Chapter Two:

Designing Communicative Acts

Daniel W. Harris

1 Communication Design, Message Design, and Signal Design

I ONCE HAD THE FOLLOWING CONVERSATION with my two-year-old daughter while we ate breakfast. I was drinking iced hibiscus tea.

OONA: What's *that*?

ME: This is my iced tea. It's like Mommy's iced coffee, but it's tea.

OONA: It's *red*.

ME: Yes, it's bright red.

OONA: Is it ice coffee?

ME: No, it's iced tea. Do you want to try? [I held the straw close to her face.]

OONA: No! *Yucky*!

There was nothing extraordinary about this conversation. I reproduce it here in order to illustrate a few of the many ways in which we normally customize both what we communicate and how we communicate it for our addressees.

First consider the process that I will call "message design"—the process by which we decide *what* to communicate and *to whom*, pairing a message with an addressee. If I had begun the exchange, I might have said any number of things to Oona, depending on my thoughts and goals at that moment. As it happened, Oona started

us off with a question, which narrowed down my options for how to respond. However, the question that Oona asked, like other *wh*-questions, would have required different kinds of answers depending on who had asked it and why. Imagine that I had been asked what was in my glass by an alien anthropologist who knew little of humanity and who had never before seen a beverage. In response, it might have been best for me to say, "This is a drink. Humans consume drinks in order to hydrate our bodies." In response to the same question posed by a tea expert, I might have said, "It's a glass of Bellocq No. 50 Hibiscus tisane, brewed hot and then poured over ice." There are many possible scenarios between these two extremes. If my wife had asked me what I was drinking, I might have assumed that she knew what hibiscus tea was but didn't know we had any, and replied, "Oh it's a hibiscus tea that I bought last week."

None of these answers would have been appropriate for Oona. Two-year-olds are like alien anthropologists in some ways, but Oona already knew that people drink things; she wanted to know *what* I was drinking. However, she would not have known what to do with the details that I might have offered to a tea snob, or even the information that I would have given to my wife. Oona barely possessed the concept of tea, but she knew about iced coffee, because she had been going with my wife to the coffee shop, and so I described my tea by analogy to something on which she had a firmer conceptual grip. Of course, the answer that I gave to Oona would itself have been inappropriate for any of the other three hypothetical addressees: too specific and dependent on unavailable background information for the alien, but patronizingly unspecific for my wife or the tea snob.

Even once I had decided *what* to communicate to Oona, I was still faced with choices about *how* to communicate it. With my message designed, I still had to design a signal with which to convey it. A year earlier, I might have just poured some of my tea into Oona's cup in order to indicate that she should try my tea and find out what it was for herself. But Oona's conversational abilities had come a long way in recent months, and so I spoke. If I had been speaking with a child of one of my Swiss friends, I might have hazarded an attempt at German or French, but Oona doesn't understand those languages, and so I went with English. I avoided vocabulary that I might have used with others but that I doubted she would understand, substituting 'iced coffee' for 'cold brew,' for example. I spoke in a way that it is typical and appropriate to address only to children: slowly, in an upbeat tone of

voice, enunciating my words more than I do with adults. Instead of referring to my wife as 'my wife', or 'Margot', I referred to her as "mommy"—something I would not do when addressing anyone other than my children. Even the nonlinguistic communication that I mixed into our conversation was specific to Oona. It would have been rude to thrust my straw into the face of an adult, or even someone else's child, but the relevant manners don't apply to interactions with my own toddler, who hadn't learned them yet anyway, and who was more likely to understand my offer with the drink brought to her lips.

So, message design is the process of deciding what to communicate to whom, and signal design is the subsequent process of deciding how to convey this message to the addressee. I will sometimes group these two processes together as "communication design"—the process of designing a communicative act.[1] This chapter is about the psychological capacities that we use for communication design. How are we able to do it? How good at it are we, and what are our limitations? And if, as I will argue, good communication design requires the marshaling of considerable cognitive capacities and resources, what benefits make it worthwhile?

My answers to these questions will revolve around the idea that our capacity for communication design is intimately tied to our capacity for intention recognition. By telling Oona that I was drinking iced tea, I intended to get her to believe that I was drinking iced tea, and I intended to achieve this by getting her to recognize what I intended. My communicative intention in this case was a complex plan with both of these smaller intentions as parts. The key to humans' special capacity for communication design, I will argue, is that communicative intentions like this one are normally embedded in even larger and more complex plans that arise from a highly flexible and informationally omnivorous kind of practical reasoning. Message design is the process by which we reason from our broader goals to arrive at a communicative intention to produce a particular effect, such as a belief or an intention, in a certain addressee. In this case, I reasoned from my goal of answering Oona's question to an intention about which belief to impart to her. Signal design is the process by which we reason from a communicative intention to a plan

[1] The most common terminology for this sort of thing in linguistics is "audience design," but it is my impression that linguists use this term in a way that emphasizes what I am calling "signal design" without quite isolating it. Since it is important for my purposes here to discuss both message and signal design while also distinguishing them, I have decided to use my own terminology instead.

to produce a signal with features that will provide our addressee with evidence of what we intend. For example, one thing that I reasoned in this case was that Oona would understand my message if I referred to what I was drinking as "iced tea"—as opposed, say, to "hibiscus tisane."

Communicative intentions are therefore the product of just one step in the communication-design process. But it is a particularly important step, because it is the step at which a communicator pairs a message with an addressee. This pairing is what sets the terms of successful communication, and this explains the Gricean intuition that what we communicatively intend is constitutive of what we mean when we speak or otherwise try to communicate. But the pairing of a message with an addressee within a broader process of practical reasoning is also important because it unlocks the power of communication design itself, allowing us to customize both messages and signals for our addressees. I will argue that our ability to do this supercharges our communicative capacities in a wide variety of ways. Moreover, the resulting communicative superpowers don't merely show up occasionally, as unusual virtuoso performances, but are features of completely normal human communication. Indeed, some of the most ubiquitous and valuable design features of natural languages presuppose that speakers are adept communication designers. Unlocking these communicative superpowers is one of the central benefits of relying on intention recognition as a communicative strategy, and should be thought of as an important part of the explanation of why we do so at all, in spite of the considerable cognitive resources required.

2 Intentions in Practical Reasoning

It should not be surprising to find communicative intentions serving as elements in complex, hierarchically organized plans. A central claim of the most influential theory of intentions—Michael Bratman's (1987) planning theory—is that it is part of the nature of intentions in general to fit into larger plans in this way. We build these plans by means of a kind of practical reasoning that proceeds in stages, each of which aims at adopting a new intention that coheres with and further fleshes out the plans laid down at previous stages.

Each episode of practical reasoning normally arises from questions about how to work out the details of the plans that we have already made. For example, sup-

pose I have decided to throw a dinner party for a few close friends. I now have an intention to do so, but this intention is partial, in the sense that it leaves the details of my dinner party unresolved, and so in order to take action toward my goal, I must address questions of means and implementation. What, for example, should I serve? There are many options. To decide between them would be to adopt a subplan of my prior intention—an intention about how to accomplish it.

We could model this particular choice using the tools of decision theory, calculating the expected utility of each practical option from the degrees to which I take various outcomes to be likely and desirable. For example, suppose that I believe that my friend's husband is disgusted by scallops, but that everyone else would greatly prefer scallops to pasta. If I take it to be sufficiently unlikely that my friend will bring her husband along, and if I am sufficiently tolerant of the small risk that I will wind up serving him food that disgusts him, then it might work out that the expected utility of scallops beats out that of pasta. We could formalize this reasoning by attaching numerical credences and utility values to the possible outcomes and crunching the numbers through a choice function that reflects, among other things, my level of risk tolerance.

One of Bratman's insights was to recognize that these decision-theoretic tools allow us to model only individual stages in broader processes of practical reasoning, and that the options that we consider at each such stage are constrained by the intentions formed as the outcomes of prior stages. The very choice of what to cook for my guests arises in the shadow of my intention to host a dinner party rather than, say, a barbecue. This prior intention, together with my beliefs about what sort of food is appropriate at dinner parties and my other background beliefs and intentions—such as, say, my intention to impress a certain guest who loves sea food—means that I never had to consider veggie burgers as one of the live options in my choice.

My new intention to serve scallops fleshes out my plan, but it in turn raises further questions that will have to be settled by further subplans: Where should I buy the scallops? How should I prepare them? What side dishes and wine will I serve? The intentions that I form in response to each of these questions will raise further questions, until ultimately I am armed with a detailed-enough plan that I can act on it with specific bodily movements. (Of course, some of this reasoning will have to wait until after I've begun acting on the plan's early stages. I won't know

how much to season the scallops until after I've bought them and seen how big they are, for example.)

It is central to the functional role of intentions that they play this role in practical reasoning, as waypoints in the cognitive journey from abstract goals to motor instructions. Intentions can play this role because they are subject to rational pressure that pushes planning agents to develop coherent plans. For example, the drive to stay rational motivates us to treat our intentions as relatively stable commitments, which persist long enough to stay consistent with and constrain downstream decisions. Likewise, it would be irrational to intend some end, to believe that a certain action would be a necessary means to that end, but not to intend the means. If I intend to make scallops, and I believe that the only store with acceptably good scallops closes at 5:00, then I'd better plan to get there by then. Finally, it would be irrational to adopt inconsistent intentions, or intentions that we believe we can't accomplish. If I think I won't be able to pull off making both scallops and a rhubarb tart in the time I have to cook, then I should plan to cook at most one of those things. These requirements push us to avoid constructing complex plans that are too sketchy or incoherent to actually be accomplished.

In order to proceed in a way that has any chance of satisfying these rational requirements, practical reasoning must be a domain-general psychological process. This is to say that its operations must be sensitive, at least in principle, to the contents of any of our beliefs, preferences, and intentions. After all, the requirements place global constraints on how our beliefs and intentions can fit together. With exotic enough beliefs hanging out in the background, there's no telling which two intentions might turn out to be inconsistent by an agent's own lights. An intention to perform only pious acts might come into conflict with an intention to serve pasta, if only one's beliefs have been shaped by a carbohydrate-averse religion. So although we no doubt rely on heuristics to avoid having to consider all background beliefs and intentions, practical reasoning must in principle have access to any beliefs and intentions that might turn out to be relevant.

Hierarchical practical reasoning has several advantages. First, it breaks down complicated, multivariate decisions into smaller, easier-to-consider choices about which cognitively limited agents like us are better able to deliberate. (It would be overwhelming to simultaneously decide on every aspect of a dinner party, from the guest list to the position my hand when sprinkling salt on the side dish.) Second,

the new intentions that I form as a result of upstream choices make downstream choices easier by constraining the options from which we have to choose. (It's easier to choose wine if I already know which food it has to pair with.) Third, insofar as the results of upstream decisions *do* successfully constrain the available options for downstream decisions in a way that is guided by rational requirements, this gives us a way to build complex plans that are coherent, in the sense that the intentions and beliefs that go into them can all be satisfied at the same time. (If all goes well, I won't come home from the store with half of the pasta ingredients and half of the scallops ingredients.) Finally, we often have to begin performing complex actions before we have all of the information needed to decide on how they will conclude, and we can do this by putting the latter decisions downstream in the reasoning process. (For example, I might need to hold off on choosing a specific recipe until I see which kinds of scallops are available at the seafood market.)

In all of these ways, hierarchical planning enables agents with limited information and cognitive resources to construct more complex plans, and to pull off more complex and coordinated actions than we could otherwise manage. The capacity to construct such complex plans opens up the possibility of pursuing certain kinds of goals that we would otherwise have no hope of achieving, and that it would therefore be pointless to intend.

Consider our ability to pursue *abstract* goals, by which I mean goals that would have to be accomplished by many different kinds of particular actions depending on the details of the circumstances. Accomplishing my intention of writing a persuasive book chapter would require moving my fingers on the keyboard in different ways depending on my beliefs about my subject matter, my audience, and my intentions about what to say. The connection between my goal and my finger movements will have to be mediated by very many intentions standing in plan–subplan relations, each of which will require input from a host of beliefs. If I had different intentions and beliefs, I would wind up with different finger movements typing different things. Without the domain-general capacity to intelligently construct complex, coherent plans in the service of such an abstract goal, there would be little point in forming an intention to pursue the goal at all.

Consider also *social* goals, whose realization depends on how others will behave, and which therefore require social coordination (Bratman, 1987, 2014). In order to pursue those goals, I need to be able to construct plans that condition

my own behavior on predictions about the activities of other agents. This requires reasoning from intention to subplan in ways that are sensitive to my beliefs about other agents' beliefs, desires, and intentions, including their beliefs about my own states of mind and future behavior.

Our capacity for hierarchical planning is what allows us to form intentions to pursue complex, abstract, social goals. It does this by allowing us to break down big, hard problems into smaller pieces, and to condition our approach to each chunk of a problem on our solutions to prior chunks and on any relevant available information, including our information about other agents' thoughts and intentions. Hopefully, the relevance of this capacity to communication design is becoming obvious. I will now try to make it more explicit.

3 Communication Design and Communicative Intentions

The only plausible explanation of our capacity for communication design is that it is an application of our capacity for hierarchical practical reasoning, informed in a domain-general way by our beliefs, including our beliefs about our addressees and their states of mind. Consider the process by which I decided how to answer Oona's question about what I was drinking, which bears all the hallmarks of hierarchical practical reasoning:

The design process was *domain-general*: It was conditioned on my beliefs about what I was drinking, about tea, about what Oona wanted to know about, and about her beliefs. If my beliefs about any of these things had been relevantly different, I would have produced a different utterance. And if Oona had asked a different question, there's no telling which of my beliefs would have been needed by the design process.

The design process was *hierarchical*: My decisions about which expressions to utter depended on my intentions about what I wanted to communicate, which in turn depended on my prior intention to answer Oona's question together with my standing intentions to be helpful and understandable to her.

The design process was *responsive to rational requirements*: My specific plan about what to say was successful because it cohered with my other intentions, poli-

cies, and beliefs. If my plan hadn't cohered with my other intentions and beliefs in this way, it would have been less likely to succeed. And if I had had different beliefs and intentions that would have conflicted with this plan (such as, say, a belief that Oona always likes to play guessing games instead of being given direct answers), I would likely have answered her question differently.

The design process unfolded in pursuit of goals that were *abstract* and *social*: As I argued in §1, I could have answered Oona's question, like other *wh*-questions, in many different ways, and which answer was appropriate depended on what I believed about Oona, her goals in asking the question, and what she already knew. Moreover, my intention to answer Oona's question in the first place was formed in pursuit of even more abstract goals, such as the goal of being a helpful parent, and these goals and their connection to my intention to answer her were themselves conditioned on my beliefs about Oona and her state of mind.

Finally, my communicative act was a *complex, coordinated action*: I uttered thirteen words in a specific order, with a specific intonational contour, presupposing some information, asserting other information, and withholding still other information. Each of these features of my utterance was conditioned on my beliefs about Oona. And this is only considering the complexity of a single utterance. If we consider the structure of the broader conversation of which this utterance was just a part, we find evidence of even more ways of packaging information for Oona's benefit, extended over multiple utterances.

These are the lines of thought that convince me to think of communication design as a special case of hierarchical practical reasoning. And from here, there is a clear argument to the conclusion that communication design has to involve at least the first component of a communicative intention. Hierarchical practical reasoning, and therefore communication design, is an iterative process of reasoning from intention from subplan, beginning with relatively abstract intentions and ending with intentions that are specific enough to serve as inputs to motor control. Given the nature of communication design, which involves both message design and signal design, there must be some point in any specific communication-design process that first pairs a message with an addressee. After all, this pairing is the point of message design. And since signal design is the process of deciding how to convey a certain message to a certain addressee, an intention to convey that message to that addressee is a necessary input to the process. So what, exactly, is an

intention that pairs a message with an addressee? It is an intention to produce a certain psychological state, such as a belief, in a certain agent. This is the first component of a communicative intention—what I will call an 'effective intention,' since it is an intention to produce a certain effect on a certain addressee.[2]

An effective intention is normally still one that represents a relatively abstract goal, which could be pursued in various ways. I could try to trick someone into believing something without realizing that this is what I was doing, for example, or I could simply present them with evidence that would guide them to the intended belief without any need for them recognize my intention or trust me. One central claim of this book, however, is that we often enough pursue a different strategy: we try to get our addressee to recognize which effect we intend to have on them, in the hope that this will cause them to enter into the mental state that we intend to produce. I will use the term "revelatory intention" to refer to an intention to reveal an effective intention—the second component of a communicative intention.[3]

Chapter 4 will be dedicated to explaining why we often pursue our effective intentions by forming revelatory intentions, but here is a preview of what I will say there. First, notice that, at least in many cases, revealing an effective intention to my addressee will be an effective strategy for pursuing that intention only given that my addressee trusts me. For example, getting Oona to recognize what I intended her to think would have been a good way to get her to think it only if she trusted me. If she had thought I was lying to her or joking around, she might have made a face and said "noooooo," as she sometimes does. This reliance on trust is a drawback of intention recognition, since it means that we can effectively use the strategy only when our addressee is at least somewhat cooperative and trusting. But there is a worthwhile trade-off between this limitation of intention recognition and a very

[2] Sperber and Wilson (1995) use the term 'informative intention' to refer to what I am calling an 'effective intention.' I have come up with my own terminology because not all effective intentions are intentions to inform. Some effective intentions are intentions to produce non-doxastic states of mind, such as intentions, rather than beliefs, for example.

[3] Sperber and Wilson (1995) inaugurated a tradition of using the term 'communicative intention' to refer to what I am calling a 'revelatory intention' (see also Scott-Phillips 2014). I prefer not to use 'communicative intention' in their way because I have followed another tradition in using that term to refer to the entire complex plan that has an effective intention and a revelatory intention as parts (this tradition includes Strawson 1969; Recanati 1986; and Neale 1992). I hope it is clear why I take it to be important to have distinct terminology for the complex plan and each of its parts.

significant strength. By revealing an effective intention to a cooperative addressee, we recruit them into an active process of understanding and coordination. If my addressee trusts me as a source of information and guidance, and believes that I am speaking with an intention of changing their mind in a certain way, they will be motivated to invest their own cognitive resources into the project of figuring out what I intended. And once they have figured that out, they will have a powerful new reason to enter the intended state of mind. So although it is often useful only in at least somewhat cooperative contexts, intention recognition is made powerful by this very reliance on cooperativity and trust.

Communicative intentions are themselves hierarchical plans, whose parts stand in a plan–subplan relationship. But, zooming out, we can also see them as elements in larger hierarchical plans that connect our extra-communicative goals to motor instructions. My effective intention to get Oona to believe that I was drinking iced tea was itself a subplan of my intention to answer her question, which was in turn a subplan of my intention to have a pleasant and cooperative conversation with her during breakfast. (We could ascend further in the hierarchy to find even more abstract parenting goals.) Message design is the process of reasoning from these broader goals to an effective intention, which pairs a message with an addressee. And, once I had a revelatory intention, this raised further questions of signal design: How could I get Oona to recognize what I was intended her to think? Answering this question involved further practical reasoning in order to settle further questions about which words to use, in which order, in which tone of voice, with which prosodic focus, with which accompanying facial expressions and co-speech gestures, and so on. Only once I had made all of these choices, each one constrained by and in service of my prior intentions, preferences, and beliefs, could my decisions be turned into specific bodily movements.

A communicative intention is therefore one of perhaps many intentions that serve as elements in a complex plan that links an agent's various goals to the bodily movements that constitute the performance of communicative acts. But they are also special, because effective intentions constitute the step in the planning process at which we pair a message with an addressee. This means that the formation of a communicative intention is the point at which the terms of successful communication are set. Everything upstream in the process is message design—the process of deciding what to communicate to whom. To count as message design, this

process has to culminate in an action-guiding representation that pairs a message with an agent. Practical reasoning that happens downstream from a communicative intention is signal design—the process of customizing an utterance or other bodily movement that will guide one's chosen addressee to recognize one's effective intention. Since signal design is the process of reasoning about how to convey a certain message to a certain addressee, it must take an action-guiding representation pairing the addressee with the message as its starting point. If I am right, then communicative intentions are the bottlenecks through which practical reasoning must flow in order to give rise to the very special kind of communication design displayed by humans.

4 From analysis to explanation

I have just given an explanation for why humans form communicative intentions, why these communicative intentions set the terms of successful communication, and, therefore, why we communicate via intention recognition. This line of thought is quite different than the one that originally led Grice to posit communicative intentions. He needed communicative intentions in order to formulate what he took to be the intuitively correct analysis of one of our ordinary concepts of meaning. Here is one formulation of his analysis of "utterer's occasion meaning" (Grice, 1969, 151):

> "*U* meant something by uttering *x*" is true iff, for some audience *A*, *U* uttered *x* intending
>
> (i) *A* to produce a particular response *r*;
>
> (ii) *A* to think (recognize) that *U* intends (i);
>
> (iii) *A* to fulfill (i) on the basis of his fulfillment of (ii).

The first and second clauses of Grice's analysis correspond to effective and revelatory intentions, respectively. Grice also added a third clause in order to capture the means–end relationship between the second clause and the first, and to rule out counterexamples in which a speaker reveals their effective intention for rea-

sons other than as a strategy for accomplishing it.[4] But it is a mistake to capture this means–end relationship by positing yet another intention. It is sufficient to say that a communicative intention is a plan consisting of an effective intention with a revelatory intention with its subplan. Although there are many kinds of complex plans about how to communicate, perhaps even including some others that involve both an effective intention and a revelatory intention, it is my view that the true power of intention recognition is unlocked only when we reveal our effective intentions as a way of fulfilling them. (I will say more about why in the next chapter.)

Grice arrived at the above analysis by considering a series of explications of statements of the form, 'by uttering *x*, *S* meant *p*,' and testing these explications against his intuitive judgments about a variety of increasingly tricky cases. In his later work, and in the work of many later authors who have tried either to amend or oppose Grice's account, this methodology became the basis for an arms race of ever more intricate analyses fueled by ever-more-complex and often recherché imagined scenarios. This methodology has never led to a consensus about the proper formulation of communicative intentions, and intuitions about imagined cases have continued to flow.[5]

I am pessimistic about the prospects of this way of justifying a theory of the nature of communication. One reason for pessimism is that this methodology has tended to be at the center of degenerating research programs, in which much en-

[4]Grice motivated his third clause using an example in which Herod gets Salome to believe that John the Baptist is dead by presenting her with his head on a platter. In this case, Herod intends for Salome to believe that Herod is dead (satisfying clause (i)) and Herod intends for her to believe that Herod has this intention (satisfying clause (ii)). "Yet [Grice] certainly [did] not think that we should want to say that we have here [a case] of meaning$_{NN}$" (i.e., the sort of meaning that Grice was trying to analyze) (Grice, 1957, 382). We could diagnose the problem by saying that Salome doesn't need to know about Herod's intentions in order to find out that John the Baptist is dead, since she can see it with her own eyes. Although Herod has a revelatory intention, it is not a subplan of his effective intention. Rather, Herod wants Salome to recognize his intention as part of his plan to taking credit for John the Baptist's murder. For Grice, this explanation serves merely to confirm his intuitive judgment. But there is a deeper point: it is only when one adopts a revelatory intention as part of a strategy for pursuing the effective intention that we pursue the communicative strategy of intention recognition.

[5]A partial bibliography of the arms race: (Bach and Harnish, 1979; Davis, 1992; Grice, 1957, 1969; Harris, 2014; Neale, 1992; Schiffer, 1972; Strawson, 1964) [ETC.]

ergy has gone into the search for intuitively correct analyses, with no convergence on the right account.[6] The literature on intention-based definitions of speaker meaning that I've just cited is a good example of this sort of degenerating research program. Attempts by epistemologists to analyze the concept of knowledge provide an even more famous case.[7] These literatures are case studies in the futility of seeking intuitive-counterexample-proof analyses of everyday concepts.

But the dismal track record of this methodology is not its real problem. Even if it were possible to give an intuitively bulletproof analysis of the concept of speaker meaning, what would be gained by this achievement? The methodology makes sense only if our goal is to precisely recapitulate the contents of our folk theory of meaning and communication. What else could our intuitive judgments about correct usage be telling us about, after all? But our aim in giving a theory of communication should not be to articulate the contents of a folk theory. Rather, we should be trying to construct an empirical theory that explains how humans in fact communicate. And we can't assume in advance that the most explanatory empirical theory will line up with our folk theories, or that our folksy concepts about communication are up to the task of playing explanatory roles in such a theory, unaltered and unaugmented.

This point is compatible with the possibility that our folk theories are on the right track in at least some ways, and this might help to explain how Grice managed to catch a glimpse of the truth using his flawed methodology. But although Grice's methodology may have afforded him such a glimpse, it can't be used to justify the claim that his theory of meaning is correct, because all his methodology tells us is that his theory sounds right by the lights of our folk-theoretic judgments. This would not be acceptable as a justification for a theory in any other empirical domain, and it should likewise be banished from the context of justification here.

Moreover, what we should want is not just an account of what it is to mean something—i.e., of what sets the terms of successful communication. We should also want an explanation of why the answer that we arrive at is the right one. Why, for example, is it communicative intentions, rather than, say, the speaker's beliefs or expectations or desires, or for that matter, some facts about linguistic conventions

[6] Fodor (1998) makes a forceful version of the general point.

[7] For a summary of this literature, see Ichikawa and Steup (2018), including Section 7 on contemporary pessimism about the project.

or social commitments, that do the job? This question is an urgent one, because each of these views has been proposed as a superior alternative to intention recognition. Grice's methodology can offer us little more on this front than the claim that his account matches our judgments as a description of whatever cases we've come up with so far. But the theory that I have offered does more than that. My claim is that we need to posit communicative intentions in order to explain how we use our capacities for practical reasoning and mindreading to customize messages and signals for our addressees. And when we use this strategy, communicative intentions set the terms of successful communication because their adoption is the step in the communication-design process at which we have settled on a message to be communicated and an addressee to which it is to be communicated.

5 The Psychological Reality of Communication Design

I have been arguing that ordinary human communication involves a lot of practical reasoning and a lot of mindreading—much more of both than we normally notice ourselves doing. In making this argument, I have steered headlong into one of the most influential objections to Grice's ideas. The criticism is that intentionalists like Grice over-intellectualize human communication, positing far more, and more sophisticated thoughts and psychological processes than are actually there. This line of thought has been repeated many times, and has often been used to motivate alternative theories of human communication that place fewer cognitive demands on communicators.[8]

Grice was aware of this worry from the beginning. The paper in which he first posited communicative intentions closed with a response to the charge that he was "peopling our talking life with armies of complicated psychological occurrences" (1957, 386). In that discussion, Grice's faithfulness to folk theory may have led him to concede too much, saying that "explicitly formulated linguistic (or quasi-linguistic) intentions are no doubt comparatively rare" (Grice, 1957, 387). Now, if Grice meant only that our communicative acts aren't guided by *conscious* plans,

[8]See, for example, Devitt (2021, ch.4); Evans and McDowell (1976, xx); Geurts (2019, 2), Millikan (1984, ch.3).

then I agree: much of the reasoning and many of the states of mind behind our communicative acts are non-conscious. This puts these thoughts and psychological processes in good company: much of our behavior is driven my non-conscious psychological states and processes. Nonetheless, many of these states and processes are perfectly "explicit" in two senses that may have little to do with consciousness: they have or traffic in specific representational contents, and they guide our actions in specific, rational ways. We are warranted in positing these mental states and processes not because we can reliably introspect them, but because they their existence is entailed by the best explanation of some aspect of human behavior. I have argued that this criterion is met by communicative intentions and the mindreading and practical reasoning by means of which we form and reason from them. We are compelled to posit communicative intentions as part of our best explanation of humans' remarkable capacity for communication design.

Suppose that you find this argument convincing. Still, you might object that the sort of practical reasoning that I have posited is only for special occasions. Perhaps humans normally rely on a less cognitively demanding way of communicating, but become more strategic and invest more resources only when they face a particularly difficult communicative challenge or want to pull off a particularly virtuosic communicative trick. Perhaps intention recognition isn't itself the normal mechanism of human communication, but is an add-on that we use to soup up the normal mechanism from time to time?

This sort of view has been defended or suggested by a number of theorists. For example, Millikan (1984, ch.3; 1998) has argued that we should understand most literal communicative acts as possessing their properties in virtue of the proper functions of the signals used to perform them. An act performed with a certain kind of signal (such as a sentence) has the proper function that it does by virtue of historical facts about which effects signals of that kind have produced in the past, such that communicators have wound up using them again.[9] For example, Millikan argues that the imperative clause type has the function of eliciting action, and it has this function because at least some past utterances of imperative clauses have elicited action, and speakers have gone on to utter more imperative clauses

[9]Millikan argues that the proper function of a syntactically complex signal-type may be determined by the functions of the signal-type's components. This is her (admittedly somewhat barebones) nod to the compositionality of natural language.

in part because of that (1984, ch.3). One supposed virtue of this account is that the psychological mechanisms involved needn't be all that sophisticated (Millikan, 1984, 58, 62). Imagine that someone has barged into your office and you want them to leave. Luckily, your own and others' past experiences of successful communication have furnished your mind with associative connections between acts of leaving and the verb 'leave' and between desires to elicit action in others and the imperative mood. These associative links—presumably together with grammatical mechanisms—cause you to utter 'leave!', which causes the interloper to leave. In simple cases of literal communication like this, Millikan thinks, we needn't posit any communicative intentions or complex reasoning; some less sophisticated psychological mechanisms connecting our inner states to certain kinds of signals, suitably installed in us by histories of communicative success, are enough to endow those signal-types with proper functions that set the terms of successfully communicating with them.

This sort of explanation doesn't make sense of non-literal communicative acts, in which "the proper function (or functions) of an expression in a public language [contrasts] with the function that a speaker intends for it on a given occasion" (Millikan, 1995, 187). These are among the cases in which Millikan *does* think that the conditions of communicative success are set by the speaker's intentions. For Millikan, then, intentions become indispensable to the theory of communication only when it comes to certain creative or otherwise function-divergent uses of signals. Intention recognition isn't the usual way that we communicate, she argues, though it is an advanced communicative skill that we sometimes fruitfully put to use.

Millikan's picture is coherent, and it is not hard to see its attraction. But I think that it dramatically underestimates the role of intention recognition in even literal communication. My reason is that communication design is a ubiquitous feature of human communication, including ordinary language use. And as I have already argued, communicative intentions inevitably arise in the course of the kind of flexible, informationally omnivorous practical reasoning that humans use for communication design.

One way to see that this is the case is to notice that certain basic features of natural languages could not be competently used by anyone other than a sophisticated communication designer. The rich system of noun phrases that we find in English, and its counterparts in other languages, provide a good example. Imagine the fol-

lowing scenario: I am reading the *Philosophical Investigations* when someone asks me who the book's author is. It would be natural to answer with a noun phrase that picks out Wittgenstein, but there are many—indeed, infinite—options that could do that job, depending on the circumstances. Here are just a few:

(1) a famous philosopher named Ludwig Wittgenstein

(2) the philosopher I told you about last week

(3) the brother of the woman in that Klimt painting you liked

(4) Wittgenstein

(5) Ludwig

(6) him

Which of these should I choose? The answer depends on what I take to be my addressee's state of mind. If I am talking to my daughter, I might utter (1): Oona isn't familiar with Wittgenstein but knows what a philosopher is, and indefinites typically function to introduce novel referents.[10] If my addressee is already familiar with Wittgenstein but might not know his name and aren't actively thinking about him, I should use a definite description that leverages whatever I think they *do* know about him to get them to think about him now. For example, (2) could be good for a student who learned about Wittgenstein a week earlier but may not recall his name. And I might utter (3) if my addressee has little interest in philosophy but has recently come back from the Neue Pinakothek Museum in Munich, raving about Gustav Klimt's portrait of Wittgenstein's sister, Margaret Stonborough-Wittgenstein. If my addressee is a colleague who knows all about Wittgenstein but just didn't recognize my book, I would probably just use his last name, (4). On the other hand, I should use his first name, (5), only in the unlikely scenario that my

[10]The mechanism by which indefinites introduce novel referents is controversial. Some have argued that this is just part of their semantic function (e.g. Heim 1983). But see Lewis (2021) for an alternative, pragmatic theory of why indefinites often introduce novel referents, which explains some notable exceptions, and which is built on the idea that we communicate by intention recognition.

addressee and I take each other to be on a first-name basis with Wittgenstein. Finally, I should use a pronoun, such as (6), if Wittgenstein is currently prominent in my addressee's thoughts, either because we have just been talking about him (and haven't discussed other males in the mean time), because Wittgenstein is currently perceptually salient to both of us (say, because we are both looking at a portrait of him on the wall), or because I know that they have been actively thinking about Wittgenstein for some other reason.

So, being a competent user of a noun-phrase system like the one in English means being a communication designer: it means making intelligent choices about which expressions to use on the basis of integrated information about the intended referent, the addressee, and the latter's perceptual and cognitive perspective on the former.

This point extends beyond noun phrases and into many other features of natural language. Many natural-language expressions can be used to communicate a range of different contents. Some of the much-studied examples include indexicals (e.g. 'here,' 'we'), demonstratives ('this,' 'that,' 'them'), quantifiers ('every beer,' 'most fish'), modals ('must,' 'can'), gradable adjectives ('tall,' 'open'), taste predicates ('fun,' 'delicious'). Using any of these expressions to communicate typically requires coordinating with your addressee on using them in at least approximately the same way.[11] If I tell my addressee that "we" should organize a conference, my addressee will have to figure out which group of people I am referring to. If I intend to convey some information about a certain group of people but I don't think that my addressee will be able to work out which group if I refer to it as 'we,' then I should find a better way to refer to the group. Likewise, if I want to tell someone that a certain basketball player is almost seven feet tall, but I know

[11] Two caveats: First, there sometimes seems to be at least some slack in how an addressee may understand us without communicative failure (Buchanan, 2010; von Fintel and Gillies, 2008; King, 2022). I take it that even in these cases, communication design is needed to ensure that the addressee will zero in on an interpretation in the correct range. Second, we sometimes use a context-sensitive expression not to communicate a specific first-order content, but to indirectly communicate information about how the expression itself ought to be understood going forward in the conversation (Barker, 2002; Ludlow, 2014; MacFarlane, 2016; Plunkett and Sundell, 2013; Stalnaker, 1978). Although these uses do exist, I take it for granted that we also often presuppose a certain way of using a context-sensitive expression and then use it to convey information, and those are the kinds of uses I am talking about in the main text.

that they don't know enough about basketball to know that this is what it takes to be a tall basketball player, I shouldn't merely describe him as "tall." In general, reliably communicating with context-sensitive vocabulary requires tracking many features of your addressee's state of mind and incorporating this information into the communication-design process.

The same point also extends to ambiguity. In order to reliably communicate with ambiguous expressions, we have to recognize and avoid situations in which our addressee will pick up on a different meaning than the one we intend, and doing this requires intelligent communication design informed by mindreading. This goes for structural ambiguities ('flying planes can be dangerous'), homophony ('bank,' 'pen'), and also—crucially—polysemy. A polysemous expression is one with two or more related meanings. For example, 'newspaper' can be used to refer to the thing you buy at a news stand or to the institution that produces those things. This polysemy gives rise to two readings of (7) with dramatically different contents, for example.

(7) Jeff Bezos bought two newspapers last week.

Polysemy is ubiquitous in natural language. Many, if not all open-class lexical items are polysemous, including most or all nouns, verbs, adjectives, and adverbs. As with other ambiguous and context-sensitive expressions, being a competent, reliable user of a polysemous expression means refraining from using that expression when your addressee is likely to understand it with a meaning other than what you intend, thereby leading to miscommunication. For example, if you want to engage your addressee in a conversation about the dangers of oligarchy in America, you might utter (7); but you should avoid doing this if think that your addressee will likely take you to have meant by (7) that Bezos bought two printed copies of newspapers last week. (Maybe you know that they're preoccupied with the decline of news literacy among elites.) But if polysemy is all over the place in natural language, and reliably communicating with polysemous expressions requires mindreading-fueled communication design, then communication design is not something that gets added on to communication in special, creative cases. Rather, you have to routinely engage in communication design in order to be a competent language user. And, if my argument in §3 is correct, then it follows that competent ordinary language use involves intention recognition.

Here is a possible objection: Miscommunication is actually pretty common. One reason for this is that we often fail at the kind of communication design that I have argued is part of our basic pragmatic competence. We often *do* utter noun phrases, context-sensitive expressions, and ambiguities that our addressees misunderstand or don't know how to interpret. Sometimes this is because we have false beliefs about our addressees that lead the communication-design process astray. (Imagine a New Yorker giving directions to a tourist whom they mistakenly believe to be a local, and who gives incomprehensible directions as a result: "It's right next to Joe's Pizza, but not the good Joe's pizza, the other one.") But in other cases, even if we have access to all of the information about our addressees that we need, we fail to integrate that information into the communication-design process. For example, teachers sometimes slip into jargon that we should know our students will find confusing. On the flip side, we sometimes annoy our colleagues by over-explaining ideas that, we should realize, they already understand perfectly well. And consider the failures of communication design that must have led the Guardian to run a 1982 story about left-wing ambivalence concerning the Falkland War with the headline, "British Left Waffles on Falklands."

Doesn't this show that humans aren't so great at communication design after all? More to the point, doesn't it show that we can and often do use language without all the communication design that I have been talking about?

My full answer to this question will come in Chapter 5, where I will describe the considerable body of empirical evidence on the role of mindreading in the design and interpretation of communicative acts. Here is a preview of some of what I will argue there:

Part of my answer is that I have claimed only that communication design is required for *competent* language use, and so pointing out lots of examples of incompetent language use is no objection. If I am right, then natural language is a tool that is sometimes difficult to use properly, because the practical reasoning and mindreading that go into the design and interpretation of communicative acts are at least sometimes effortful and resource-intensive, and we aren't always willing or able to invest the effort and cognitive resources that would be ideal. When this happens, miscommunication can result. But, I will argue in §6 that the very features that make competent language use cognitively demanding also confer enormous communicative advantages. And so even if language is an unwieldy tool that we

sometimes struggle to use properly, its power makes this struggle worthwhile.

A second part of my answer is that there is a difference between doing communication design badly and not doing it at all, and I think it is best to describe most failures in the former terms rather than the latter. When we fail at communication design, we are normally cutting corners, falling back on resource-saving heuristics. These heuristics are more likely to lead to miscommunication than putting in the work would, but they work out often enough that they're better than nothing. And they are best understood as less-good methods for communication design than as failures to engage in communication design.

Again, these are empirical claims about human psychology for which I will defend in Chapter Five. For now, I will move on to the question of why we would bother to make the kinds of costly cognitive investments that, I have just claimed, are involved in communication design.

6 Why Communication Design?

I have argued that humans are prolific communication designers, that communication design is a necessary feature of competent language use, that to do it is to engage in domain-general practical reasoning, and that when we do this kind of reasoning, we form communicative intentions along the way. I have also argued that communication design is not easy. It expends scarce cognitive resources, and we sometimes miscommunicate because the relevant resources run short. So what makes communication design worth it? Why do we do it? And why would natural language be organized so as to force us to do it?

My answer is that communication design makes us much better communicators, both by increasing the flexibility and efficiency with which we communicate, and also by allowing us to wield natural languages with much greater expressive power than we otherwise could.

Some of these benefits of communication design are straightforward applications of the more general benefits of our capacity for practical reasoning. Just as hierarchical practical reasoning in general allows us to act in the service of abstract, social goals in ways that are not mere reactions to our immediate situations, it likewise allows us to communicate in the service of these goals. If I send you a text message inviting you to my dinner party, this action can be explained only by ap-

peal to a complex plan that connects my abstract goal of hosting a dinner party to the specific movements of my fingers in this moment. The reasoning that led to these finger movements needn't have been initiated by any feature of the environment in which I typed, and the specifics of my action would have been informed by my knowledge of your beliefs and plans and by my anticipation of how my message would alter them. Our ability to create plans that allow us to act on goals of these kinds is obviously extremely valuable in general, but they also make us much better communicators, in part because they allow us to say more kinds of things in the service of more kinds of goals, and in part because they allow us to say those things in ways that are more likely to be understood and have the effects that we intend.

By accounting for what I know about my addressee's background information, I can steer between the twin hazards of incomprehensibility and redundancy, communicating more efficiently with those who have lots of relevant background information and vocabulary, but sacrificing brevity for the sake of filling in lacking background information for those who otherwise wouldn't understand. (For example, I can avoid telling Oona that I'd brewed my tisane hot and then poured it over ice, but I can also avoid telling a tea expert that I am drinking tea.) Conditioning what I say on my addressees' preferences and goals gives me flexible ways to be both cooperative (acting in ways that serve their preferences and goals) and strategic (selectively frustrating them).

Our capacity to recognize and adapt to each others' goals also allows us to conduct our conversations as cooperative joint activities, which comes with enormous benefits. This will be one of the main topics of Chapter Four, and so I won't belabor the point here. But the basic idea is that when we are able to organize conversations in a coherent way around shared plans, this allows us to work efficiently toward our shared goals by investing cognitive resources in mutually supportive and beneficial ways. Interlocutors could not hold up their end of this arrangement if they could not condition their decisions about what to say and how to say it on what they knew about their addressees' beliefs and plans.

So far I have discussed benefits of communication design that apply to non-linguistic and linguistic communication alike. But there are also some specifically linguistic benefits. Many basic design features of natural languages allow us to communicate more information more efficiently, but can only be used in this way

by adept communication designers. I have already discussed some of these features in Section 5: perspectival noun-phrase systems, context-sensitive expressions, and polysemous expressions. Until now, I have presented these features of natural language as obstacles that we need communication design to overcome. But they also grant us significant benefits, and the existence of these benefits helps to explain why we would have these features in natural language at all, given the significant cognitive demands on their use.

Consider the ways in which the noun-phrase systems give us many different ways of leveraging our addressee's perspective on a referent in order to get them to think a new thought about it. Pronouns are efficient tools for communicating about a referent that is already salient (or that can somehow be made salient in the moment). But if pronouns were our only noun phrases, we would be unable to refer to non-salient entities, and this would severely limit our topics of conversation. Names, on the other hand, can be efficient tools for communicating about a referent who is buried deeper in the addressee's long-term memory, but only if the speaker and addressee know the referent by the same name. Definite descriptions give an option for when we can't rely on salience or mutually known proper name, allowing us to leverage the addressee's specific prior beliefs about the referent in order to impart some new representations of it. But just as descriptions can be useful when pronouns and names would not, the reverse is also true: we can sometimes use names and pronouns to successfully refer to a referent about which neither we nor our addressees possess enough information to uniquely pick them out with a definite description (Kripke, 1980). Finally, we can use indefinites (e.g. 'a guy I know') in order to communicate about a referent with whom the addressee is, from what we know, entirely unfamiliar, thereby indicating that they should create a new representation—"start a new file," so to speak (Heim, 1982). Of course, there are also many options within each of these categories, which allow us to further refine our attempts to activate and add to our addressees' representations of our intended referents. In sum, this system gives us a versatile set of options for getting our addressees to think about a referent, in spite of the very different kinds of perceptual or cognitive perspectives that they might have on it.[12] To take advantage of these

[12] There is of course much more to say about the formal-semantic, pragmatic, and psycholinguistic properties of noun phrases that allow them to achieve their diverse functions. I should acknowledge the influence of Neale (2008) on the thoughts in this paragraph, as well as Heim (1982), Prince

options, we need to track and infer our addressees' states of mind, and this expends scarce cognitive resources. Those who are short on these resources, and who must fall back on a heuristic, are more likely to misuse the tools and miscommunicate as a result. But the tools so greatly expand and diversify our communicative abilities that they are worth the trouble when we can manage them.

As I argued in Section 5, much the same can be said about our ability to use other context-sensitive vocabulary. For example, it is sometimes possible to convey relatively specific information by using a gradable adjective in the positive form—for example, to convey information about someone's height by saying 'he is tall.' But this is possible only to the extent that the speaker and addressee take a similar standard for tallness to be operative—something that may vary between occasions. If, in uttering 'he is tall,' I am trying to say that a certain man is over $6'5''$, and you take me to be saying that he is over $5'11''$, then we miscommunicate. In order for gradable adjectives to be useful, we need to be able to anticipate uncoordinated background assumptions of this kind. But in exchange for this cost, we get a versatile lexical item that expands our expressive resources, allowing us to predicate many different heights of individuals.

This same pattern of costs and benefits can be found across a range of other context-sensitive expressions. If I say, 'every student is here,' I could be talking about every student in my class, in my department, at my college, in America, or really any other contextually restricted group of students.[13] If I say, 'the students must be here,' I could be talking about what the students are obligated to do by the law, by our social norms, by the categorical imperative, or by any other body of prescriptions—or I could be saying that it follows that they are here from what I know, from what we have been collectively assuming, or from some other body of information.[14]. If I say, 'The students were here,' I could saying that they were here five minutes ago, earlier today, yesterday, or last week.[15] Quantifier phrases (e.g.

(1981), Roberts (2002; 2003; 2005), and Lewis (2021). I will return to this topic in more detail in a later chapter.

[13]Some influential works on quantifier domain restriction include (Barwise and Perry, 1983; von Fintel, 1994; Neale, 1990; Stanley and Szabó, 2000; Westerståhl, 1984).

[14]The classic work on the context sensitivity of modal auxiliaries is due to Kratzer (1977; 1981); for a recent overview, see Matthewson (2016).

[15]The classic text on the context-sensitivity of tense markers is Partee (1973); for a recent overview, see Grønn and von Stechow (2016).

'every student'), modal auxiliaries (e.g. 'must'), tense markers (e.g. the past-tense marker attached to 'were') are among the many natural-language expressions that can be used to communicate an open-ended range of messages, thereby greatly expanding the expressive power of natural language. But our expressions can be versatile in this way only because we are decently good at anticipating when our addressees will focus on the same messages that we intend to convey. And so a language with such versatile expressions presupposes speakers who are adept at communication design.

Finally, consider the fact that natural-language speakers are all lexical specialists. There are millions of words in the Oxford English Dictionary, but each English speaker has at least an order of magnitude fewer words than that in their own lexicon, with lots of variation from one speaker to the next. There are big advantages to this arrangement: Each of us gets to have specialized tools at our disposal for efficiently communicating about things that we have particular needs to talk about. I get specialized terminology that allows me to talk precisely and efficiently about the nature of communication, for example, but equally, a postal worker gets to have specialized terminology that allows them to more efficiently solve the communicative problems that they face. It would be difficult to exaggerate how valuable this is, particularly when we consider the role that specialists play in introducing new words into our languages that may turn out to be useful enough to others that they spread to other populations. But this way of distributing lexical resources can work only if language users are at least fairly good at making intelligent decisions about which addressees are likely to have which vocabulary. In order to be understood, we have to avoid using too much philosophical terminology at the post office or too much postal terminology in a philosophy seminar, to take just one silly example. Our system of lexical specialization could only exist among speakers who design their utterances for their addressees.

7 Looking Forward

In this chapter, I have argued that human communication is made powerful by our capacity for practical reasoning, and that intentions about what to convey to whom ("effective intentions") are important steps in the process of designing our utterances for particular addressees. In the next chapter, I will turn my intention

to our capacity to treat conversations as cooperative joint activities, which is another communication-supercharging human endowment. Among other things, I will argue that cooperativity is what explains why we bother to form intentions to reveal our effective intentions to others ("revelatory intentions"). Doing so, I will argue, makes it much easier to efficiently coordinate our communicative efforts with others. After that, I will spend a chapter making good on my claim that we really do all of the complex mindreading that I have posited in this chapter (and more of which I will posit in the next chapter). Then I will turn my attention to language.

References

Bach, K. and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, Mass.

Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25:1–36.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press.

Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press, Cambridge, Mass.

Bratman, M. (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford University Press, Oxford.

Buchanan, R. (2010). A puzzle about meaning and communication. *Noûs*, 44(2):340–371.

Davis, W. (1992). Speaker meaning. *Linguistics and Philosophy*, 15(3):223–253.

Devitt, M. (2021). *Overlooking Conventions: The Trouble with Linguistic Pragmatism*. Springer, Cham.

Evans, G. and McDowell, J. (1976). Introduction. In Evans, G. and McDowell, J., editors, *Truth and Meaning: Essays in Semantics*, pages vii–xxiii. Oxford University Press, Oxford.

von Fintel, K. (1994). *Restrictions on Quantifier Domains*. PhD thesis, University of Massachusetts, Amherst, Amherst, MA.

von Fintel, K. and Gillies, A. S. (2008). CIA leaks. *Philosophical Review*, 117(1):77–98.

Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press, Oxford.

Geurts, B. (2019). Communication as commitment sharing. *Thoeretical Linguistics*, 45(1–2):1–30.

Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3):377–388.

Grice, H. P. (1969). Utterer's meaning and intention. *The Philosophical Review*, 78(2):147–177.

Grønn, A. and von Stechow, A. (2016). Tense. In Aloni, M. and Dekker, P., editors, *The Cambridge Companion to Formal Semantics*, chapter 11, pages 313–341. Cambridge University Press, Cambridge.

Harris, D. W. (2014). *Speech Act Theoretic Semantics*. PhD Dissertation, City University of New York Graduate Center.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD Dissertation, University of Massachusetts at Amherst.

Heim, I. (1983). File change semantics and the familiarity theory of definiteness. In Bäuerle, R., Schwarze, C., and von Stechow, A., editors, *Meaning, Use and Interpretation of Language*, pages 164–189. de Gruyter, Berlin.

Ichikawa, J. and Steup, M. (2018). The analysis of knowledge. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>, summer 2018 edition.

King, J. C. (2022). *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. Oxford University Press, Oxford and New York.

Kratzer, A. (1977). What *Must* and *Can* must and can mean. *Linguistics and Philosophy*, 1(3):337–355.

Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H. J. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. de Gruyter, Berlin.

Kripke, S. (1980). *Naming and Necessity*. Harvard University Press, 1980.

Lewis, K. (2021). Anaphora and negation. *Philosophical Studies*, 178(5):1403–1440.

Ludlow, P. (2014). *Living Words*. Oxford University Press.

MacFarlane, J. (2016). Vagueness as indecision. *Proceedings of the Aristotelian Society, Supplementary Volume*, 90(1):255–283.

Matthewson, L. (2016). Modality. In Aloni, M. and Dekker, P., editors, *The Cambridge Companion to Formal Semantics*, chapter 18, pages 525–559. Cambridge University Press, Cambridge.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. MIT Press, Cambridge, MA.

Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical Perspectives*, 9:185–200.

Millikan, R. G. (1998). Proper function and convention in speech acts. In Hahn, L. E., editor, *The Philosophy of Peter F. Strawson*, The Library of Living Philosophers, pages 25–43. Open Court, LaSalle, Illinois.

Neale, S. (1990). *Descriptions*. MIT Press.

Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15:509–559.

Neale, S. (2008). Term limits revisited. *Philosophical Perspectives*, 22(1):375–442.

Partee, B. (1973). Some structural analogies between tense and pronouns in english. *Journal of Philosophy*, 70(18):601–609.

Plunkett, D. and Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13(23).

Prince, E. (1981). Toward a taxonomy of given-new information. In Cole, P., editor, *Syntax and Semantics, Vol. 14: Radical Pragmatics*, pages 223–255. Academic Press, New York.

Recanati, F. (1986). On defining communicative intentions. *Mind & Language*, 1(3):213–242.

Roberts, C. (2002). Demonstratives as definites. In van Deemyer, K. and Kibble, R., editors, *Information Sharing*, pages 1–48. CSLI.

Roberts, C. (2003). Uniqueness in definite noun phrases. *Linguistics and Philosophy*, 26:287–350.

Roberts, C. (2005). Pronouns as definites. In Reimer, M. and Bezuidenhout, A., editors, *Descriptions and Beyond*, pages 503–543. Oxford University Press, Oxford.

Schiffer, S. (1972). *Meaning*. Oxford University Press, Oxford.

Scott-Phillips, T. (2014). *Speaking Our Minds: Why Human Communication is Different, and How Language Evolved to Make it Special*. Palgrave Macmillan.

Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell, Oxford.

Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and Semantics 9*, pages 315–332. Academic Press, New York.

Stanley, J. and Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15(2–3):219–261.

Strawson, P. F. (1964). Intention and convention in speech acts. *The Philosophical Review*, 73(4):439–460.

Strawson, P. F. (1969). *Meaning and Truth*. Inaugural Lecture at the University of Oxford. Oxford University Press, Oxford.

Westerståhl, D. (1984). Determiners and context sets. In van Benthem, J. and ter Meulen, A., editors, *Generalized Quantifiers in Natural Language*, pages 45–71. Foris, Dordrecht.