

Luck and the Value of Communication

Megan Hyska

megan.hyska@northwestern.edu

in *Synthese*, 201(96), 2023. Please cite published version.

Abstract

Those in the Gricean tradition take it that successful human communication features an audience who not only arrives at the intended content of the signal, but also recognizes the speaker's intention that they do so. Some in this tradition have also argued that there are yet further conditions on communicative success, which rule out the possibility of communicating by luck. Supposing that both intention-recognition and some sort of anti-luck condition are correctly included in an analysis of human communication, this article asks what the *value* of events satisfying these conditions is. I present a puzzle concerning the value of intention-recognition which is analogous to the Meno Problem in epistemology, but ultimately argue that this puzzle is solveable: the signaling-relevant value of intention recognition can be vindicated. However, I argue that the version of this puzzle that concerns the further proposed luck-proofing conditions on communication can not be answered. I argue therefore that communication, as analyzed by many, is no more valuable qua signal than a proper subset of its conditions. Human communication is then not a uniquely valuable signaling event.

Acknowledgments:

For their help with the development of this paper, I would like to thank Josh Dever, Ray Buchanan, Henry Schiller, and the audiences at the 2019 Ontario Meaning Workshop, the 2019 Mentoring Workshop for Pre-Tenure Women in Philosophy, and the 2020 Southern Society for Philosophy and Psychology Conference. I would also like to thank my anonymous reviewers, whose feedback made the paper immeasurably better.

1 Introduction

Humans are far from being the only animals to engage in signaling; vervet monkeys deploy distinct calls to warn one another of predators from the ground and from the air (Cheney and Seyfarth 1990; Ducheminsky et al. 2014); octopi raise their mantles and extend their tentacles to intimidate would-be aggressors (Godfrey-Smith 2016); honey bees use a “dance language” to indicate the location of food to one another (von Frisch 1967). But the way that humans typically signal to one another is often taken to be distinctive, featuring very large repertoires of signs whose combinations feature a compositional semantics and are generated by a recursive syntax. Nor is this where the distinctiveness of human signaling ends; those in the Gricean tradition take it that human signaling is also typically mediated by complex inferences about each other’s mental states, where the theory of mind required in order to do this belongs, many contend, to us alone¹. If the Griceans are correct, and paradigmatic human communication does require exceptional additions like mind-reading, it is natural to suppose that this contributes to human communication being somehow better *as a type of signaling*.

This paper presents two puzzles which challenge this supposition, the first of which I offer a solution to, and the latter of which I suggest may be interestingly unsolvable. The puzzles arise for anyone who accepts the following premises, as many working in the Gricean tradition now will²:

¹There is a massive amount of literature on the question of non-human animals’ capacity for metarepresentation; a relatively recent review of this literature is Burge (2018). Important recent philosophical work on the relationship between human and non-human communication includes Armstrong (2021); Bar-On (2013a,b).

²Strictly speaking, the first puzzle is aptly posed even to someone who accepts just the second and fourth of these premises, but for the solution that I offer to be available, one must also accept the first (intentionalism). In any event, those who accept that signaling success requires intention recognition will

Intentionalism: What an instance of paradigmatic human signaling means is (at least partly) determined by what the signaler intends.

Signaling success requires intention attribution: for an instance of paradigmatic human signaling to succeed, the receiver must arrive at the content the signaler intended to get across and moreover see that the signaler intended her to do this.

Signaling success is luck-sensitive: for an instance of paradigmatic human signaling to succeed, the receiver's arrival at the content intended by the signaler can't be a matter of mere luck.

Human communication is one sort of signaling transaction among others: while humans may characteristically communicate via transactions accurately characterized by the above premises, there are other phenomena in the natural world that are aptly described as signals/ signaling, which don't satisfy these conditions.

For the sort of theorist who accepts these premises, the following is a natural question: what is so special about intention-recognition, or the avoidance of luck, such that they feature in the standard of successful communication that we hold one another to? In other words, why are transactions in which the audience meets this standard— successful human communication³— uniquely valuable signaling events? The two puzzles that I raise for communication concern variations on just this sort of question.

likewise generally accept intentionalism. The second puzzle is aptly posed to anyone who accepts even just the third and fourth premises below.

³While I will speak, at times, of “successful communication” and “successful understanding”, I take this to be strictly speaking redundant, as I treat both “communicate” and “understand” as success terms. Moreover, in what follows, the term “communication” is intended to be understood as referring to the kind of signaling transaction that humans paradigmatically engage in.

The subject matter of this paper is the value of communication, but much of our discussion will focus specifically on the receiver’s half of a successful communicative transaction, which we refer to as *utterance-understanding*. In §2 I begin by presenting in a bit more detail the analysis that those in the Gricean tradition give of utterance-understanding. In §3 we’re then positioned to introduce the first puzzle for the signaling-relevant value of such utterance-understanding and for the variety of communication that requires it. I argue that this puzzle is solveable, and in so arguing I offer a novel argument for the value of the intention-recognition condition on successful communication. In §4 I then move on to a discussion of the literature suggesting that successful utterance-understanding is undermined by luck, and of the luck-proofing conditions that have been, or might be, suggested in order to provide an analysis commensurate with this observation. In §5 I then present a puzzle concerning the value of such anti-luck conditions, and conclude that this puzzle isn’t solveable.

2 What is Utterance-Understanding?

In order to get to our puzzles about the value of utterance-understanding, and with it communication, we first need attend to the Gricean picture of what successful utterance-understanding *is*. According to the broadly Gricean⁴ analysis of utterance-understanding, the receiver of a signal understands it just in case she satisfies the following three conditions:

ATTITUDE TO CONTENT: as a result of the signal, she comes to bear some

⁴This presentation is not precisely Grice’s own, as it borrows language from the subsequent tradition he inspired, notably Sperber and Wilson (1986, 2015) who, with others in the cognitive sciences, often refer to the sort of communication we’re targeting here as “ostensive-inferential communication”. These theorists are concerned with analyzing human communication in general, rather than “speaker meaning” in particular, as Grice was.

attitude, α , to some content, p .

ACCORDANCE WITH SPEAKER INTENTION: p is (consistent with) the content the speaker intended to get across with the signal.

ATTRIBUTION OF SPEAKER INTENTIONS: it is evident to her that the speaker intended the signal to cause her to bear α to (some content appropriately related to) p .

Call these conditions ATTITUDE, ACCORDANCE, and ATTRIBUTION, for short.

ATTITUDE gets at the condition which, on the original Gricean formulation (Grice 1957), would require the receiver to come to *believe* the relevant content. However, the subsequent literature is well-populated by reasons to resist the idea that belief is the right attitude to specify here, and by Grice (1969) he talks merely of the speaker creating a desired “response” in the audience. The question of precisely which attitude is the right one here will be irrelevant for our purposes, and for this reason we abstract away from the matter, leaving α as a stand-in. Particularly in cases of non-declarative utterances, it is likewise not clear that a proposition is even the sort of content that an audience comes to grasp, so we abstract away from this issue as well; while I use ‘ p ’ as a stand-in for the content to which α is borne, I do not mean this to commit us to any particular view of what utterance content is. The above is then compatible with the idea that utterance-understanding may sometimes involve having an appropriate attitude toward *sets* of propositions, if this is indeed the content of interrogatives (Hamblin 1973; Groenendijk and Stokhof 1984), or properties, if this is the content of imperatives (Portner 2004, 2007; Charlow 2014; Starr 2020).

ACCORDANCE says that communication involves some kind of conformity of the object of the audience’s α attitude to what the speaker wanted to express in the first place. Included in ACCORDANCE is a parenthetical hedge against the possibility that communication can occur even where there is no one particular content that the speaker and audience both

have in mind⁵; if this causes any consternation, the reader can suppose that the relation between the propositions targeted by speaker and audience is just that of identity.

And ATTRIBUTION rules out cases where a receiver comes to bear α to p but doesn't see that this is what the signaler intended⁶.

While those in the Gricean tradition would agree that the above are necessary conditions on successful utterance-understanding, they are also widely taken not to be jointly sufficient—counterexamples to their sufficiency will be discussed in §4 below, and are what will lead us to our second puzzle. Already though, an interesting question about the value of utterance-understanding can be presented which focuses on the value of ATTRIBUTION.

⁵Bezuidenhout (1997) for instance suggested that signalers and receivers will often represent objects under different modes of presentation, and that this is entirely compatible with successful communication, arguing that “we need recognize only speaker-relative content and listener-relative content and a relation of similarity holding between these two contents” (1997, 198). Buchanan (2010) argues that, for linguistic signals involving quantification, there might be multiple ways of restricting the domain of this quantification which yield different propositions, but which a signaler would find it equally acceptable for the receiver to arrive at. He suggests in fact that signaler's informative intentions might be aimed at something more like a “propositional template” than a specific proposition. For further discussion of the idea that communication doesn't require coordination on identical contents, see also Carston (2002) and Onofri (2018, 2019).

⁶Note that ATTRIBUTION doesn't contain what would be needed in order to restrict cases of successful communication to those of successful *speaker meaning*, as articulated in Grice (1957). Speaker meaning requires that the receiver bear α to p *because of* the recognition that the speaker intends this, whereas the above allows that the receiver merely see that the speaker did intend this, without requiring that this stand in a causal relation to her bearing α to p

3 The First Puzzle: What is the Value of Intention-Attribution?

Supposing that the above has made some headway in correctly saying what utterance-understanding (and thus communication) *is*, it remains to ask why we should care about the state that satisfies the above conditions. In asking this question we turn from matters of analysis to matters of value. Why is a signaling event that satisfies ATTITUDE, ACCORDANCE, *and* ATTRIBUTION better than one that merely satisfies ATTITUDE and ACCORDANCE? That is, what is the added value of the receiver attributing an informative intention to the signaler over and above the value of the receiver having arrived at an attitude toward a proposition consistent with what the speaker intended?

Here's a natural first answer: satisfying ATTRIBUTION makes satisfying ATTITUDE and ACCORDANCE more likely. It is often claimed, after all, that intention-inferencing is what enables the complexity and efficiency of human signaling acts. But immediately apparent is that, while an audience's recognizing the speaker's intentions might be *conducive* to her entertaining the content the speaker had in mind (i.e. to satisfying ATTITUDE and ACCORDANCE), this instrumental value will not be able to vindicate the ATTRIBUTION condition. That's because, if the audience already satisfies ATTITUDE and ACCORDANCE, then there's no need to stack further conditions on top that will probabilify their satisfaction—they're already satisfied! This is what is sometimes called a *swamping problem*: where we stipulate that ATTITUDE and ACCORDANCE are already satisfied, the value of ATTRIBUTION is "swamped" by the already-secured value of the other two conditions.

If we want to vindicate the value of ATTRIBUTION over and above that of ATTITUDE and ACCORDANCE, we'll need to come up with some other story about why it is valuable, one that doesn't rely on its being instrumental in satisfying the other conditions. Maybe, for instance, in tracking the signaler's intentions, one has secured something that is valuable for the larger task of modeling conspecifics' mental states and thus predicting and explaining their actions down the line. This may well be true. But here is the next worry: this story

wouldn't suggest that *ATTRIBUTION* contributed any further value that was relevant to the task of *signaling* per se.

The point here is a delicate, but not unfamiliar, one: we can adjudicate the value of tokens relative to different domains of value, and certain traits may contribute to the token's value in one domain of evaluation but not in another. How well a person runs would be irrelevant to ranking their ability qua weightlifter, for instance. And how well an interaction augments the receiver's future behavior-predicting abilities is, I contend, irrelevant to assessing its value qua signaling event.

This makes sense so long as we regard "signaling" as a distinct domain of value—but what would this mean, exactly? The first step in characterizing a domain of value is an axiology that enumerates what is intrinsically valuable in that domain, or, if one likes, what the central *goal* of that domain is. So what, in the signaling domain, is intrinsically valuable? It won't do to characterize this goal in a way that is narrowly appropriate only to the sort of communication that (those in the Gricean tradition think) is distinctive of humans, since such communication is meant to be only one event-type within this domain. If one takes a naturalistic view of signaling, as we have committed to doing, then the goal of signaling per se can't be specified in a way that excludes consideration of the way that non-human animals can fill either sender or receiver nodes of a signaling transaction. Indeed, on the naturalistic view even non-sentient entities, like plants or storm clouds, can stand in the signaler node of a signaling transaction, and any creature capable of responsiveness to its environment can stand in the receiver node. So the goal of signaling, per se, can not be to induce a mental state, such as an attitude toward some content, even if a particular *kind* of signaling (e.g. human communication) does require this. Finally, the goal of signaling per se can not be specified in a way that only makes sense for declarative utterances. There are those who have argued that knowledge is the essential goal of communication (e.g. Evans 1982; Heck 1995; Peet 2019) and while I will leave it open whether this might be apt when it comes to human declarative communication specifically, I take it is not plausible as an

account of the distinctive goal of signaling generally. Questions and commands, after all, won't as a rule give rise to worldly knowledge⁷ in their receivers— while it's plausible that questions do distinctively serve a goal related to inquiry, which is a close relative to the goal of knowledge (see Friedman 2020), this case can not be made for commands— when someone says “shut the door” they are concerned with a world-mind direction of fit, not the distinctively epistemic mind-world direction. It therefore is not plausible that knowledge production is the goal of signaling per se.

What then is intrinsically valuable within the signaling domain? Familiar from the literature advancing naturalistic perspectives in epistemology, is a way of talking about signals as instigating the “flow,” “spread,” or “transmission” of information between systems whose states bear some (often asymmetric) dependence relation to one another (e.g. Lewis 1969; Dretske 1981; Millikan 1984; Skyrms 2010). Regarding the familiar questions that this characterization raises—what exactly is information? What is it for a system to bear some information? How are systems to be individuated? And is the dependence we're talking about here causal, or else a matter of metaphysical or logical necessitation, or something else?— answers have of course varied. But the discrepancies between particular naturalistic accounts will be unimportant for our purposes: I think we can say something helpful, even while leaving terms like *system*, *information*, and *dependence* less than fully interpreted. We will proceed with the bare idea that what is fundamentally valuable within the domain of signaling is the *spread of information*, where this minimally involves information present in System 1 coming to be present in System 2 as a consequence of some specific relation of dependence between the two⁸. And we will construe “information” broadly, to include

⁷Of course questions and commands may still give rise to some knowledge, e.g. that a signal occurred, or even that the signal meant what it did. But this is knowledge about the signal itself, not about the world, and I take it is not the sort of knowledge that a theorist like Evans (1982) thinks communication aims at.

⁸Note that ATTITUDE already builds in a relation of dependence, because it requires that the receiver

not just propositional contents but also alternative semantic objects, so as to include the contents of, say, questions and commands⁹.

The foregoing positions us to get back to our puzzle. This puzzle specifically concerns whether paradigmatic human communication is uniquely valuable qua signaling event. And it can now be summarized as a dilemma: either *ATTRIBUTION* derives its value from the fact that it is conducive to satisfying *ATTITUDE* and *ACCORDANCE*, in which case its value is swamped, or else it derives its value from something unrelated to information transfer—and so irrelevant to signaling as a domain of value. Either horn—the swamping problem or the domain irrelevance problem—would entail that *ATTRIBUTION* added no signaling-relevant value over and above that of the other conditions, and thus that, if the Griceans are right about what successful human communication requires, human communication is no more valuable as a signaling transaction than other events satisfying a proper subset of its conditions.

This dilemma about the value of communication is a novel one, but its structure is not new: it is precisely analogous to a problem that has historically been raised for the value of bear α to p because of the signal.

⁹What about signals that don't seem to provide new information, even of the kind carried by questions or commands? Consider what Malinowski, thinking of “free, aimless social intercourse” such as comments on weather, called “phatic communion” (Malinowski 1936, 313)—does such intercourse initiate any sort of information transfer? On its face it may not seem like it, since comments about the weather, say, are often already known to be true. Due to lack of space, I will simply gesture at two options for pushing back against this conclusion: first, one might engage in further interpretation of the notion information transfer so that cases where System 1 brings about a heightened *salience* of a piece of information in System 2 counts, perhaps drawing upon the discussion of such cases in Sperber and Wilson (2015). Secondly, in the spirit of Malinowski's own view, we might insist that phatic communion *does* pass on some sort of novel information having to do with affirmation of the addressee's continued personhood in the community.

knowledge, where the question has been what value justification could add over and above the value of true belief: if justification merely makes true belief more likely, then its value is swamped where true belief is already secured. But if the value of justification doesn't have *something* to do with true belief, it's hard to say why it would have any distinctively *epistemic* value. This puzzle dates back to Plato's *Meno*, and has been updated in ways that take aim at a number of different contemporary theories of justification (see e.g. Zagzebski 2004; Kvanvig 2003).

In any case, whatever may be the prospects for answering the epistemic versions of the puzzle, I want to present a method for solving our puzzle concerning the signaling-relevant value of *ATTRIBUTION*. *ATTRIBUTION*, one might argue, does bear a relation to information transfer, just not an instrumental relation. Rather, it represents a kind of *respect for the goal of information transfer* on the receiver's part¹⁰.

This strategy for solving the puzzle implicitly raises a fundamental question for theories of value: where s_1 is among the states countenanced by a theory's axiology as valuable in itself— we may speak of such a state also as *intrinsically* or *fundamentally* valuable within that domain— what relation to s_1 is sufficient to render some further state also valuable? For any domain of value—be it moral, epistemic, or signaling— we may ask what, in that domain, holds intrinsic value, and what relation to that intrinsic value allows further states to hold extrinsic or derivative value. We are all familiar with the different answers

¹⁰This proposed solution is inspired by the attempt in Kvanvig (2003) to solve the analogous problem concerning knowledge. However, one needn't find Kvanvig's solution to that problem persuasive in order to accept the solution I sketch here to the first value problem concerning communication. For instance, Kvanvig's solution to the value problem for knowledge requires accepting a variety of subjective internalism as the correct standard of epistemic justification, but the solution I give here neither requires this nor makes any analogous move to allow the problem of communication's *value* to put pressure on communication's *analysis*.

that consequentialists and deontologists give to this question in moral philosophy¹¹, but in other domains of value, a consequentialist answer is sometimes prematurely assumed to be correct— as Berker (2013) notes for instance, epistemological theories have almost always said that the right relation to the fundamental epistemic good is the *conducting* relation¹². The strategy I present here, which is inspired by a version of epistemic non-consequentialism recently advocated for by Sylvan (2020), relies on the idea that a conducting relation is at least not the *only* relation to information transfer that allows a state or event to bear extrinsic signaling-relevant value, and that a further such relation is something like that of *respecting*. The notion of respect we aim to get at here is something like Darwall’s “recognition respect,” which “consists in giving appropriate consideration or recognition to some feature of its object in deliberating about what to do” (Darwall 1977, 38). To respect the goal of information transfer is then to bear in mind certain important features of what *constitutes* information transfer when deliberating about which attitudes to form.

We saw that, at first glance, it might look like *ATTRIBUTION* is a substantive suggestion about how to probabilify the satisfaction of *ATTITUDE* and *ACCORDANCE*. Were this true, *ATTRIBUTION* would bear an instrumental relation to *ATTITUDE-ACCORDANCE* and its value would be swamped. But I don’t think this is how *ATTRIBUTION* is best understood. For one thing, *ATTRIBUTION* by itself doesn’t specify that the audience proceeded in a rational or otherwise prescribed way from the signal to a conclusion about the speaker’s communicative intentions, just that she attributed to them *some* communicative intention—so the

¹¹Some deontologists will disagree with this framing, which supposes that the good is prior to the right. However, I take it that there are at least some deontologists who won’t. See Sylvan (2020) for discussion.

¹²Discussion of epistemic consequentialism and its alternatives has abounded in the last few years (see e.g. essays in Grajner and Schmechtig 2016; Ahlstrom-Vij and Dunn 2018). While a rejection of epistemic consequentialism isn’t uncontroversial, nor is it without recent support (e.g. Berker 2013, 2015; Andow 2017; Littlejohn 2018; Snow 2018; Roberts et al. 2018; Sylvan 2020).

satisfaction of *ATTRIBUTION* actually *doesn't* by itself make it more likely that the audience satisfies *ACCORDANCE*, at least.

I think that we should instead, or at least in addition, see *ATTRIBUTION* as requiring that *the audience grasps the constitutive norms of a signaling act*, where a constitutive norm is a rule that defines an activity¹³. And signaling as an activity is defined by an orientation, at the audience node, toward the recovery of the information encoded in the signal. The satisfaction of *ATTRIBUTION* may be seen as part of grasping this norm where one is committed to intentionalism, and so takes the speaker's informative intentions to be *determinative* of what the signal's meaning is. The audience's attributing intentions to the speaker is a way of respecting information transfer itself, in the sense described above: it amounts to letting the determinants of the signal's meaning weigh appropriately in one's deliberation about what attitude to form.

In the communicative case, the audience's allowing what content they arrive at to be dependent on the speaker's intentions is like a would-be knower's acknowledging that having true beliefs is a matter of their mental state's fitting the world; it is just part of understanding the nature of the activity. The speaker's information state (including their intentions), together with the signal itself, is the first half of a signaling transaction, without which the notion of information transfer is incoherent. The foregoing is consistent with the naturalistic paradigm we committed to above, on which signaling acts per se don't require

¹³Some might be resistant to the idea that understanding is rule-guided in any way—so-called pragmatic particularists (Buchanan and Schiller 2021) for instance, will say that audiences understand speakers by making inferences to the best explanation, but that there are no principles more specific than that that guide interpretation. Properly understood though, my claim about signaling having constitutive norms is compatible with the particularists' line—I have said merely that signaling as an activity is defined by an orientation, at the audience node, toward the recovery of the information encoded in the signal. I have said nothing about *how* that recovery is to be executed.

intentions. It's true that grasping the project of signaling doesn't require an acknowledgement of intentions, because signaling can take place with a system incapable of intentional states at the signaler node. But this is correct in the same way it is correct that grasping the project of belief does not require the acknowledgement of mountains; it is correct in general but not when it comes to beliefs *about the height of the world's tallest mountain*, which do of course require such an acknowledgment. If intentionalism is true, then signaling event-types that are intentionally initiated by minded individuals are such that an audience can't respect the constitutive norm of signaling without taking the speaker to have intentions, and treating these intentions as decisive of her success.

If the foregoing is correct, then `ATTRIBUTION` derives its value from the value of achieving information transfer (dodging the domain irrelevance horn) without being merely instrumental in it (dodging the swamping horn). So states satisfying `ATTITUDE`, `ACCORDANCE`, and `ATTRIBUTION` have more signaling-relevant value than those only satisfying `ATTITUDE` and `ACCORDANCE`.

I've presented a first puzzle concerning the value of communication, and offered a solution to it available to those willing to countenance non-consequentialist moves in value theory. But, I will go on to argue next, there are further vexing questions when it comes to vindicating communication as a uniquely valuable signaling event. And to arrive at the second puzzle concerning communication's value, we need first to dip back into its analysis.

4 Communication and Luck

While those in the Gricean tradition agree that `ATTITUDE`, `ACCORDANCE`, and `ATTRIBUTION` are necessary conditions on successful utterance-understanding, they are also widely taken not to be jointly sufficient. To see why, we begin with the literature responding to the following case from Brian Loar (1976):

Smith and Jones are unaware that the man being interviewed on television is

someone they see on the train every morning and about whom, in that latter role, they have just been talking. Smith says ‘He is a stockbroker’, intending to refer to the man on television; Jones takes Smith to be referring to the man on the train. (p. 357)

Loar and subsequent commentators have the intuition that this case doesn’t look like a successful case of communication. Loar brought this datum to bear on the debate between direct reference theorists and Fregeans: that Jones’ arriving at the right referent is insufficient for successful communication, Loar argues, shows that some “‘manner of presentation’ of the referent is, even on referential uses, essential to what is being communicated” (Loar 1976, 357).

But some commentators subsequent to Loar have focused on a different feature of the case: that Jones identifies the right referent seemingly by *luck*. Utterance-understanding is undermined by luck, just as it’s been widely taken since Gettier (1963) that knowledge is¹⁴¹⁵.

¹⁴Here, as throughout, I use “luck” as a pre-theoretically intuitive description of what is going on here, but with no technical conception of luck or, relatedly “deviance,” in mind. Of course, there are certainly kinds of luck that do not undermine communication, just as there are some that do not undermine knowledge (a point made nicely in Pritchard 2005a). It being a matter of luck that the interlocutors met up when they did does not, for instance, seem to undermine the possibility of communication between them.

¹⁵The literature does include at least one suggestion that utterance-understanding *can* be achieved by luck, in Byrne and Thau (1996, 148). Someone without the intuition that luck ever undermines communication does not accept one of the premises that this article began by listing, and so the second puzzle here may just not be a puzzle that applies to them, just as the first puzzle didn’t apply to theorists who reject the necessity of ATTRIBUTION. However, along the lines of what Heck (1996, 155) noted in response to Byrne and Thau, if one is inclined to argue that human declarative communication is essentially aimed at enabling testimonial knowledge and that knowledge itself is underminable by luck, this adds a further

And providing an analysis of communication that rules out lucky cases appears to be challenging in ways very familiar from the post-Gettier epistemology literature. Below I'll give a rundown of some options for luck-proofing the analysis, along with their shortcomings. This discussion will then be in service of our asking: what signaling-relevant value, if any, would the satisfaction of some luck-proofing condition add over and above the value of satisfying ATTITUDE, ACCORDANCE, and CONTRIBUTION?

4.1 Attempts to luck-proof an analysis of communication

Buchanan (2013) charged that Loar's case is not actually evidence against direct reference views about the semantics of singular terms. Rather, per Buchanan, Smith did want Jones to arrive at the information that an object, *a* (where *a* is the man who happens to be both on television and the train), had the property of being a stockbroker—but he wanted him to do so on the basis of the observation that the salient man in the signaling context was the man on the television. Instead, Jones reached the same proposition on the basis of taking the man on the train to be the salient individual.

Buchanan is developing a point also made in Grice (1969), a general gloss on which is that, on top of the standardly required intention to have the audience arrive at some content (the *informative* intention) and the intention to have the audience see that the speaker has this informative intention (the *communicative* intention), the speaker also intends that they arrive at the content on the basis of some particular feature of the signal (e.g. on the basis of the signal being tokened in the presence of the television screen displaying the relevant man, *not* on the basis of the signal being tokened just after discussion of the, coincidentally identical, man on the train). Using a term favored by Schiffer (2017), Buchanan refers to reason, beyond mere intuition, to suspect that communication is undermineable by luck as well. This commitment about the purpose of human declarative communication is one I take no stance on here but, as discussed in §3, it is compatible with my account.

these features as *inference base features*, or *ib-features* ¹⁶.

We should observe that this ib-feature requirement (IBFR) is related to the ‘no-false-lemmas’ (NFL) response to the Gettier problem (see e.g. Armstrong 1973; Clark 1963). Both suggest that a satisfactory analysis of their respective target phenomena can be reached by requiring that the *steps of inference* involved in each, not just these inferences’ *outcomes*, satisfy a fit condition: in the case of knowledge that of fitting the world (i.e. truth), and in the case of utterance-understanding, that of fitting the speaker’s intentions. What’s more, the IBFR view is vulnerable to objections closely related to those that plague the NFL condition on knowledge.

Peet (2017) points out that the IBFR, while working nicely to exclude some lucky cases, still admits others, since “we can always construct a case in which the audience recognizes all of the speaker’s intended ib-features, and follows the intended inferential path as far as it goes, but then deviates wildly in such a way that they could easily have failed to recover the correct referent were it not for some coincidence” (Peet 2017, 381). The audience’s inferential path may display the elements of any posited set of ib-features, but so long as that set specifies *an inferential path that is less than maximally demanding*—that is, so long as it stops short of specifying an exact and exhaustive set of inferential steps that the audience is to pass through—it will not be enough to insulate them against a kind of luck that undermines utterance-understanding. And further, the IBFR *can’t* specify an exact and exhaustive set of inferential steps, it is argued, because speakers don’t typically have such maximally demanding ib-feature intentions ¹⁷.

In short, the most demanding IBFR consistent with the fact that speakers have less than maximally specific intentions still won’t provide a condition which, together with the stan-

¹⁶Note that Buchanan doesn’t actually claim that the addition of the ib-feature requirement makes for a jointly sufficient set of conditions on utterance-understanding, just that it is a further necessary one.

¹⁷Peet (2017) makes this argument with reference to a case from Heck (2014).

standard conditions on successful utterance-understanding, makes for a jointly sufficient set. Bearing this in mind, consider that the standard objection against the NFL analysis of knowledge is that there are Gettier-style cases that don't involve the would-be knower's passing through any false beliefs on their way to their conclusion (Clark 1963; Goldman 1976). The NFL condition puts some constraints on how the would-be knower arrives at their conclusion but, even taken together with a separate justification condition, is not maximally specific about how they reach any given belief, and leaves room for ways of doing so that are incompatible with knowledge. IBFR and NFL flow from the same instinct applied to different target phenomena, and they fail in neatly analogous ways.

Having seen that our analysis of utterance-understanding can't be luck-proofed by the addition of an analog to the No False Lemmas Condition, it would be natural to wonder if analogs to the modal conditions on knowledge historically proffered in the epistemology literature might do any better. Here is a communicative variant of one of these, the sensitivity condition (Nozick 1981) plus the original epistemic version, for reference:

SENSITIVITY_{know}:if p were false, then S would not believe that p.

SENSITIVITY_{com}:if p were not (compatible with) the content intended by the speaker, then S would not have arrived at p.¹⁸

On first glance, SENSITIVITY_{know} had the resources to deal with Gettier variants like Goldman (1976)'s False Barn case. In this famous case, an agent, Henry, has a true justified belief that the structure in front of him is a barn, but, because he is in an area otherwise mostly populated by fake barns which he might just as easily have judged incorrectly to be a real barn, he intuitively seems not to have knowledge. After all, if, as is surely the case

¹⁸This proposal is similar in spirit, if not in exact letter, to some modal conditions on communication suggested by Peet (2019), though Peet takes himself to be offering conditions on "knowledge-yielding communication" rather than communication per se, as we are.

in some nearby possible world, the structure in front of Henry had been a barn facade, it seems like he would still have believed it was a barn; actual Henry's lack of knowledge is explained then by a failure to satisfy $\text{SENSITIVITY}_{\textit{know}}$.

Similarly, it might on first glance look like $\text{SENSITIVITY}_{\textit{com}}$ can help us deal with Loar's case, above. The issue in that case, we might be inclined to say, is that we can easily imagine a nearby world where Jones arrived at the referent (and so the content) that he did but where this wasn't what Smith had in mind (i.e. a world in which the man on the train, even if resembling the man on the television, was not the same person).

However, consider now a variation of Loar's case in which the train conductor, a big fan of this (alleged) stockbroker character, will allow only him to fill that particular seat on the train. Where this is the case, had the relevant man (the one who is in fact on television) not been sitting next to Smith and Jones, Jones also would not have identified the individual in that seat as the referent of Smith's utterance, because there would have been no individual there at all. So Jones' half of the signaling transaction in fact satisfies $\text{SENSITIVITY}_{\textit{com}}$ —nonetheless, the transaction fails to look like successful communication to at least the degree that the original Loar case did. This, incidentally, is a counterexample to $\text{SENSITIVITY}_{\textit{com}}$ that is directly analogous to the classic objection to $\text{SENSITIVITY}_{\textit{know}}$, from Kripke (2011).

The case against the analog of the epistemic safety condition (Sosa 1999) is even more straightforward:

$\text{SAFETY}_{\textit{know}}$: In all nearby worlds where S believes that p, p is not false.

$\text{SAFETY}_{\textit{com}}$: In all nearby worlds where S arrives at the proposition p, p is not inconsistent with what the speaker intended.

The standard objection to $\text{SAFETY}_{\textit{know}}$ is to ask which worlds count as nearby; unless there is a non-arbitrary and non-circular way of saying what standard of proximity is in question,

the condition hasn't given us a helpful analysis. And this objection persists verbatim except for the substitution of "understanding" for "knowledge" in the communicative case.

To review, we've seen that utterance-understanding is vulnerable to being undermined by luck. We've also seen some of the attempts to luck-proof the analysis of utterance-understanding, along with their shortcomings.

What we've neglected is consideration of one popular response to the foregoing issues in the analysis of knowledge. For the last twenty years, a popular way out of the frustrations of post-Gettier knowledge analysis has been the Knowledge-First approach (Williamson 2000). On this approach, it is a mistake to think that knowledge can be analyzed into more basic conditions, and so a mistake to think that the value of knowledge is derived from the value of such conditions. I'm not sure, however, that there is a viable "Utterance-Understanding-First" position to which the foregoing frustrations with the task of luck-proofing the analysis of utterance-understanding might drive us. One reason is that I think this strategy is in tension with a commitment to metaphysical intentionalism about meaning (i.e. the view meaning is grounded by or otherwise metaphysically derived from intentions), which many of those who accept a broadly Gricean analysis of communication will endorse. This is because, on this analysis, what a speaker intends is to bring about a certain kind of response in the audience—the response we're calling utterance-understanding. But if utterance-understanding were deemed unanalyzeable, questions might arise as to whether the intentionalist had really reduced meaning to intentions, since these intentions would embed reference to the further unanalyzeable semantic notion, utterance-understanding.

5 The Second Puzzle: What is the Value of an Anti-luck Condition?

What we turn now to considering, for the various candidate anti-luck conditions, is not whether their addition makes for a jointly sufficient analysis of utterance-understanding

(as we've seen, none clearly do), but rather whether a state that satisfies them will be a more valuable one. In other words, our second puzzle asks what signaling-relevant value any anti-luck condition would add, over and above the value of ATTITUDE, ACCORDANCE, and CONTRIBUTION.

Consider the IBFR; why would arriving at the intended content on the basis of the intended features matter from the point of view of signaling value? As before, it won't do to appeal to the satisfaction of this condition making the satisfaction of the others more probable, as this runs straight into the swamping problem; and it won't do to allow it to have a value that isn't somehow related to the one distinctive of the signaling domain, which is to spread information. Now at first glance it may seem like there is an escape hatch here similar to the one we just used to vindicate CONTRIBUTION; allowing one's response, as an audience, to depend on speaker intentions is after all something we regarded as valuable in a way that dodged both horns of the dilemma. But the intentions invoked in IBFR aren't the right kind. The reason that acknowledging the speaker's intention that the audience arrive at the proposition p counted as respecting the goals of the signaling activity was that this intention—the informative intention—was what was determinative of the meaning of the signal. But the speaker's other intentions concerning *the way* the receiver is to arrive at p are not determinative of the content of the signal; so a condition like IBFR can't be vindicated this way¹⁹.

Consider next the modal conditions like SENSITIVITY_{com} and SAFETY_{com}. If what we care about, in an assessment of signaling-relevant value, is the *actual* spread of information, why should we care about whether this spread also takes place in the next world over? Certainly,

¹⁹This conclusion might be troubling for someone who thinks that IBFR is a plausible condition on communication, even if one doesn't think that its addition makes for a jointly sufficient set of conditions. Note that I am offering no argument against IBFR's being part of the analysis of utterance-understanding—I remain agnostic on this issue. I'm only alleging that it contributes no further *signaling-relevant value*.

neither of the proposed modal conditions represents any further individual respect for signaling success— as will be discussed in §5.1.2 below, satisfying a modal condition does not appear to be under agential control— so they can't be vindicated in the way that *ATTRIBUTION* was.

One thought might be that a certain modal robustness of the correlation between what the signaler intended and what the audience arrived at could just be constitutive of information transfer. It's not unheard of to build a modally robust correlation between sign and signified into what it means to *be* a signal (as in e.g. Dretske 1981), though that is different from building modal demands into what it is for information to have been *transferred* by a signal. Notably, this suggestion would mean that the modal condition applied not merely to paradigmatic human communication, but to all instances of signaling. Our modal conditions as articulated above would not then be exactly fit for purpose, since they allude explicitly to a signaler with informative intentions— something that rain clouds and rabbit footprints for instance don't possess. No trouble— one wanting to run with this suggestion might tweak, say, *SENSITIVITY_{com}* to make it more general:

*SENSITIVITY_{com}**:if p were not (compatible with) the content of the signal, then
S would not have arrived at p.

But this revised condition isn't plausibly a necessary condition on information transfer, because this would falsely entail that information transfer in the natural world is rare. Suppose that a predator creeping through the brush inadvertently causes the leaves to rustle, and its prey, a rabbit, bolts. Rabbits are skittish— the rabbit would have bolted even if it was the wind that rustled the leaves. And if the wind had rustled the leaves, the rustling of the leaves would not have constituted a sign that there was a predator nearby. The rabbit would have arrived at the information “there's a predator nearby, (better run)” even if that's not what the rustling signal meant. So if *SENSITIVITY_{com}** were true, then even in the case where the rustling *is* caused by the predator, we couldn't

explain the rabbit's bolting by saying that the rabbit had received this information from its environment. And this is problematic, for reasons that Millikan (2004) rehearses in her critique of Dretske (1981)²⁰: in order to provide "a theory of natural information that will help to explain how real animals manage to obtain useful information," (2004, 34) it won't do to give an account of which makes information transfer much rarer than animals' apparent responsiveness to their environments. All this to say, the attempt to vindicate the modal conditions' signaling-relevant value by maintaining that information transfer per se constitutively requires their satisfaction is not plausible from the point of view of analysis. But this leaves me, at least, without any further ideas for how the modal conditions' signaling-relevant value could be defended.

This second puzzle concerning the value of communication clearly involves a complexity that wasn't present in the case of our first puzzle: the previous section has made it clear that the literature has not yet offered a fully satisfactory anti-luck condition. And how can we assess the value of a condition that we haven't yet identified? Why think that the status of any of the failed anti-luck conditions vis-à-vis signaling-relevant value is dispositive concerning the value of the not-yet-discovered anti-luck condition that will feature in a fully sufficient analysis of utterance-understanding?

Kvanvig (2003) drew a pessimistic conclusion concerning the value of an anti-luck condition on knowledge very much like the pessimistic conclusion I have suggested concerning the value of an anti-luck condition on communication. Acknowledging that in the case of knowledge too, it's not clear that a successful luck-proofing condition has been found and so the object of critique is in a sense unknown, Kvanvig observes that, in the succession of candidate conditions that have been put forward:

²⁰Millikan here is concerned that Dretske's view makes natural signs themselves too rare— but her argument extends nicely to showing why we shouldn't want information transfer *from* these signs to be too rare either.

Counterexamples to the initial formulation of the approach force alterations in the approach, and the alterations are guided exclusively by concern over the nature of knowledge, resulting in emendations of the original suggestion that appear entirely ad hoc from the point of view focusing on the question of the value of knowledge...When we look carefully at the variety of failed attempts to satisfy the twin desiderata concerning the nature and value of knowledge, we do not find signs of progress. We find, instead, a repeated pattern in which progress with respect to one desideratum is balanced by greater weakness with respect to the other (Kvanvig 2003, 138-139, italics mine).

I suggest analogously that innovations designed to offer a better *analysis* of communication will continue to make no improvement concerning the signaling-relevant *value* of communication. The argument here doesn't pretend to be deductive: that there are other ways in which extant or future anti-luck conditions on utterance-understanding might navigate the dilemma is not impossible, for all that's been said here. But lacking such a candidate, or a demonstration that the mentioned candidates contribute more to the value of the event than has occurred to us here, the remaining conclusion deserves consideration: that whatever the further conditions on utterance-understanding are, they add no signaling-relevant value beyond that already present in ATTITUDE, ACCORDANCE, and ATTRIBUTION. Utterance-understanding then, and communication by extension, is no more valuable than the state satisfying only a proper subset of its conditions. Communication is just not a uniquely valuable signaling achievement.

5.1 Some Objections and Replies

5.1.1 Why not center on the speaker?

In the search for a unique value possessed by communication, we might be tempted to turn to the idea that utterance-understanding uniquely satisfies the speaker's intentions.

Whereas when we attempt to reach true beliefs about the world, our doing so is not prefigured by any intention by the world to be known, what the audience side of a successful communicative transaction achieves is that which was intended by the speaker. And so an alternative response to the value question for communication emerges where an analogous response wouldn't have been possible in the case of knowledge: perhaps the unique value of utterance-understanding lies in the fact that it uniquely satisfies speaker intentions. I don't embrace this possibility.

The first thing to note is that, just as no fully successful audience-side anti-luck criterion has been articulated by the literature, nor is there a promising speaker-side anti-luck condition (i.e. one articulating a speaker's intention that the signal be transmitted non-luckily) on offer. But even supposing that such a speaker-side condition were articulated, and supposing that it was satisfactory to allow the audience-side event to inherit its value from the value of speaker intentions²¹, we would only have referred the second value problem over to the speaker's side— it would still remain to be explained why it was valuable that the speaker had this anti-luck intention; and if the intention's value couldn't navigate our now-familiar dilemma, nor could the value of the audience's satisfying it.

5.1.2 Why do we aim for communication?

It's natural to ask why, if utterance-understanding, and so communication, is not a uniquely valuable kind of signaling event, it is the one we aim for—or at least take ourselves to be aiming for—and the standard we hold others to. Roughly I think: a) it's not totally clear that we really do aim for, or hold others to the standard of, utterance-understanding rather than merely ATTITUDE, ACCORDANCE, and CONTRIBUTION and b) even if we do, there are explanations of this fact available that don't rely on communication being uniquely valuable

²¹One might have further doubts on the grounds that it is a very odd story on which value flows from intentions to their satisfactions, rather than vice versa.

qua signaling event.

Part a) of my response takes some further inspiration from epistemology, specifically Kaplan (1985) where he notes that “given our conception of knowledge, all we can do by way of seeking knowledge is seek justified belief and hope that this justified belief will satisfy whatever other conditions a justified belief must satisfy in order to qualify as knowledge” (361). That is, particularly if the successful fourth condition on knowledge turns out to be a more sophisticated cousin of the modal conditions we’ve seen offered before, it’s really not clear what a person could *do* to aim at knowledge beyond aiming for justified belief.

Imagine now a person who, aiming to be a good understander, is attentive to the intentions of agents in her environment, which involves always trying to be diligently alert to the presence of signals and to the informative intentions that accompany them. We of course might say more about what this diligent alertness will require, and no doubt a lot of it will have to do with being a good epistemic agent (e.g. making good inferences about signaler intentions on the basis of evidence). But all that would be by way of unpacking the actions already listed, rather than enumerating more. And I take it that’s because there aren’t really more to be listed— there’s nothing more we could ask an agent to do, by way of being a good signal audience, beyond trying to infer the signaler’s signaling-relevant attitudes and their contents. This seems to indicate that all that’s encompassed by the standard of utterance-understanding to which we hold ourselves and others, is that we do our darndest to satisfy ATTITUDE, ACCORDANCE, and CONTRIBUTION. Communication per se may require more than this, but that simply suggests that we aren’t really *aiming* at communication, even if we *hope* for it— if utterance-understanding requires the satisfaction of something like SENSITIVITY_{com} or SAFETY_{com}, the gap between satisfying merely the three original conditions and achieving utterance-understanding itself is not one that can be closed by anything under the audience’s agential control.

Part b) of my response starts by noting that nothing I have argued for has strictly ruled

out the possibility that communication’s anti-luck criterion might have extrinsic value in relation to purposes beyond the goals of the signaling domain. For instance, one might revert to the position discussed in §3, that communication’s incompatibility with luck is useful because it allows communication to give rise to a further luck-incompatible state like knowledge²².

Of course, if one accepts the value-skeptical argument not only about communication, but about knowledge too, then this move will raise the related question of why knowledge, if not uniquely epistemically valuable, is what we aim for. Very generally: why would we prefer unlucky states, even when the luck makes no difference concerning any positively articulable value?

Papineau (2019, 2) has suggested that because “the concept of knowledge is both phylogenetically and ontogenetically primitive”, our retaining a practice of privileging this anti-luck concept will be explainable by a story about our cognitive evolution, without being *justifiable* in terms of any value we now have. Pritchard (2005b, 194) and others connect luck aversion to our attachment to the idea of epistemic achievement by the agent, whereas “luck undermines genuine achievement”— even while arguing that a certain kind of luck is inescapable, and so a certain aspiration for agential achievement futile (see also Pritchard 2003, 2005a). I take no position on which of these is the correct explanation, but I do think it is interesting to consider the extent to which these explanations of luck aversion might apply to communication as well— if so, they provide a direct explanation of why we might retain the concept of communication that needn’t route through the premise that the phenomenon has value vis-à-vis the signaling domain, or any other.

²²What *is* incompatible with my view is the idea that knowledge production is the characteristic goal of signaling per se—but this doesn’t mean there aren’t signaling event-types that are knowledge-conducive.

6 Conclusion

If we have offered reason to think that communication has no distinctive value qua signaling transaction, there remains the prospect, which I have not argued against, that it is unique among signaling phenomena in the value that it has vis-à-vis *other* domains of value. For instance, as mentioned above, satisfaction of the ATTRIBUTION condition on communication (as well, perhaps, as the IBFR) is clearly valuable for the purpose of mind-reading and all the social facility it enables. And further anti-luck conditions, as we have said, may be what allows communication to give rise to the further luck-proof state, knowledge, which may or may not be uniquely epistemically valuable itself.

While not then clearly motivating an abandonment of communication as a target phenomenon of lay significance and theoretical interest, our conclusion remains significant. If human communication really has among its success conditions any of those that the luck-proofing literature has enumerated, then caring about communication per se can not be justified merely by an interest in information transfer. If communication is special among signaling events, it is perhaps special because of its connections to knowledge production or to human sociality— the value of which register only outside the domain of signaling per se.

References

- Ahlstrom-Vij, K. and Dunn, J., editors (2018). *Epistemic Consequentialism*. Oxford University Press.
- Andow, J. (2017). Do non-philosophers think epistemic consequentialism is counterintuitive? *Synthese*, 194(7):2631–2643.
- Armstrong, D. M. (1973). *Belief, Truth, and Knowledge*. Cambridge University Press.
- Armstrong, J. (2021). Communication before communicative intentions. *Noûs*, pages 1–25.
- Bar-On, D. (2013a). Expressive Communication and Continuity Skepticism. *The Journal of Philosophy*, 110(6):293–330.
- Bar-On, D. (2013b). Origins of Meaning: Must We ‘Go Gricean’? *Mind & Language*, 28(3):342–375.
- Berker, S. (2013). The rejection of epistemic consequentialism. *Philosophical Issues*, 23(1):363–387.
- Berker, S. (2015). Reply to goldman: Cutting up the one to save the five in epistemology. *Episteme*, 12(2):145–53.
- Bezuidenhout, A. (1997). The communication of de re thoughts. *Noûs*, 31(2):197–225.
- Buchanan, R. (2010). A puzzle about meaning and communication. *Noûs*, 44(2):340–371.
- Buchanan, R. (2013). Reference, understanding, and communication. *Australasian Journal of Philosophy*, 92(1):55–70.
- Buchanan, R. and Schiller, H. I. (2021). Pragmatic particularism. *Philosophy and Phenomenological Research*, pages 1–17.
- Burge, T. (2018). Do infants and nonhuman animals attribute mental states? *Psychological Review*, 125(3):409–434.

- Byrne, A. and Thau, M. (1996). In defence of the hybrid view. *Mind*, 105(417):139–149.
- Carston, R. (2002). *Thoughts and Utterances: the Pragmatics of Explicit Communication*. Blackwell.
- Charlow, N. (2014). Logic and semantics for imperatives. *Journal of Philosophical Logic*, 43:617–664.
- Cheney, D. L. and Seyfarth, R. M. (1990). *How Monkeys See the World: Inside the Mind of Another Species*. University of Chicago Press.
- Clark, M. (1963). Knowledge and grounds. a comment on mr. gettier’s paper. *Analysis*, 24(2):46–48.
- Darwall, S. (1977). Two kinds of respect. *Ethics*, 88:36–49.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. MIT Press.
- Ducheminsky, N., Henzi, P., and Barrett, L. (2014). Responses of vervet monkeys in large troops to terrestrial and aerial predator alarm calls. *Behavioral Ecology*, 25(6):1474–1484.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Clarendon Press.
- Friedman, J. (2020). The epistemic and the zetetic. *Philosophical Review*, 129(4):501–536.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6):121–123.
- Godfrey-Smith, P. (2016). *Other Minds: The Octopus, the Sea, and the Deep Origins of Consciousness*. Farrar, Straus & Giroux.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20):771–791.
- Grajner, M. and Schmechtig, P., editors (2016). *Epistemic Reasons, Norms and Goals*. De Gruyter.
- Grice, H. P. (1957). Meaning. *Philosophical Review*, 66(3):377–388.

- Grice, H. P. (1969). Utterer's meaning and intention. *Philosophical Review*, 78(2):147–177.
- Groenendijk, J. and Stokhof, M. (1984). *Studies on the semantics of questions and the pragmatics of answers*. PhD thesis, Institute for Logic, Language & Computation (ILLC), University of Amsterdam.
- Hamblin, C. L. (1973). Questions in montague english. *Foundations of Language*, 10(1):41–53.
- Heck, R. K. (1995). The sense of communication. *Mind*, 104(413):79–106. Originally published under the name “Richard G. Heck, Jr.”.
- Heck, R. K. (1996). Communication and knowledge: Rejoinder to byrne and thau. *Mind*, 105(1996):151–156. Originally published under the name “Richard G. Heck, Jr.”.
- Heck, R. K. (2014). Semantics and context dependence: Towards a strawsonian account. In Burgess, A. and Sherman, B., editors, *Metasemantics: New Essays on the Foundations of Meaning*, pages 327–365. Oxford University Press. Originally published under the name “Richard G. Heck, Jr.”.
- Kaplan, M. (1985). It's not what you know that counts. *The Journal of Philosophy*, 82(7):350–363.
- Kripke, S. (2011). Nozick on knowledge. In *Philosophical Troubles: Collected Papers*, volume 1. Oxford University Press.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lewis, D. (1969). *Convention*. Harvard University Press.
- Littlejohn, C. (2018). The right in the good: A defense of teleological non-consequentialism. In Ahlstrom-Vij, K. and Dunn, J., editors, *Epistemic Consequentialism*, pages 23–47. Oxford University Press.

- Loar, B. (1976). The semantics of singular terms. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 30(6):353–377.
- Malinowski, B. (1936). The problem of meaning in primitive languages. In Ogden, C. K. and Richards, I. A., editors, *The meaning of meaning: a study of influence of language upon thought and of the science of symbolism*, volume Supplement I, pages 296–336. Kegan Paul.
- Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. MIT Press, Cambridge, Massachusetts.
- Millikan, R. (2004). *Varieties of Meaning: The 2002 Jean Nicod Lectures*. MIT Press, Cambridge, Massachusetts.
- Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.
- Onofri, A. (2018). The publicity of thought. *Philosophical Quarterly*, 68(272).
- Onofri, A. (2019). Loar’s puzzle, similarity, and knowledge of reference. *Manuscripta*, 42(2):1–45.
- Papineau, D. (2019). The disvalue of knowledge. *Synthese*, pages 1–22.
- Peet, A. (2017). Referential intentions and communicative luck. *Australasian Journal of Philosophy*, 95(2):379–384.
- Peet, A. (2019). Knowledge-yielding communication. *Philosophical Studies*, 176(12):3303–3327.
- Portner, P. (2004). The semantics of imperatives within a theory of clause types. In Young, R. B., editor, *Proceedings of SALT*, volume 14, pages 235–252. CLC Publications.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15:351–383.

- Pritchard, D. (2003). Virtue epistemology and epistemic luck. *Metaphilosophy*, 34(1/2):106–130.
- Pritchard, D. (2005a). *Epistemic Luck*. Oxford University Press.
- Pritchard, D. (2005b). Scepticism, epistemic luck, and epistemic angst. *Australasian Journal of Philosophy*, 83(2):185–205.
- Roberts, P., Andow, J., and Schmitdtke, K. A. (2018). Lay intuitions about epistemic normativity. *Synthese*, 195(7):3267–3287.
- Schiffer, S. (2017). Intention and convention in the theory of meaning. In Hale, B., Wright, C., and Miller, A., editors, *A Companion to the Philosophy of Language*. John Wiley & Sons, 2nd edition.
- Skyrms, B. (2010). *Signals: Evolution, Learning, & Information*. Oxford University Press.
- Snow, N. (2018). Adaptive misbeliefs, value trade-offs, and epistemic responsibility. In Ahlstrom-Vij, K. and Dunn, J., editors, *Epistemic Consequentialism*, pages 48–69. Oxford University Press.
- Sosa, E. (1999). How to defeat opposition to moore. *Nous (Supplement: Philosophical Perspectives, Epistemology)*, 33(s13):141–153.
- Sperber, D. and Wilson, D. (1986). *Relevance: Communication and Cognition*. Blackwell.
- Sperber, D. and Wilson, D. (2015). Beyond speaker’s meaning. *Croatian Journal of Philosophy*, 15(44):117–149.
- Starr, W. B. (2020). A preference semantics for imperatives. *Semantics and Pragmatics*, 13:1–60.
- Sylvan, K. J. (2020). An epistemic non-consequentialism. *The Philosophical Review*, 129(1):1–51.

von Frisch, K. (1967). *The Dance Language and Orientation of Bees*. Harvard University Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Zagzebski, L. T. (2004). The search for the source of epistemic good. *Metaphilosophy*, 34(1-2):12–28.