notion of signal meaning that draws on adaptationist theories of biological communication. Such a criterion has interesting consequences, such as that just because language use is vague does not necessarily entail that language is vague.

"Wenn sich alles so verhält als hätte ein Zeichen Bedeutung, dann hat es auch Bedeutung." (Wittgenstein, TLP 3.328)

Communication takes place between two separable entities. These need not necessarily be two different individuals, for we want to allow for communication to be possible between, say, neurons in the brain. Communicating entities also need not be of the same species, for we want to allow for communication between, say, a flower and a bee. Stereotypically, we might think of communication as taking place between human speakers and listeners, but that is just a special case. Opinions diverge on how special it is, and why it should be so. In any case, communication often has a certain directedness from one entity to the other. The former will be called sender, the latter receiver.

We expect communication to arise when the utilities (think: happiness, longevity, replication or procreation probability, biological fitness, ...) of sender and receiver are dependent in appropriate ways. A comprehensive, versatile and highly expandable model of such a situation is found in signaling games as introduced by David Lewis (1969). The sender shows or performs some observable pattern, henceforth called a message m. This may possibly be conditional on some state of nature t. The receiver shows some kind of reaction a to the observable messages m. (The receiver may have his own informational resources, but in the simplest case that we assume here the receiver is ignorant of which state of nature obtains.) The triple t, m and a is an outcome of the game. How valuable an outcome is to sender and receiver is given by a utility function that maps outcomes onto reals (the higher the number, the better).

In the simplest, non-trivial case there are two states, two messages and two acts. There is a one-to-one mapping between states and acts in the sense that sender and receiver both get a high payoff when the receiver chooses the act which is right given the state and a low one otherwise. We assume that a_1 is appropriate for t_1, and a_2 for t_2. There are four non-probabilistic patterns of sender behavior:

(1) a. $t_1 \mapsto m_a$, $t_2 \mapsto m_a$ b. $t_1 \mapsto m_a$, $t_2 \mapsto m_b$ c. $t_1 \mapsto m_b$, $t_2 \mapsto m_a$ d. $t_1 \mapsto m_b$, $t_2 \mapsto m_b$

Similarly, there are four principled types of non-probabilistic receiver behavior:

(2) a. $m_1 \mapsto a_1$, $m_2 \mapsto a_1$ b. $m_1 \mapsto a_1$, $m_2 \mapsto a_2$ c. $m_1 \mapsto a_2$, $m_2 \mapsto a_1$ d. $m_1 \mapsto a_2$, $m_2 \mapsto a_2$

Two combinations of sender and receiver behavior are special; Lewis called them signaling systems:

(3) 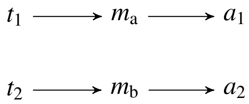

(4) 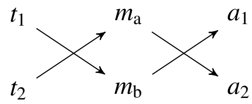

Lewis suggested that messages can be said to have meaning in signaling systems. The indicative meaning of a messages is the set of all states in which that message is sent. The imperative meaning of a message is the set of acts that it triggers. For example, the indicative meaning of m_1 in (3) is t_1 and its imperative meaning is a_1. This verdict seems *prima facie* plausible. As Lewis put it:

> I have now described the character of a case of signaling without mentioning the meaning of the signals (. . .). But nothing important seems to have been left unsaid, so what has been said must somehow imply that the signals have their meanings.
>
> (Lewis, 1969, pp. 124–125)

But Lewis did not just shoot from the hip like so. He also showed that signaling behavior in signaling systems can be justified by practical reasoning in a way that fits the characterization of non-natural meaning given by Paul Grice (1957).

Non-natural meaning. Grice introduced an important distinction that still reverberates in contemporary thought. On the one hand, a systematic correlation between events, if known, can be used to acquire information that is not directly observed. Smoke may mean fire, certain red spots may mean measles, and so on. These are examples of what Grice called natural meaning. On the other hand, non-natural meaning is of the type instantiated prototypically in human communication. A certain gesture at a party may mean that Smith thinks that the event falls short of his expectations. The hallmark of non-natural meaning, according to Grice, is that the speaker's intention to induce a certain response in his listeners (most palpably: a belief that such-and-such) be recognized and that this recognition be effective in triggering the intended response:

> "*A* meant$_{\text{NN}}$ something by *x*" is (roughly) equivalent to "*A* intended the utterance of *x* to produce some effect in an audience by means of the recognition of this intention."
>
> (Grice, 1957, p. 385)

Lewis argued that "meaning$_{\text{NN}}$ is a consequence of conventional signaling" (Lewis, 1969, p. 154). The reason, if couched in more modern terms, is the following. Signaling systems are Nash equilibria. Nash equilibria are rationalizable, i.e., can be supported by an infinite chain of what each player does and believes the other player does and believes, so that each combination of actions and beliefs is rational given the utilities of the game. At each step, the choices, preferences and beliefs of each agent also legitimize the ascription of appropriate intentions to the agents. This scaffolding of choices, desires, expectations and intentions need not be actualized, but should merely be viewed as a potential justification of the practical reasoning that may lead to behavior in conformity with a signaling system. But most importantly, the justification of the signaling system by practical reasoning does fulfill Grice's definition of non-natural meaning, as Lewis showed.

Informational content. This reconstruction of Gricean non-natural meaning is viable only when ascriptions of potential practical reasoning to the signaling agents is sensible. Animals presumably do not qualify (see Lewis' remarks about horses, (Lewis, 1969, p. 130)), flowers and neurons most certainly do not. For that reason, Brian Skyrms has arranged a beautiful marriage of Lewis' signaling models with information theoretic notions to assess, if not the meaning of signals, then at least the information contained in the use of signals (Skyrms, 2010).

Each states of nature t occurs with an objective probability $\Pr(t)$. In each state t, the average sender behavior can also be expressed probabilistically, so that $\Pr(m \mid t)$ is the probability with which (on average) message m is sent in state t (by senders in the population). This is enough to make the following assessment: how much information (about which state is actual) would we lose if we only looked at the prior probability of states and not also at the signal that the sender sent? More concretely, we turn to the likelihood of states given a message, computed by Bayesian update:

$$\Pr(t \mid m) = \frac{\Pr(t)\Pr(m \mid t)}{\sum_{t'} \Pr(t')\Pr(m \mid t')} = \frac{\text{probability that } t \text{ occurs and } m \text{ is sent}}{\text{probability that } m \text{ is sent at all}},$$

and compare this to the prior probability of $\Pr(t)$. The amount of information lost when looking at the prior $\Pr(\cdot)$ instead of the likelihood $\Pr(\cdot \mid m)$ is given by the Kullback-Leibler divergence and yields Skyrms' notion of *informational content in a message about the states*:

$$\mathrm{I}(m) = \sum_{t \in T} P(t \mid m) \log \frac{P(t \mid m)}{P(t)}.$$

But we can also look specifically at the information in the use of m about a particular state t. This is Skyrms' informational content vector:

$$\mathrm{ICV}(m) = \left\langle \log \frac{P(t \mid m)}{P(t)} \mid t \in T \right\rangle.$$

Here, we track for each state how much its likelihood given m differs from its prior. If desired, the propositional content of a message is then just another short step away, obtained by abstracting from the quantitive information:

$$\mathrm{Prop}(m) = \left\{ t \in T \mid \log \frac{P(t \mid m)}{P(t)} > -\infty \right\}.$$

For a signaling system of a Lewisean signaling game (exactly one action right for each state; equal number of states, messages and acts), the propositional content of a message as defined by Skyrms is equivalent to the indicative meaning of that message as defined by Lewis. But Skyrms' information-theoretic purism has a number of advantages over Lewis' reconstruction of non-natural meaning. The former does not require strategies to be in equilibrium or evolutionarily stable; it applies to any sender or receiver strategy, no matter whether they have co-evolved to be mutually optimal or just arrived from planet Random Intelligent Design. What is more, the information-theoretic approach applies just as well to signaling among neurons, flowers, bees and horses as to signaling among human beings.

Misinformation & deception. An information-theoretic assessment of signal content groups under Grice's notion of natural meaning. Given a probabilistic sender strategy, we read the signaling game like a probabilistic flow chart or a Bayesian causal network. The event that a given signal occurred carries information about latent events, namely the actual state. The same logic is applied to justify saying that fire means smoke and that those spots mean measles.

Skyrms is content with this; he does not aim for defining fancier notions like communication or meaning. We should bear this in mind. But even so errors and cases of improper signal use deserve special attention.

Consider errors first. If senders make mistakes, even with low probability, this will have an effect on informational content. If error rates are uniform, each message is used in each state and so the propositional content of all messages is simply tautologous, so to speak. Of course, informational content vectors contain quantitative information as well. It might be possible to subtract the impact of systematic errors and focus on the use-induced informational content, but that is not what informational content in Skyrms' sense is after.

Consider what we might call improper signal use. Cases of putative deception among animals are of great interest to biologists. Skyrms discusses the case of Photinus and Photuris, two species of fireflies. Female Photuris, the *femme fatale* firefly, mimics the mating flashes of female Photinus, thereby attracting male Photinus fireflies which they eat. A simpler case to the same effect can be constructed as a signaling game that adds an additional (unlikely) state to a Lewisean signaling game, like so:

	$P(t)$	a_1	a_2
t_1	2/5	1,1	0,0
t_2	2/5	0,0	1,1
t_3	1/5	1,0	0,1

In states t_1 and t_2, sender and receiver agree that a_1 and a_2 are best respectively. State t_3 marks a case of diverging interests, where sender and receiver prefer opposite acts. Here is an equilibrium of this game, roughly similar to the signaling behavior of the fireflies:

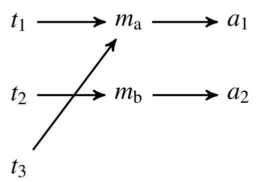

It is appealing to think that when sending m_a in t_3 the sender free-rides on the "proper meaning" of m_a, which "properly means" t_1 and "properly triggers" a_2. The case of the fireflies abusing a mating call as a food delivery call is similar in this sense.

Still, the informational content of m_a is not meant to reflect this, as it deals with plain objective statistical correlation. To be able to talk about improper uses nonetheless, Skyrms introduces additional notions on top of informational content. If the use of a message m increases the probability of a state t that is not actual ($\Pr(t \mid m) > \Pr(t)$), or if it decreases the probability of the actual state t^* ($\Pr(t^* \mid m) < \Pr(t^*)$), message m contains *misinformation*. A misinformative signal that is systematically sent to the benefit of the sender and the detriment of the receiver is *deceptive* (Skyrms, 2010, Chapter 6).

When Photuris females mimic the mating calls of Photinus females to order dinner, that is deception in this sense. The message m_a in the toy model above is deceptive as well. If the actual state is t_1 it raises the probability of the non-actual state t_3. So, it carries misinformation. In equilibrium, it is used systematically in t_3 and this clearly is to the benefit of the sender, who wishes to induce a_1, and to the detriment of the receiver, who would rather want to do a_2 in t_3.

Notice that this definition of deception parts with the austere approach of looking just at statistical correlations. To judge whether behavior is to the sender's benefit and the receiver's detriment we need to consider the wider context of signaling, in particular information in the agents' utilities. Below I will go down that road a few more steps by trying to define "proper meaning" directly in terms of behavior and the wider signaling context.

Under-informativity. By Skyrms' definition, signals that are under-informative, in an intuitive sense, may also be deceptive. A case in point is discussed by Skyrms himself (Skyrms, 2010, pp. 80–82):

	$P(t)$	a_1	a_2	a_3
t_1	1/3	2,10	0,0	10,8
t_2	1/3	0,0	2,10	10,8
t_3	1/3	0,0	10,10	0,0

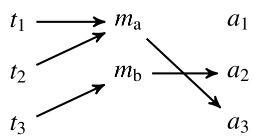

Notably, interests of sender and receiver are largely aligned, except for which act is best in the first two states. There are fewer signals than states, so that, in the given equilibrium, the sender uses the same message m_a in the first two states. The best response of the receiver is to go for a_3. This is deception, according to Skyrms' definition. No matter whether state t_1 or t_2 is actual, message m_a raises the probability of the respective other non-actual state. Hence, it is misinformation. It is also deceptive, because in the given equilibrium the sender sends it systematically to his own benefit and to the receiver's detriment.

Maintenance. Some feel that there is a difference between the two examples of deceptive signal use that we looked at, and that it is valuable to make a distinction. Peter Godfrey-Smith seems to share this sentiment and he tries to identify the source of this intuition thus:

> There is a difference between the *maintaining* and the *non-maintaining* uses of the signal. Some uses contribute to stabilization of the sender-receiver configuration and some, if more common, would undermine it. (Godfrey-Smith, 2011, p. 1295)

The mimicking of the Photinus mating signals by Photuris is non-maintaining in an intuitive sense: increased uses of deceptive signaling of this type would soon lead to the signal being ignored. This is clear to see in the toy signaling game given in this context. If the probability of t_3 would increase beyond that of t_1, the given equilibrium would break as the receiver would rather choose act a_2 in response to m_a.

But the under-informative signal m_a in the other example we looked at, though deceptive in Skyrms' sense, is a maintaining use. That, at least, is what Godfrey-Smith (2011) intuits. But it is not clear, then, what a maintaining use is. For if we apply the same reasoning as before and imagine that the frequency of t_1 increases, eventually the receiver would also abandon his part of the reference equilibrium and choose a_1 in response to m_a. Consequently, it is not obvious what it means exactly for a signal's use to be maintaining or not.

Biological meaning. There is a different way of demarcation. I suggest that the difference between the deceptive signals in the previous examples could be seen as a difference in their meaning: the latter under-informative use gives us a disjunctive meaning, but not the latter genuinely deceptive case. What notion of meaning warrants this conclusion?

A terminological caveat. It makes no sense to ask whether a flower meant to inform a bee of its presence by the display of its colors. Maybe we can say that the flower's colors mean to the bee that the flower is over there. On any account, the use of the term "meaning" in biological contexts is shaky. This said, let's shake it some more.

Biologists debate whether meaningfulness in the organic world should be accounted for in terms of information transfer or behavioral adaptation. The latter position is in keeping with Lewis' and Skyrms' tradition, and so I stick with it as well. The goal is to derive meaningfulness and information transfer from evolved behavioral patterns. But adaptationist definitions of biological meaningfulness are not all equally suitable. A recent attempt that might fit is due to Thomas Scott-Phillips (2008):

> A signal is any act or structure that (i) affects the behaviour of other organisms; (ii) evolved because of those effects; and (iii) which is effective because the effect (the response) has evolved to be affected by the act of structure. (Scott-Phillips, 2008, p. 388)

We need to bear in mind an unfortunate but harmless overload of terminology here. “Signal” in the sense used here is a type-level term for something that carries biological meaning. The above definition then comes down to requiring that for something to carry biological meaning it should be part of an exchange between two parties, and be supported by a history of evolutionary selection both on the sender and receiver side, because of its effects on the receiving party.

This does not yet define what the content of a biological signal is. It is hard to work this out in full generality; it is much easier if we stick to the signaling game framework. Following the adaptationist’s view of animal communication, we could then try to define a message’s meaning content as follows:

(5) Let $\tilde{s}$ and $\tilde{r}$ be the average sender and receiver behavior in a signaling game. Message m means state t iff

- (i) m is used in t;
- (ii) sending m in t evolved because of $\tilde{r}$; and
- (iii) $\tilde{r}$ evolved because of m being sent in t.

The crucial notion in this definitional attempt is that of a “being evolved for a reason.” I have nothing better to offer right now. But the vagueness in this locution is not even entirely unwelcome, and a number of important implications are clear enough. Firstly, the definition can in principle apply to signaling among neurons, flowers, bees, horses and humans, because no indecent mentalistic abilities are required of signaling agents for their signaling behavior to be constitutive of meaning in this sense. Secondly, unlike Lewis’ notion of meaning and Skyrms’ informational content vectors, my take on meaning in signaling games requires more than just a momentary snapshot of sender and receiver behavior. It does not require signaling behavior in equilibrium, but it requires to look at the evolutionary history of strategies, their evolutionary *raison d’être*. This it shares with the approach of Harms (2010).

By this definition, the mating signals of our fireflies most plausibly means “here is a Photunis female ready to mate” and not “here is a Photunis female ready to mate or a Photuris female ready to eat.” Check this in the corresponding toy signaling game example. Let us assume that we are at this equilibrium, but that we know, or may reconstruct with sufficient certainty, the most recent evolutionary path that led to the equilibrium behavior. We then consider a probabilistic sender and receiver behavior $\tilde{s}$ and $\tilde{r}$ as it was shortly before evolutionary pressure drove the signaling behavior to the hypothesized equilibrium state. This would look something like this:

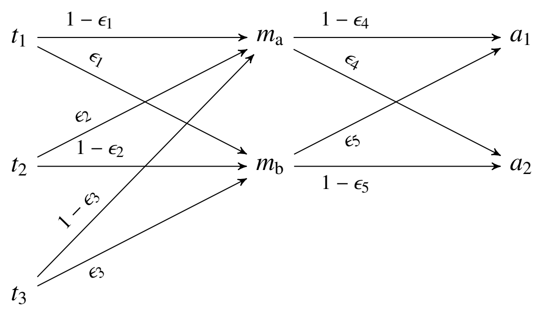

Message m_a plausibly means t_1 because (i) given the receiver’s pre-equilibrium behavior $\tilde{r}$, there is selective pressure for the sender to use m_a in t_1, and (ii) the sender’s use of m_a in t_1 in $\tilde{s}$ provides evolutionary pressure for the receiver strategy to converge to the equilibrium state. The latter reasoning does not go through for the association of m_a and t_3. That the speaker uses m_a in t_3 in the pre-equilibrium state is not a reason why the receiver’s equilibrium behavior evolved. This should hold true no matter how the vague notion of “being a reason for evolution” is spelled out precisely.

The case of under-informative signals is different. In this case, m_a means both t_1 and t_2 because, speaking with respect to an obvious evolutionarily ancestral state close to the reference equilibrium, there is evolutionary pressure for sending m_a in either state, given the receiver strategy at that time. Also, the hypothesized ancestral receiver strategy would be driven towards the equilibrium strategy because of the dominant but roughly equal use of m_a in these states.

This reasoning may be closely related to Godfrey-Smith's (2011) intuitions about maintaining and non-maintaining uses, but I suggest that my wording makes things a little more clear, and highlights the obvious connection to adaptationist accounts of biological meaningfulness.

Sim-max games. My proposed notion for capturing the meaning that the use of a signal bestows on it has interesting consequences for the emergence of meaning in what has been called similarity-maximizing games or *sim-max games*, for short. Sim-max games are a generalization of Lewis' signaling games to account for the formation of linguistic categories of contiguous (or continuous) perceptual input (e.g. Jäger and van Rooij, 2007). These games have recently attracted attention because they enable modeling the conditions under which vague meaning can evolve (e.g. Franke, Jäger, and van Rooij, 2011; O'Connor, 2013; Correia, 2013). What is interesting is that, unlike what has been suggested previously, the meaning that evolves in the suggested models would not be vague at all under the present notion of behavioral meaningfulness.

Sim-max games have many states and few messages. For simplicity of parlor, we may assume that the receiver's acts are identical to the states. Then, if the actual state is t and the receiver selects t' as his best guess of what the actual state is, the payoff for sender and receiver is no longer a matter of black-or-white but of shades of gray. In particular, it will be proportional to how similar t' is to t. To give an example, if the speaker wants to communicate how tall Smith is, he could say "Smith is tall" and the receiver guesses Smith's body length. If Smith is actually 181 cm, but the receiver guesses 182, that's better than guessing 185, but worse than guessing 181, which is the absolute maximum (assuming we are only dealing with unrounded centimeter measures).

It will help the subsequent discussion to have a concrete sim-max game at hand. Let there be six states and two messages. With flat priors of states and similarity-based utilities as specified below, we are interested in the following game and its equilibrium:

	$P(t)$	a_1	a_2	a_3	a_4	a_5	a_6
t_1	$1/6$	5	4	3	2	1	0
t_2	$1/6$	4	5	4	3	2	1
t_3	$1/6$	3	4	5	4	3	2
t_4	$1/6$	2	3	4	5	4	3
t_5	$1/6$	1	2	3	4	5	4
t_6	$1/6$	0	1	2	3	4	5

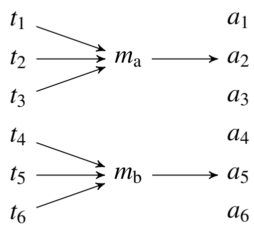

Intuitively, this looks like a case of category formation mediated by language use. We may be tempted to say that the category extension induced by message m_a is $\{t_1, t_2, t_3\}$ and its prototype is t_2. Based on this intuition, sim-max games have been heralded because of their ability to explain why natural categories are convex and where their prototypes are to be located: it can be shown that the evolutionary stable states of sim-max games are such that (i) the indicative meaning of a signal, in Lewis' sense, yields a (quasi-)partition of the state space whose cells are convex regions based on the underlying similarity metric, and (ii) the imperative meaning, in Lewis' sense, specifies the Bayesian estimator of the corresponding category (given the priors and the similarity metric) (Jäger, 2007; Jäger, Metzger, and Riedel, 2011).

Vagueness in signaling behavior. The use-based meaning of signals in evolutionary stable states of sim-max games is always sharp and crisp (barring aberrant cases that can occur due to finiteness of the state space). But exactly those words and concepts in the human inventory to which sim-max games would most plausibly apply are vague: think of color terms, gradient natural properties like extension in space, weight, temperature, duration in time, intensity of tastes etc. To stick with the recent example, the sentence "Smith is tall" would not normally induce a belief in the listener that Smith is exactly n cm, let alone that very same belief in all listeners' who understand the meaning of that sentence, and the use of speakers of English is not uniform in calling everybody of at least m cm "tall."

There are many potential explanations why the categories and prototypes in sim-max games may be roughly convex, well-behaved and conducive of successful information flow, while also showing the hallmarks of fuzziness and vagueness. Some suggest individual differences or individual errors in language use (Franke, Jäger, and van Rooij, 2011; Correia, 2013), some suggest putative (but eventually useful) imperfections in learning (O'Connor, 2013). Irrespective of the exact origin of the noisy element, essentially the ensuing sender and receiver behavior would gravitate to something like the following:

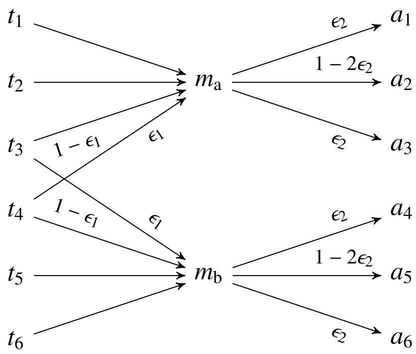

In intuitive terms, the categories induced by message use will show a gradual and smooth transition (if there are enough states to show this) and the prototypes will be extended regions with more and more weight in their center.

Noisy use, crisp meaning. What do these models tell us about the use-induced meanings of signals? Do these models lend credence to the idea that human language is vague? Franke, Jäger, and van Rooij (2011) suggest so, when they extend Lewis' notion of meaning to cover probabilistic information about use. For example, the probabilistic descriptive meaning of a message m is defined as a function from sender behavior to a probability distribution over states:

$$F_{\tilde{s}}(m, t) = \frac{\text{probability that } m \text{ is used in } t \text{ under } \tilde{s}}{\text{probability that } m \text{ is used at all under } \tilde{s}}\,.$$

By this criterion the meaning of m_a in the previous example would be a probability vector, namely $\langle 1/3, 1/3, 1-\epsilon/3, \epsilon/3, 0, 0\rangle$. This view licenses the rough-and-ready conclusion that m_a has a vague meaning because its denotation gradually (within the limits of having only very few states in this example) declines as we travers the state space along a gradient of similarity.

But what if we take into account the reasons why a particular signaling behavior evolved? Under the definition of use-induced meaning suggested above, the meaning of m_a is not vague. Its content is $\{t_1, t_2, t_3\}$ and that's it. Message m_a does not mean t_4, although it is used in that state as well. This is because we should plausibly say that the ϵ-use of m_a in t_4 is not evolved (or maintained) by evolutionary pressure given the receiver's (pre-equilibrium) behavior to respond with mostly a_1 to m_a.

More generally, any putative meaning component, probabilistic or otherwise, that is supported by slips, errors or noise alone will not count as part of a signal's meaning by a functional meaning definition, like the one proposed here, that asks for the wider evolutionary purpose of the relevant behavior. This picture points in the direction of epistemicism (Williamson, 1994), or positions that maintain that language use is noisy, and (hence) knowledge of language imperfect for every language user, but still meaning conventions as abstractions over noisy and uncertain use are not vague:

> [L]anguages themselves are free of vagueness but (...) the linguistic conventions of a population, or the linguistic habits of a person, select not a point but a fuzzy region in the space of precise languages. (Lewis, 1970, p. 64)

Many will not like this implication. Some will deny the usefulness of abstracting over language use like that. Others will say that there may be an evolutionary reason why language is vague after all. I cannot rule out that there is a purpose to vagueness for which it evolved, but, at present, I find it quite improbable that such a purpose exists that also pays the double bill of benefiting sender and receiver.

Conclusion. Signaling games are widely employed as models for the contextualized use of language or any other type of potentially communicative behavior or trait. The insinuation is that signaling behavior is constitutive of signal meaning. But only little effort is placed in specifying more concretely what exactly the meaning of signals is that their use should bestow on them. I have tried to add to this discussion by situating, formulating and exploring (albeit ever so briefly) an adaptationist account of signal meaning. My suggestion is that the identification of a signal's meaning requires inspecting the reasons why the signaling behavior evolved, on the part of both sender and receiver.

References

Correia, José Pedro (2013). "The Bivalent Trap: Vagueness, Theories of Meaning, and Identity". MA thesis. Universiteit van Amsterdam.

Franke, Michael, Gerhard Jäger, and Robert van Rooij (2011). "Vagueness, Signaling & Bounded Rationality". In: *JSAI-isAI 2010*. Ed. by T. Onoda, D. Bekki, and E. McCready. Springer, pp. 45–59.

Godfrey-Smith, Peter (2011). "Signals: Evolution, Learning, and Information". In: *Mind* 120.480, pp. 1288–1297.

Grice, Paul Herbert (1957). "Meaning". In: *Philosophical Review* 66.3, pp. 213–223.

Harms, William F. (2010). "Determining truth conditions in signaling games". In: *Philosophical Studies* 147, pp. 23–35.

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 (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 (2007). "Language Structure: Psychological and Social Constraints". In: *Synthese* 159.1, pp. 99–130.

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

— (1970). "General Semantics". In: *Synthese* 22.1–2, pp. 18–67.

O'Connor, Cailin (2013). "The Evolution of Vagueness". In: *Erkenntnis*. To appear.

Scott-Phillips, Thomas C. (2008). "Defining biological communication". In: *Journal of Evolutionary Biology* 21, pp. 387–395.

Skyrms, Brian (2010). *Signals*. Oxford: Oxford University Press.

Williamson, Timothy (1994). *Vagueness*. London and New York: Routledge.