Animal metasemantics

Luca Gasparri | luca.gasparri@cnrs.fr

AAM, July 2025. Published in *Biology & Philosophy*, https://doi.org/10.1007/s10539-025-09993-1

Abstract

Multiple scholars in animal communication research have proposed that models of animal signals should appropriate theoretical terminology and concepts from natural language semantics. However, the application of semantic descriptors to nonhuman communication is a delicate undertaking. The goal of this paper is to consider whether there are any basic methodological principles that can help us determine whether the appropriation of a semantic descriptor by an animal model is justified or unwarranted. I begin by drawing a parallel between types of theories of meaning and types of theories of signal content, and propose to categorize the metatheoretical issues surrounding the application of semantic descriptors to animal signals under the banner of an animal metasemantics. I then develop three simple conditions on the operation: a condition of epistemic fertility, one of intensional depth, and one of grammatical congruity. Finally, I showcase how these heuristics can assist theory production and help identify strengths and opportunities for improvement in emerging projects in animal linguistics.

Keywords: animal communication; animal signals; meaning; semantics; metasemantics; animal linguistics.

Acknowledgments

A preliminary version of this material was presented in a PhD seminar on nonhuman communication held between 2023 and 2024 at the University of Lille, co-organized by Sylvain Billiard, Nina Joffard, Cédric Patin, Gerhard Schaden, and myself. Thanks to my co-organizers, to the students, and to the invited speakers (Pius W. Akumbu, Mélissa Berthet, Bart de Boer, Virginie Cuvillier, Piera Filippi, Albain Lemasson, Violaine Llaurens, Magali Proffit, and Fanny Rybak) for the stimulating discussions we had during the seminar. Thanks to Mélissa Berthet, Piera Filippi, Hanjo Glock, and Stuart Watson for comments on an earlier draft, and to the anonymous reviewers for very helpful feedback on the submitted manuscript. All errors and inaccuracies are mine.

1. Introduction

Every attempt to reflect, whether informally or scientifically, on the parallels and non-parallels between human language and animal communication inevitably runs into the following question: can animals *mean* things with their signals the same way humans do with words and sentences?[1] For a long time, the consensus response has been negative: humans were the only creature with a technology of communication complex and powerful enough to convey genuine semantic currency (Hockett 1960). Nowadays opinions have shifted: we are not the only species to profit from core components of the miracle machine of language, and many of the properties once considered unique to it (e.g., syntax, compositionality, form-function decoupling) can be found in animal signals, at least in attenuated forms (Cartmill 2023).

This swing of the pendulum has had, among others, two implications. The first is a change in public ideology. In popular science and in the general press alike (e.g., Shah 2023), the perception is veering from the received wisdom that meaning is a uniquely linguistic achievement, to the notion that the "meaning of animal signals" is a legitimate, non-figurative topic of inquiry. The second is a new trend in the research on animal communication, which has seen a surge of pleas for the application of technical terminology and concepts borrowed from linguistic semantics (such as reference, denotation, or entailment) to the description of animal systems – or at least seriously hypothesizing that the operation should be pursued (examples include Zuberbühler 2020; Narbona Sabaté et al. 2022; Amphaeris et al. 2023; Berthet et al. 2023; Patel-Grosz et al. 2023; Schlenker et al. 2023; Franke, Bohn and Fröhlich 2024; Suzuki 2024).

Yet, the generalization of descriptors co-opted from theories of meaning to animal communication is a delicate undertaking, and the topic is unsurprisingly surrounded by controversy in the literature (e.g., Manser 2013; Scott-Phillips 2015; Scott-Phillips and Heintz 2023). If feasible,

1 Throughout the paper, "animal communication" designates nonhuman communication in general, and the label is not intended to prejudge matters against the signaling systems of, say, plants and mono-cellular organisms. "Animal" is used in this all-encompassing sense following standard terminological practice.

an “animal semantics” (Kuhn and Berthet 2025) would alleviate the terminological divide between animal communication research and linguistics, and facilitate comparative work across these two domains. However, there are potential downsides lurking. The generalization of (parts of) the descriptive apparatus of theories of linguistic meaning to animal communication could bias towards spurious analogies that are imposed by the conceptual priors of the theorist but do not emerge organically from the empirical data; it could lead to intellectualized representations of known properties of animal signals whose description would be stated more accurately using the “lower-order” theoretical vocabularies of evolutionary biology and classical signaling theory (Maynard Smith and Harper 2003); or it could conceal properties of animal signals that have no clear equivalent in human language and do not fit the conceptual grid of semantic models. But then: can we say anything principled – and acceptable across the spectrum of existing positions on the subject – about the constraints that the appropriation of a semantic descriptor by an animal model should respect in order to constitute a sensible scientific bargain?

I offer an argument that we can fruitfully zero in on some such basic principles. I begin by drawing a parallel between types of theories of meaning and types of theories of signal content, and propose to categorize the metatheoretical issues surrounding the application of semantic concepts to animal signals under the banner of an “animal metasemantics”. I then develop three simple conditions on the operation: a condition of epistemic fertility, one of intensional depth, and one of grammatical congruity. Finally, I showcase how these heuristics can assist theory production and help identify strengths and opportunities for improvement in emerging projects in animal linguistics.

Two clarifications before we begin. First, some of the critical points I will develop have already been anticipated in the literature. This is an intended feature, not a blind spot, of the approach on offer. The goal is to contribute to the theoretical landscape by systematizing under the banner of a clear, legible nomenclature a set of actionable heuristics that researchers would find worth enforcing, not – or not necessarily – to develop constraints that interested parties have

overlooked. Part of the exercise will involve drawing attention to desiderata of good epistemic conduct that have been contemplated by others, and giving them an explicit formulation in the hope that this may help theorists navigate the intricacies at stake in these transplants.

The second clarification is that none of this should be interpreted as an exercise in epistemic luddism. Though difficulties and problematic implications will often take center stage, the goal is not to encourage a segregationist or dogmatic distrust of conceptual swaps from theories of meaning to animal communication.[2] The goal is also not to deny that experimenting with the application of familiar descriptive tools outside their established domains can suddenly shine a spotlight on unrecognized hypotheses we might not have considered otherwise. Nor is it to reiterate the (correct, but somewhat trite) concern that appealing to linguistic concepts is wont to infest animal communication research with anthropomorphism. Many of the tenets of formal semantics may well seem to "fly in the face of much of what we already know about animal communication" (Fitch 2016: 98). Yet, resisting this reflex of incredulity may be precisely what we need to discover unnoticed junctures of commensurability between two domains.

An "animal metasemantics"?

I begin with some conceptual cartography. The purpose of the preliminary is to assign a clear theoretical niche to the questions that will occupy us for the remainder of the paper. My objective in this section is minimal. I want to: a) argue that we can draw a parallel between types of theories of linguistic meaning and types of theories of signal content; b) define the domain of an "animal metasemantics" against that backdrop; and c) file the issues surrounding the appropriation of semantic descriptors by animal models under that heading.

Philosophers of language distinguish between two basic types of theories of meaning: *semantic* and *metasemantic* theories (Speaks 2024). A *semantic* theory (also called a "descriptive

2 Not least because there is robust historical evidence that conceptual swaps between linguistics and biology have been tremendously productive. See Bromham (2017).

theory of meaning" or a "first-order theory of meaning") is a type of theory whose job is to specify the semantic properties of the expressions of an object language. For instance, a semantic theory would tell us that in Italian 'gatto' refers to cats, that in French 'courir' is atelic, that in English demonstrative 'that' denotes a certain Kaplanian function from contexts to contents, and so on. By contrast, a *metasemantic* theory (also called a "foundational theory of meaning" or a "metatheory of meaning") is a type of theory whose job is to specify the facts by virtue of which the expressions of an object language have the semantic properties they have in that language, and state the constraints over a proper semantic theory of the object language.[3] For instance, a metasemantic theory would adjudicate whether 'that' denotes the Kaplanian function it denotes by virtue of the implicit beliefs about 'that' shared by speakers of English (a "mentalist" analysis) or by virtue of facts about previous uses of 'that' in the history of English (a "causal" analysis). Furthermore, a metasemantic theory would adjudicate whether a semantic theory of the compositional properties of 'that' should be couched in the dominant approach for first-order semantic theorizing (a truth-conditional framework with variables, lambda abstraction, and the usual set of algebraic instruments) or in some alternative descriptive framework.[4]

3 As Speaks notes, the terms 'metasemantic', 'foundational' and 'metatheoretical' are used variably in the analytic literature on meaning and content. A more fine-grained way to map the landscape might combine a distinction between "semantic" and "metasemantic" theories, followed by a partition of "metasemantic" theories into two additional subtypes: "foundational theories of meaning", i.e., metasemantic work over the facts in virtue of which words and other content-bearing vehicles (see, e.g., Landrum, forthcoming) have the semantic properties they have, and "metatheories of meaning", i.e., metasemantic work over the epistemology, the ontology and the methodology of semantic theories. The simplified bipartition I have provided cuts across these points.

4 Ball and Rabern (2018: 27–28) divide metatheoretical work on meaning into five main areas: "what is semantic theorizing meant to explain?"; "does semantic theorizing aim to state facts, or is it merely an instrument?"; "what is the nature of languages and other entities discussed by semantic theories?"; "what is the relationship between languages and speakers?"; "how can we know what is the right semantic theory for a language?".

The division of labor can be generalized to theories of animal signals. For abstraction's sake, let σ be a signal type produced by an arbitrary nonhuman animal. Next, following Stegmann (2017), let us adopt a "superficial" (in Queloz's 2022 sense) notion of "signal content" tracking, in an intuitive and noncommittal fashion, whatever σ is about. As in the linguistic case, we can distinguish between two theoretical projects: a *first-order* theory of signal content and a *metatheory* of signal content. A *first-order* theory of signal content would, simply put, state the content of σ. It would tell us whether σ's content is information about a change in the probability that something (say, a predator) is present in the environment, whether σ's content is a volitional or affective state in the sender, whether σ's content is some type of social influence that tokenings of σ typically seek to exert, and so on. In contrast, a *metatheory* of signal content would specify the facts by virtue of which σ has the content it has, and state the constraints over a proper first-order theory of σ. Suppose σ's content is information increasing the probability that a predator is present in the environment. A metatheory of signal content will determine whether σ encodes probabilistic information about predator presence because σ is reliably preceded in senders, and followed in receivers, by representations of the predator in question (a "mentalist" analysis), or because signal contents are fixed by evolutionary success conditions and the production of σ in predator contexts led the species to evolutionary success (a "causal" analysis).[5] Furthermore, a metatheory of signal content would determine whether a first-order theory of σ's content should be couched in traditional information-theoretic terms, build on resources as close as possible to the vocabulary of evolutionary biology, draw from Bayesian games, and so on.[6]

5 For discussion on this, see, *inter alia*, Millikan (2004), Skyrms (2010), Stegmann (2013), and O'Connor (2020).

6 Here is, for instance, a face-value transposition of Ball and Rabern (2018: 27–28) from fn. 4: what is a first-order theory of signal content meant to explain? Do first-order theories of signal content aim to state facts, or are they merely instruments? What are the nature of signals and other entities discussed by first-order theories of signal content? What is the relationship between signals and signalers/receivers? How can we know what is the right first-order theory of content for a signaling system?

So, we can distribute theories of meaning in a bipartite cabinet comprising semantic and metasemantic theories; and we can organize projects about the contents of animal signals in a similar fashion, distinguishing first-order theories of signal content and metatheories of signal content. I noted that part of the job of metasemantic inquiry is to reflect on the tools and theoretical devices we (should) use to formulate semantic theories. And I added that the metatheory of signal content can be similarly concerned with the reflection on the conceptual devices we (should) use in the pursuit of first-order theories of signal content. One chapter of this, as was mentioned, concerns the controversies on the application of semantic notions to animal communication. Barring revisionist inclinations, whether the pursuit of a linguistic semantics is a legitimate epistemic aspiration is not an open metasemantic question. Things are different with animal signals. Whether the pursuit of an animal semantics (a first-order theory for signal contents feeding on semantic descriptors) would be a sound epistemic aspiration, and, if yes, what parts of the (dominant) metalanguage of first-order theories of meaning should be incorporated in such a theory, are open questions. These questions coalesce naturally in the domain of an "animal metasemantics": the branch of the metatheory of signal contents concerned with the conditions on the transplant of semantic descriptors into first-order theories of signal content.

Fertility, depth, congruity

Now let us turn to the question itself and begin developing the conditions. Let *s*, short for "source", be a semantic concept or term, and let *t*, short for "target", be a signaling behavior of a nonhuman animal (or animal signaling behavior in general for projects with general ambitions). I propose three simple heuristics: for every *s* and every *t*, *s* is a viable descriptor of *t* just in case *s*'s application to *t* is *epistemically fertile, intensionally deep*, and *grammatically congruous*.[7]

7 For reasons that will become clear as the discussion unfolds, the heuristics are *contiguous*, *non-exhaustive*, and *domain-specific*. They are *contiguous* in that they safeguard against neighboring epistemic harms, and the violation of one condition can generally be expected to increase the likelihood that the others will also be violated. They are

First, *s* is a viable descriptor of *t* just in case *s*'s application to *t* is *epistemically fertile*. Meaning: in its application to *t*, *s* must be able to provide a potent explanatory or descriptive purchase on *t* (potency that would vanish if the theorist adopted, in lieu of *s*, a semantically non-committal alternative) and must not pull theory production towards bad hypotheses about *t* or about phenomena surrounding *t*.

An example from the primate literature. Campbell's monkeys have two basic alarm calls, *krak* and *hok*, used in two different predatory contexts: *krak* for predatory disturbances on the ground, *hok* for predatory disturbances in the canopy. Ouattara, Lemasson and Zuberbühler (2009a, b) observed that *krak* and *hok* can be followed by a low-frequency -*oo* sound that attenuates the force of the call. For instance, while *hok* is reliably emitted only in the presence of predatory disturbances in the canopy, *hok-oo* occurs in the presence of both predatory and *non*-predatory disturbances in the canopy. For its post-call position and its effects on call content alike, this is reminiscent of the behavior of bound word-final morphemes. Based on this, Ouattara, Lemasson and Zuberbühler classified the -*oo* coda as a combinatorial suffix, and proposed that Campbell's monkey use affixation to alter call content.

Now, 'suffix' and 'affixation' are morphologically charged terms,[8] and in a subsequent comment, Kuhn et al. (2018: 171) are careful to point out that "this use of terminology should not be interpreted […] as committing to any deeper analogy with spoken language, such as postulating -*oo* as a sublexical morpheme or a sentence-final particle". Approximating, the real empirical question is whether -*oo* is able to combine with *krak* and *hok* compositionally; morphemic jargon is just a convenient shorthand. The hedging may be indicative of an issue of epistemic urgency. First,

non-exhaustive in that the conditions do not rule out the existence of undercutting factors that may cancel the initial epistemic incentive provided by their satisfaction. Lastly, they are *domain-specific* in that they are not "one-size-fits-all": I offer no argument that similar desiderata should be invoked for unrelated cross-disciplinary swaps (say, concept borrowing from fluid mechanics to population science).

8 Whose proper definition, incidentally, is itself discussed. See, e.g., Haspelmath (2020).

it is unclear whether there are unique descriptive or explanatory benefits the model of Campbell's monkeys' calls stands to gain by adopting morphologically charged terminology as opposed to more conservative notions such as that of a call-final unit with regular (perhaps compositional) effects on call content. Second, the choice of morphologically charged terminology could encourage overly favorable expectations about Campbell's monkeys' proximity to systems with genuine morphological capacities – expectations that the field may then struggle to counter because of positive bias. Finally, terminological stipulations often take on a life of their own. Over time, the field could shift from descriptions on which Campbell's monkeys' calls pattern the way they do *as if* they had a morphological structure, to descriptions on which such calls pattern the way they do *because* they a morphological structure, reifying the shorthand into a causally efficacious feature.[9] Should these reservations be correct (an issue I will not discuss further), the recruitment of morphological terminology in this context would not be epistemically fertile. Modeling the call system of Campbell's monkeys *as if* it featured suffixes does not incur obvious empirical issues; however, unless the application is demanded by the data or leads to unique insights into the behavior of the call components it is meant to label, parsimony would be preferable.

The second heuristic: *s* is a viable descriptor of *t* just in case *s*'s application to *t* is *intensionally deep.* Meaning: in its application to *t*, *s* must be associated with binding, transparent conditions of application that prevent the risk of shallow appeals to the notion in *t*'s domain. In other words, in describing *t*, watch out for metaphorical or casual appeals to *s*, and make sure your prior theory of *s* incorporates conceptual components you are sufficiently prepared to account for.

Here is a high-contrast example to clarify. Imagine the following is the true, complete definition of what an "arbitrary" word is: a word *w* is arbitrary just in case *w* lacks perceptual resemblance with its referent and *w* could have evolved to denote something else.[10] The English noun 'gift' is then arbitrary. Its perceptible contour does not guide language learners towards the

9 Cf. Owren, Rendall and Ryan (2010: 766) on the "information metaphor" in models of animal signals.

10 The issue is obviously more complex. See Gasparri et al. (2023).

belief that the word refers to gifts (unlike ‘boom’ and explosions), and, historically, ‘gift’ could have developed an alternative denotation (it could have meant poison, as /gɪft/ does in German). Now suppose we wish to establish whether some animal A is capable of arbitrary communication. We discover that A has a repertoire of learned calls bearing no correspondence between form and content. We have a rough-and-ready grasp of what arbitrariness consists of, and we stipulate that ‘arbitrary’ simply means accidental or devoid of purpose. We then say: since the calls of A are learned, they must be a cultural accident. So we conclude that A should be ascribed a capacity to produce arbitrary signals.

It should be clear that in this scenario the application of the notion of arbitrariness would be shallow. There are at least three issues. First, the extension of the property that analyzes ‘arbitrary’ cannot be decided absent a sustained definition of what ‘accidental’ and ‘devoid of purpose’ mean, and there are a few non-trivial options in theoretical space. Second, whether A should be ascribed a capacity to produce arbitrary signals is a substantive question; answering it with a casual application of the “arbitrary” descriptor threatens to shift the focus to a simplistic caricature of the original problem. Third, settling arbitrariness ascriptions on the basis of this vague checklist exposes them to the risk of opportunistic applications and verbal disagreements. The problem is not, or not necessarily, that the operationalization underlying the ascription does not correspond to the “true” definition of arbitrary word stipulated earlier (one horn of the initial definition has gone missing: the possibility that the signal could have evolved to bear a different content). The problem is that the definition is simply too loose, that its commitments are not specified in a sufficiently binding manner, and that by incorporating conceptual gears whose contours have not been defined with the rigor we would expect of a non-superficial treatment of the problem, the terminology can be opportunistically modulated to fit one’s prior ideological affiliations – e.g., those of a theorist antecedently biased towards multiplying, or reducing at all costs, the amount of signature features of human language with a counterpart in animal systems.

Finally, the third heuristic: *s* is a viable descriptor of *t* just in case *s*'s application to *t* is *grammatically congruous*. Meaning: in its application to *t*, *s* cannot take on non-shallow conditions of application with revisionary extensional consequences in semantic theory. In other words, during the swap into its new domain, *s* must not morph into a construct that, swapped back into first-order theories of meaning, would generate descriptive results incongruent with those generated by its initial incarnation.[11]

Another example to clarify. Engesser et al. (2015) observed that chestnut-crowned babbler vocalizations comprise two tonal call units, A and B, acoustically differentiated by pitch contour and capable of combining into more complex call sequences. They were further able to demonstrate that BAB sequences stimulate changes in listener behavior that are not observed in response to calls comprising only the B element or in response to AB calls. They inferred that the addition of the B element at the beginning of the sequence AB changes the content of the call in a manner comparable to phoneme additions to natural language words: appending B at the beginning of AB causes call content to change just as appending 'p' at the beginning of the preposition 'at' turns 'at' into (the noun, verb, or adjective) 'pat'. They therefore concluded that certain elements of chestnut-crowned babbler vocalizations behave like human phonemes.

The reasoning appears sound, but as subsequent discussion was quick to acknowledge (see, *inter alia*, Bowling and Fitch 2015, Engesser and Townsend 2019, Huybregts 2020) the heuristic risks giving us phonemic status on the cheap.[12] In English, the unaspirated phone [t] and the aspirated phone [t^h] are allophones of the phoneme /t/. The reason is that the lexicon of English lacks minimal pairs of words whose canonical phonetic implementations differ only in the

11 Note the logical independence of intensional depth and grammatical congruity. A theoretical construct could be intensionally shallow but allow for grammatically congruous specifications, or it could be intensionally deep but have commitments that make it grammatically incongruous.

12 For a recent discussion of parallel difficulties for a completely different taxon, see Sharma et al. (2024) and Beguš et al. (2024) on the application of phonological descriptors to sperm whale codas.

substitution of [t] with [t^h], and vice versa. In Thai, in contrast, [tam] means "to pound", whereas [t^ham] means "to do", so [t] and [t^h] must be allophones of two distinct phonemes, /t/ and /t^h/. As this shows, the comparison between two phonetic strings reveals a phonemic difference only when there is a single-segment contrastive difference between *non-empty* phonetic positions coupled with a difference in meaning (Dresher 2011). Now look at what happens if we follow Engesser et al.'s proposal to the letter. Suppose we start with a canonical phonetic implementation of the preposition 'at': [æt]. We now append [p^h] at the beginning of the sequence, and obtain [p^hat]. [p^hat] is a canonical phonetic implementation of 'pat', and the addition of [p^h] has brought about a difference in meaning. So English has a /p^h/ phoneme? It does not. But then we have a false positive and there is something wrong with the initial classification. That something is that the proper test for phonemic status is "discriminative" (Bowling and Fitch 2015), and that phonemic status cannot be adjudicated just by noting that the addition or the elision of a sound unit from a sequence causes a change in content. Irrespective of whether or not the call system of chestnut-crowned babblers has phonemic contrasts after all, the point should be clear: there is an argument that Engesser et al.'s treatment runs on a heuristic for phonemic status that, applied to human language, would grant phonemic status to units that are *not* regarded as phonemes in consensus models of English phonology. This is grammatical incongruity: the operationalization has paid for the application in the new domain in the coin of a micro-insurrection against what the concept was supposed to be about in the first place.

So, I have described three simple metasemantic conditions on the recruitment of semantic descriptors by first-order theories of signal content: a condition of *epistemic fertility*, one of *intensional depth*, and one of *grammatical congruity*. The remainder of the paper will do two things. First, I will rehearse some skeptical arguments against the notion of "functional reference" in primate communication research, and propose that these arguments can be reconstructed as pointing to violations of the conditions outlined. The purpose of the exercise is to show that the conditions

are not exotic concoctions and track methodological desiderata that others have found reasonable to invoke. Second, I will showcase how the heuristics can help us assess the strengths and weaknesses of emerging proposals in animal linguistics, taking as an example Berthet et al.'s (2023) case for the introduction of a notion of "signal denotation" in animal communication research.

Functional reference

According to the standard definition, an animal signal σ is "functionally referential" if and only if a) σ's production is specific to a context, where the relevant contextual feature defines the referent of σ, and b) σ triggers appropriate receiver-side inferences about the context of production even if σ's referent is absent. In other words, a signal σ is functionally referential if and only if σ is context-specific and elicits relevant stimulus-independent responses in signal perceivers (Townsend and Manser 2012).

The notion of functional reference originated from the observation of the call system of East African vervet monkeys. These monkeys were observed to produce distinct alarm calls in response to leopards, eagles, and snakes, calls that elicited predator-specific behaviors regardless of the presence of the corresponding predator. Even in the absence of actual predator stimuli, "leopard alarms" prompted listeners to climb trees, "eagle alarms" caused them to look up, and "snake alarms" made them scan the ground. Initially, the combination of predator specificity and stimulus independence was thought to warrant an ascription of full-blooded referentiality and, with it, the conclusion that these call types discriminated the predator types much like kind terms in natural language. However, the fact that these calls were in all likelihood not amenable to intentional regulation, and lacked the "arbitrary", symbolic character of human words and sentences, was considered problematic.[13] To attenuate these issues, Marler et al. (1992) and Macedonia and Evans (1993) introduced the notion of "*functional* reference". The job of the notion was to signal that vervet alarm calls achieved something ostensibly equivalent to linguistic reference, but did so via

13 For more on arbitrariness and symbolic status, see Planer and Kalkman (2021).

different underlying mechanisms; they were "not exactly like human words" and yet "appear[ed] to function in the same way" (Hauser 1996: 509).

Since its introduction more than three decades ago, the construct has enjoyed remarkable success in the literature on animal communication. It has since been applied to several other primate and non-primate species, and has been taken to track an evolutionary precursor of linguistic reference (Cheney and Seyfarth 1996; 2005). However, there have also been dissenting voices and, over time, the construct's appropriation of referential terminology has drawn increasing criticism. Let me ease into the gist of the problem with an armchair example of my own construction.

Suppose Josh is a 9-year-old native speaker of English with normal linguistic abilities. Further, suppose that Josh's linguistic behavior takes an idiosyncratic but perfectly regular turn whenever he is hungry: whenever Josh is hungry, he begins repeating "School is boring" in a loop. Josh is perfectly aware of what "School is boring" means in English (in Josh's mental lexicon, in the syntactic module of his grammar, and so on, everything is as it should) and, being a zealous and enthusiastic student, has never been heard saying "School is boring" in any contexts other than those where he is hungry. Finally, suppose that Josh's parents are aware of Josh's peculiar linguistic propensities. Every time they hear Josh say "School is boring", they infer that Josh must be hungry and rush to the kitchen to prepare some food.

Let us examine the situation. Josh's utterances of "School is boring" occur exclusively in Josh-is-hungry contexts. So they meet the requirement of context-specificity. They are also stimulus-independent for the relevant perceivers: if an experimenter hid a set of speakers in Josh's house and played a recording of Josh saying "School is boring" in Josh's absence, Josh's parents would form the thought that Josh is hungry and should be fed. Both criteria for functional referentiality are met. So Josh's utterances of "School is boring" functionally refer to Josh's hunger.

The verdict is puzzling, and we can see why through the heuristics described in the previous section. To begin with, the application faces a macroscopic issue of grammatical congruity. The

conjunction of context-specificity and stimulus-independence leads to the conclusion that Josh's utterances of "School is boring" functionally refer to something entirely orthogonal to the semantic reference of "School is boring" (the proposition that school is boring). We were told that "functional reference" was supposed to be semantic reference stripped of the confounds stemming from the symbolic character of human words; instead, we seem to have changed the subject. Used as a diagnostic tool for the distribution of referential relations in language, the construct ends up putting words and sentences in referential relation with contextual features that nobody would regard as part of their actual reference (cf. Wheeler and Fischer 2012).

The second point is intensional depth. The adoption of referential vocabulary to categorize the relation between Josh's utterances of "School is boring" and Josh's hunger holds up only if referentiality is interpreted in a non-technical sense, akin to the correlational meaning the verb 'refer' has in sentences like "These symptoms refer to an abnormal mental function". But this is not the sense of referring that matters to theories of linguistic reference. Furthermore, the relationship between this correlational phenomenon and semantic reference is not antecedently clear. Imagine a variant of the scenario above where Josh, instead of saying "School is boring" whenever he is hungry, says "I am hungry". Based on the standard diagnostics for functional reference, all we can infer is that "School is boring" and "I am hungry" are functionally-referentially synonymous across these two scenarios; we cannot provide even the beginning of an account of why "I am hungry" refers to what it semantically refers to in English, and of why only in the scenario where Josh says "I am hungry" if and only if he is hungry, semantic reference and functional reference coincide. On these matters, the notion is in full "conceptual retreat" (Rendall et al. 2009: 236).

Finally, there is an argument that the notion is not epistemically fertile. First, the appeal to referential vocabulary does not seem indispensable or uniquely productive when it comes to characterizing the regularities at work in Josh's case (and, retrospectively, in the monkey cases). We can capture all the necessary facts with the more parsimonious explanation that Josh's utterances of

"School is boring" are expressive *indices* of hunger that trigger reliable inferences about the situation of production (i.e., that the context is a Josh-is-hungry context) because the receivers (the parents) know what the utterances of the sentence "naturally" mean, in Grice's sense (Fischer 2020). This stands in contrast to cases of genuine semantic reference (say, the relationship between occurrences of the proper name 'Mars' in subject position of simple clauses about the planet Mars, and the planet Mars itself), which cannot be paraphrased away in a similar manner. The second problem is that because the concept captures something whose relationship with semantic reference is unclear, it offers "a poor diagnostic to establish whether [given] calls might constitute preadaptations to human speech" (Fischer and Price 2017: 26). If the construct is so wanting on the front of grammatical congruity, and faces difficulties of this magnitude in reproducing the referential relations *our* words entertain with the world, how can it be expected to produce sound hypotheses about reference in nonhuman systems?

I recognize that I am painting a somber picture here, but it is important to be straightforward. More than a decade ago, Wheeler and Fischer (2012: 204) urged the field to "drop the term 'functionally referential signals' from the animal communication literature in favor of more accurate, and linguistically neutral, descriptions such as 'context-specific signals', 'predator-specific alarm calls', or 'food-specific calls'." If they were right and if, as I have argued, the notion of functional reference has substantial metasemantic weaknesses, it is unclear what to make of the body of literature that has continued to proliferate under the banner of the original construct. LaPergola et al. (2023) argue that greater anis, a bird in the cuckoo family, are capable of "referential signaling" (*sic*, without the "functional" qualifier), using the classical checklist for functional reference by Macedonia and Evans (1993): "referential signals are reliably associated with specific objects or events in the environment and can convey information to receivers about these referents". Metasemantic rigor may seem pedantic; it need not therefore be discretionary.

Signal denotations

I now turn to showcasing how the conditions can be used to evaluate emerging proposals in animal linguistics, discussing Berthet et al. (2023)'s argument for the introduction of a notion of "signal denotation" in animal communication research (see also Kuhn and Berthet 2025). Berthet et al.'s paper contributes to the animal linguistics program pioneered by Schlenker et al. (2016), and offers one of the most carefully argued proposals on the market about the generalization of semantic concepts to animal communication. Berthet et al.'s goal is to "establish strong basic foundations for animal linguistics" (p. 82); that is, "to provide linguists with the tools to study animal communication, and to provide biologists with basic linguistic notions applicable to the study of animal communication, using concepts and criteria compatible with modern linguistic thinking" (p. 82). This involves "[providing] precise definitions of the main linguistic concepts [...] using general principles that can be applied equally well to human and non-human communication" (p. 83), thereby individuating "minimal criteria to be fulfilled to claim that a given species displays a particular linguistic capacity" (p. 81).

Berthet et al. set the stage for their proposal by rehearsing a standard division of labor between semantic and pragmatic meaning. In semantics, the denotation of an expression is "its stable semantic contribution" (p. 85). For instance, "It's 5 pm" may pragmatically imply that it is too early for dinner, that we are late for our 5:15 appointment, that it will soon be dark outside, and so on. However, the denotation of the sentence remains constant: across these pragmatic uses, the sentence invariantly denotes the proposition that, at the place of utterance, the time of utterance is 5 pm. They then suggest that applied to animal signals, denotation can be defined as "the largest set of meaningful features of circumstances that appear across all occurrences of the signal" (p. 85), where a feature is "meaningful" just in case it "[merely] appears at a rate greater than chance across the signal's occurrences" (p. 84).[14] For instance, "if a signal is produced 70% of the time in response to

14 I add "[merely]" to account for the fact that Berthet et al. exclude from the category of "meaningful" features those that "appear at a rate greater than chance" because they are *always* instantiated. Example: the sender-side possession

leopards and the other 30% of the time in response to eagles, the denotation of the signal is the set of features common to all occurrences: presence of a predator" (p. 85).

Now, let us review the intensional depth and the grammatical congruity of the maneuver, starting with the elephant in the room.[15] The cornerstone of formal treatments of natural language meaning is the notion of truth. The noun 'hamster' denotes the sortal property of being a hamster; it is true of entities that belong to the set of hamsters in the domain, and false of those that do not. Similarly, the sentence "Mary's hamster is on the Eiffel Tower" denotes the proposition that Mary's hamster is on the Eiffel Tower; the sentence is true of worlds where Mary's hamster is on the Eiffel Tower and false of worlds where Mary's hamster is somewhere else (Heim and Kratzer 1998). Now, humans are capable of forming truth-evaluable thoughts. Truth-conditional treatments of natural language meaning appeal therefore to a property within reach of the representational capacities of ordinary speakers. However, the prospects of packing truth into a credible reductive or instrumental picture of the representational capacities available to many animals are less clear.

Berthet et al. acknowledge the complication and offer a preemptive fix, due to Schlenker et al. (2016): retain the standard intensional apparatus of the notion of denotation but anchor signal denotations, rather than to conditions of signal truth, to conditions of signal "applicability" or "adequacy".[16] That is, instead of giving the semantics for animal signals via formulas of type $[\![x]\!] = 1$

---

of a functioning articulatory apparatus appears across all occurrences of a primate alarm call. However, it does not "merely" appear at a rate greater than chance across the signal's occurrences. So it is not "meaningful".

15 I will not comment on matters of epistemic fertility, except to note that Berthet et al. acknowledge their importance. They say, for instance, that by modeling animal signals using constructs borrowed from first-order theories of meaning, the theorist can "overcome the great differences between biological and linguistic methodologies" (p. 82), "[improve] the relevance of comparisons between human and animal communicative systems" (p. 82), and "draw precise theories about the use and structure of animal signals" (p. 94).

16 From Schlenker et al. (2016: 4): "While we could say that a [monkey] sentence is *adequate* or *inadequate* in a certain situation, we prefer to employ traditional logical terminology and use the terms *true* and *false*, which are more familiar to linguists. These terminological moves are just intended to facilitate the discussion". And later: "at a

just in case P, 0 otherwise – where *x* is the signal, the double square brackets are the function to the denotation of *x*, and '1' and '0' are truth and falsity – use formulas of type ⟦*x*⟧ = A just in case P, N/A otherwise, where 'A' and 'N/A' are applicability and non-applicability.[17]

However, the substitution raises concerns of its own. Two specific observations. First, absent some argument to the contrary, the bipartition between "applicable" and "inapplicable" may reinstate a successor to the problem raised by truth. Applicability is a normative property, and it is unclear that a model of nonhuman signaling appealing to (tacit) representations of applicability conditions would be significantly less contentious than one appealing to signal truth.[18] Perhaps there is an argument to the contrary for species that – pardon the approximation – are close enough to humans on the phylogenetic tree or have signals in modalities analogous to the ones exploited by language. But the intuitive burden increases the moment these features disappear from the picture. Common minnows have been found to produce chemical signals that reliably occur in the presence of predators and trigger coordinated group defense in familiar conspecifics (Bairos-Novak, Ferrari, and Chivers 2019). Based on what we have said, it would seem we have three options: a) project applicability conditions and knowledge of applicability conditions onto freshwater fish and their chemical signals; b) dismiss the theoretical apparatus that generated the projection; c) reject the projection but hold the line and insist that the apparatus remains viable for monkeys and other animals (the problem would be *which*).

The second observation is that the bipartition between applicable and inapplicable cannot casually replace the distinction between true and false. A truth-conditional semantics can be rebranded as an applicability-conditional semantics provided we can replicate with 'A' and 'N/A' the

---

minimum, monkeys must know under what conditions a call is or isn't applicable – and the bipartition between *applicable* and *inapplicable* is just the distinction between *true* and *false* under a different name" (p. 80). I think there is a tension between the stipulative tone of the first quote and the second one. Back to this in a moment.

17 The notation is my responsibility.

18 For more on normativity in non-human animals, see the essays in Roughley and Bayertz (2019).

operations that '1' and '0' generated in the original system. And here there are some loose ends. One problem is that some of the patterns of entailment licensed by the truth predicate become unstable under applicability and adequacy. Suppose a monkey M has a specialized alarm call *e* produced only in the presence of eagles, and a general call *p* produced in the presence of all kinds of predator threats, including occasionally eagles. On the approach on offer, *e* denotes that an eagle is present, and *p* that a predator of some kind is present. Eagles are predators, and $[\![e]\!] = 1$ entails $[\![p]\!] = 1$, a perfectly fine result given what we have said about the distribution of the two signals. Now let us try to replicate things with A and N/A. Were this a benign replacement, we should find it unproblematic to conclude, from the premise that $[\![e]\!] = 1$ entails $[\![p]\!] = 1$, that $[\![e]\!] = \text{A}$ guarantees $[\![p]\!] = \text{A}$. But this goes against the premise that the eagle-contexts are not contexts where M sends *p*. Perhaps we made an unfriendly assumption about the denotation of the general call to begin with, and the real extension of *p* is logically complex: the set of all the predator contexts *minus* the eagle-contexts. Or, following again Schlenker et al. (2016), *p* is in fact semantically applicable in the eagle-contexts but M represents it as inapplicable under a pragmatic constraint of informativity or urgency. In such a case, the risk is that applicability would morph into a hybrid predicate tracking speech-act-level conditions of situational felicity that are orthogonal to strict *semantic* truth and falsity – on par with ordinary uses of the applicability jargon in linguistic reasoning: assertions of declarative sentences can be "applicable" even if the sentences themselves are semantically false (and, for instance, flout a Gricean maxim), and "inapplicable" even if the sentences themselves are semantically true (but, for instance, violate a Gricean maxim).[19]

19 Think of scalar implicatures with bare numerals (Spector 2013). "Two cats arrived" is strictly true if five cats arrived, but any rational speaker hearing "Two cats arrived" would reliably infer that *exactly* two cats arrived, and would say that the sentence is "inadequate" if five cats arrived. So, "Two cats arrived" is applicable only in contexts where exactly two cats arrived. But the fact remains that "Two cats arrived" is semantically true of contexts where *at least* two cats arrived. Applicability and semantic truth are separate matters.

This is not to suggest that there are no conceivable repairs or workable paths to stabilize things on the horizon. One would be rejecting sentences as the appropriate point of comparison for alarm calls and other signals, and assimilating predator calls to pre-propositional vehicles devoid of any covert predicative structure. Under this approach, for instance, a predator call would be analogous to exclaiming "predator!" upon sighting a predator. The signal would semantically denote the predator predicate, its production would introduce the set of predators in the conversational scoreboard, and nothing more. The call would be truth-conditionally empty, and would have purely pragmatic conditions of applicability supervening on the fact that the perception of tokens of this non-propositional signal reliably leads listeners to form a thought *that* a predator is near. This way, one would obtain a hybrid treatment that preserves a use-invariant communicative contribution for the call type, while regarding call tokens as subject to conditions of situational applicability and inapplicability that do not involve truth and falsity. Be that as it may, unless the specifics of the repair are worked out in sufficient detail, the concern that the framework is threatened by a shallow insistence on denotational terminology remains.

I now turn to the grammatical congruity of the proposal. Berthet et al. say that the definition of denotation under discussion should "[apply] equally well to human and non-human communication" (p. 83). The standardization should therefore generate non-revisionary results if employed as an algorithm to individuate natural language denotations. Recall the definition: "the semantic denotation of a signal is the largest set of meaningful features of circumstances that appear across all occurrences of the signal", and a feature is meaningful if it "[merely][20] appears at a rate greater than chance across the signal's occurrences". So the semantic denotation of an expression *e* should be the largest set of meaningful features of circumstances that appear across all occurrences of *e*. This, however, does not seem right.

One problem echoes familiar objections to "correlational" accounts of signal meaning (Stegmann 2009; Scarantino 2015). Natural languages have expressions for circumstantially

20 See above, fn. 14.

infrequent objects, properties and events. The noun ‘pangolin’ denotes pangolins, and pangolins are a rare mammal. If a Martian civilization landed on Earth and tried to reconstruct the denotation of ‘pangolin’ in English based on the distribution of situational features found at utterances of the English word ‘pangolin’, it would, *prima facie*, have a hard time getting at the actual denotation of ‘pangolin’.[21] *Mutatis mutandis* for the denotation of features that are reliably present whether or not the “linguistic signal” is produced: for instance, expressions for environmental constants like breathable air. If I understand correctly, features of this sort cannot be “meaningful” in the sense invoked by the approach, since they do not merely occur at a rate greater than chance across utterances. They should therefore be denotational impossibilities, which, however, is not the case.

A related concern is that the machinery of meaningful features seems to lead to a view in which the denotations of signal types are maximal collections of sets of worldly features reliably found across contexts of signal tokening. But in natural language, the properties of utterance contexts do not trickle down into lexical and sentential denotations in a comparable manner. Regularities of situations of utterance can be orthogonal to signal “meaning”, and even when linguistic signals do reliably co-occur with given sets of contextual features, these can have nothing to do with what the signals involved are about in the sense at issue in semantics. Suppose K is a population of speakers of English who know the meaning of the word ‘elephant’ and reliably use ‘elephant’ correctly. Suppose, further, that members of K rarely think of elephants, but that in K utterances of ‘elephant’ are quasi-invariably accompanied by the formation of a thought about elephants. Thinking about elephants is then a meaningful feature in the largest set of features that appear at a rate greater than chance across occurrences of ‘elephant’ in K. So having a thought about

21 From another angle: whether (some) animal signals are capable of “denoting”, “having in their extension” or “singling out” things that reliably fail to occur in contextual partnership with the signals themselves should be an empirical question; the definition risks ruling that possibility out on conceptual grounds.

elephants should be part of what 'elephant' denotes in K. But 'elephant' does not denote thoughts of elephants (in K or elsewhere); it denotes elephants.[22]

To clarify: none of this is to argue that Berthet et al.'s notion of "signal denotation" is incoherent or empty; the remarks I have offered are perfectly compatible with the idea that the apparatus of meaningful features could lead to important experimental advancements in animal communication research. Nor is it to argue that it would not be possible to develop an iteration of the framework that manages to address the issues above. The point is that in doing so we should be prudent with the appropriation of semantically charged parlance. As things stand, there is reason to doubt that the "signal denotations" of Berthet et al. are about what denoting is in natural language. Calling them "denotations" anyway risks doing the opposite of facilitating impartial comparisons between natural language and animal communication.

Conclusion

The goal of this paper has been to contribute to the ongoing discussion about applying constructs from linguistic semantics to animal communication. I began by drawing a parallel between types of theories of linguistic meaning and types of theories of signal content, and suggested filing the metatheoretical issues surrounding the incorporation of semantic descriptors in animal models under the banner of an animal metasemantics. I then described three simple conditions on these incorporations. Specifically, I suggested that conceptual and terminological swaps from first-order theories of meaning to first-order theories of signal content should be epistemically fertile, intensionally deep, and grammatically congruous. Finally, I offered evidence that requirements analogous to the ones developed here have already been enforced in the debate, and showcased how

22 At best, utterances of 'elephants' "denote" thoughts of elephants in the non-technical, "Grice-natural" sense in which we say that an angry tone of voice "denotes" a marked displeasure, which, however is not the sense of "denoting" at issue in natural language semantics. If the observation rings a bell, recall the comments to the example "These symptoms refer to an abnormal mental function" in the discussion of functional reference.

the conditions can help identify strengths and opportunities for improvement in emerging projects in animal linguistics.

One final point. My focus in this paper has been on semantic descriptors, but I have noted that parallel questions apply to the appropriation of terms and tools from other branches of linguistic inquiry. As Scott-Phillips and Heintz (2023: 97) report, recent contributions in animal communication research have appealed, just to name a few, to "phonocoding", "minimal units", "hierarchical dependencies", "temporal / stochastic / readout / proportional structures", "combinatoriality", "segmental concatenations", and "merged compounds". In each case, we deal with a theoretical construct borrowed from linguistic theory and applied to animal communication, whether with or without adjustments, with or without explicit adjustments, and with or without an argument for the adjustments involved. Reflecting on the expediency of these attempts through versions of the heuristics described in this paper could help us stay on course.

References

Amphaeris, J., Blumstein, D. T., Shannon, G., Tenbrink, T., & Kershenbaum, A. (2023). A multifaceted framework to establish the presence of meaning in non-human communication. *Biological Reviews*, 98, 1887–1909.

Ball, D., & Rabern, B. (Eds.) (2018). *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics*. Oxford: Oxford University Press.

Bairos-Novak, K. R., Ferrari, M. C. O., & Chivers, D. P. (2019). A novel alarm signal in aquatic prey: Familiar minnows coordinate group defences against predators through chemical disturbance cues. *Journal of Animal Ecology*, 88, 1281–1290.

Beguš, G., Sprouse R., Leban, A., Silva, M., & Gero, S. (2024). Vowels and diphthongs in sperm whales. *OSF Preprints.* doi: 10.31219/osf.io/285cs.

Berthet, M., Coye, C., Dezecache, G., & Kuhn, J. (2023). Animal linguistics: A primer. *Biological Reviews*, 98, 81–98.

Bowling, D. L., & Fitch, W. T. (2015). Do animal communication systems have phonemes? *Trends in Cognitive Sciences*, 19, 555–557.

Bromham, L. (2017). Curiously the same: Swapping tools between linguistics and evolutionary biology. *Biology & Philosophy*, 32, 855–886.

Cartmill, E. A. (2023). Overcoming bias in the comparison of human language and animal communication. *Proceedings of the National Academy of Sciences*, 120, e2218799120.

Cheney, D. L. & Seyfarth, R. M. (1996). Function and intention in the calls of non-human primates. *Proceedings of the British Academy*, 88, 59–76.

Cheney, D. L. & Seyfarth, R. M. (2005). Constraints and preadaptations in the earliest stages of language evolution. *Linguistic Review*, 22, 135–159.

Dresher, B. E. (2011). The phoneme. In M. Oostendorp, C. J. Ewen, E. Hume, & K. Rice (Eds.), *The Blackwell Companion to Phonology* (pp. 241–266). London: Wiley.

Engesser, S., Crane, J. M. S., Savage, J. L., Russell, A. F., & Townsend, S. W. (2015). Experimental evidence for phonemic contrasts in a nonhuman vocal system. *PLOS Biology*, 13, e1002171.

Engesser, S., & Townsend, S. W. (2019). Combinatoriality in the vocal systems of nonhuman animals. *WIREs Cognitive Science*, 10, e1493.

Fischer, J. (2020). Nonhuman primate alarm calls then and now. *Animal Behavior and Cognition*, 7: 108–116.

Fischer, J., & Price, T. (2017). Meaning, intention, and inference in primate vocal communication. *Neuroscience & Biobehavioral Reviews*, 82, 22–31.

Fitch, W. T. (2016). Why formal semantics and primate communication make strange bedfellows. *Theoretical Linguistics*, 42, 97–109.

Franke, M., Bohn, M., & Fröhlich, M. (2024). Latent meaning representations in great-ape gestural communication. In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society*, 633–640.

Gasparri, L., Filippi, P., Wild, M., & Glock, H.-J. (2023). Notions of arbitrariness. *Mind & Language*, 38, 1120–1137.

Haspelmath, M. (2020). The morph as a minimal linguistic form. *Morphology*, 30, 117–134.

Hauser, M. (1996). *The evolution of Communication*. Cambridge, MA: MIT Press.

Heim, I., & Kratzer, A. (1998). *Semantics in Generative Grammar.* Malden, MA: Blackwell.

Hockett, C. F. (1960). The origin of speech. *Scientific American*, 203, 88–96.

Huybregts, R. (2020). Babbling birds. *Inference: International Review of Science*, 5. doi: 10.37282/991819.20.39.

Kuhn, J., & Berthet, M. (2025). Animal semantics. In *International Encyclopedia of Language and Linguistics*, 3rd edn. Amsterdam: Elsevier. doi: 10.1016/B978-0-323-95504-1.00566-4.

Kuhn, J., Keenan, S., Arnold, K., & Lemasson, A. (2018). On the *-oo* suffix of Campbell's monkeys. *Linguistic Inquiry*, 49, 169–181.

Landrum, K. (forthcoming). Good explanations and the causal metasemantics of inductive concepts. *Mind.* doi: 10.1093/mind/fzaf024.

LaPergola, J. B., Savagian, A. G., Smith, M. G., Bennett, B. L., Strong, M. J., & Riehl, C. (2023). Referential signaling in a communally breeding bird. *Proceedings of the National Academy of Sciences*, 120, e2222008120.

Macedonia, J. M., & Evans, C. S. (1993). Essay on contemporary issues in ethology: Variation among mammalian alarm call systems and the problem of meaning in animal signals. *Ethology*, 93, 177–197.

Manser, M. B. (2013). Semantic communication in vervet monkeys and other animals. *Animal Behaviour*, 86, 491–496.

Marler, P., Evans, C. S., & Hauser, M. D. (1992). Animal signals: Motivational, referential, or both? In H. Papoušek, U. Jürgens, & M. Papoušek (Eds.), *Nonverbal Vocal Communication: Comparative and Developmental Approaches* (pp. 66–86). Cambridge: Cambridge University Press.

Maynard Smith, J., & Harper, D. (2003). *Animal Signals*. New York: Oxford University Press.

Millikan, R. G. (2004). *The Varieties of Meaning*. Cambridge, MA: MIT Press.

Narbona Sabaté, L., Mesbahi, G., Dezecache, G., Cäsar, C., Zuberbühler, K., & Berthet, M. (2022). Animal linguistics in the making: The Urgency Principle and titi monkeys' alarm system. *Ethology Ecology & Evolution*, 34, 378–394.

O'Connor, C. (2020). *Games in the Philosophy of Biology.* Cambridge: Cambridge University Press.

Ouattara, K., Lemasson, A., & Zuberbühler, K. (2009a). Campbell's monkeys concatenate vocalizations into context-specific call sequences. *Proceedings of the National Academy of Sciences,* 106, 22026–22031.

Ouattara, K., Lemasson, A., & Zuberbühler, K. (2009b). Campbell's monkeys use affixation to alter call meaning. *PLoS ONE*, 4, e7808.

Owren, M. J., Rendall, D., & Ryan, M. J. (2010). Redefining animal signaling: Influence versus information in communication. *Biology & Philosophy*, *25*, 755–780.

Patel-Grosz, P., Mascarenhas, S., Chemla, E., & Schlenker, P. (2023). Super linguistics: An introduction. *Linguistics and Philosophy*, 46, 627–692.

Planer, R. J., & Kalkman, D. (2021). Arbitrary signals and cognitive complexity. *The British Journal for the Philosophy of Science*, 72, 563–586.

Queloz, M. (2022). The essential superficiality of the voluntary and the moralization of psychology. *Philosophical Studies*, 179, 1591–1620.

Rendall, D., Owren, M. J., & Ryan, M. J. (2009). What do animal signals mean? *Animal Behaviour*, 78, 233–240.

Roughley, N., & Bayertz, K. (Eds.). (2019). *The Normative Animal? On the Anthropological Significance of Social, Moral, and Linguistic Norms*. Oxford: Oxford University Press.

Scarantino, A. (2015). Information as a probabilistic difference maker. *Australasian Journal of Philosophy*, 93, 419–43.

Schlenker, P., Chemla, E., Schel, A. M., Fuller, J., Gautier, J.-P., Kuhn, J., … Zuberbühler, K. (2016). Formal monkey linguistics. *Theoretical Linguistics*, 42, 1–90.

Schlenker, P., Coye, C., Leroux, M., & Chemla, E. (2023). The ABC-D of animal linguistics: Are syntax and compositionality for real? *Biological Reviews*, 98, 1142–1159.

Scott-Phillips, T. C. (2015). Meaning in animal and human communication. *Animal Cognition*, 18, 801–805.

Scott-Phillips, T., & Heintz, C. (2023). Animal communication in linguistic and cognitive perspective. *Annual Review of Linguistics*, 9, 93–111.

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

Shah, S. (2023). The animals are talking. What does it mean? *The New York Times Magazine*, September 20th, 2023.

Sharma, P., Gero, S., Payne, R. Gruber, D. F., Rus, D., Torralba, A., & Andreas, J. (2024) Contextual and combinatorial structure in sperm whale vocalisations. *Nature Communications*, 15, 3617.

Speaks, J. (2024). Theories of meaning. *In The Stanford Encyclopedia of Philosophy (Fall 2024 Edn.)*. https://plato.stanford.edu/archives/fall2024/entries/meaning/

Spector, B. (2013). Bare numerals and scalar implicatures. *Language and Linguistics Compass*, 7, 273–294.

Stegmann, U. E. (2009). A consumer-based teleosemantics for animal signals. *Philosophy of Science* 76: 864–75.

Stegmann, U. E. (Ed.) (2013) *Animal Communication Theory: Information and Influence.* Cambridge: Cambridge University Press.

Stegmann, U. (2017). The content of animal signals. In K. Andrews & J. Beck (Eds.), *The Routledge Handbook of Philosophy of Animal Minds* (pp. 324–332). New York, NY: Routledge.

Suzuki, T. N. (2024). Animal linguistics. *Annual Review of Ecology, Evolution, and Systematics*, 55, 205–226.

Townsend, S. W., & Manser, M. B. (2013). Functionally referential communication in mammals: The past, present and the future. *Ethology*, 119, 1–11.

Wheeler, B. C., & Fischer, J. (2012). Functionally referential signals: A promising paradigm whose time has passed. *Evolutionary Anthropology: Issues, News, and Reviews*, 21, 195–205.

Zuberbühler, K. (2020). Syntax and compositionality in animal communication. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375, 20190062.