

Gricean Communication, Natural Language, and Human Evolution

Daniel W. Harris*

Draft of July 25, 2023

Final version to appear in *Evolutionary Pragmatics*, edited by Bart Geurts and Richard Moore for Oxford University Press

Abstract

I argue that our capacity to use natural language depends on our capacity for Gricean communication—i.e., our capacity to communicate by revealing and recognizing intentions to change addressees' states of mind. I defend this claim by showing that communicative intentions play central roles in our capacity to customize what we say and how we say it for our addressees, and in our capacity to organize our conversations around shared plans. Our ability to take advantage of ubiquitous and communicatively valuable features of natural language depend on these capacities, and so our capacity for Gricean communication is part of what it takes to be a competent language user. But, contrary what both Griceans and their critics have maintained, it does not follow from this either that Gricean communication had to evolve before language or that all linguistic communication takes the ideal Gricean form. Not all language use is competent, and not everything that deserves to be called a language need have features that only Gricean communicators can competently use.

1 The Gricean model and the evolutionary challenge

According to a view originating in the work of Paul Grice (1957; 1969), much ordinary human communication, including linguistic communication, takes the form of revealing and recognizing intentions. A communicative act is an act done with a

*Thanks to Ray Buchanan and Richard Moore for comments on earlier drafts.

communicative intention, which consists of an intention to produce a state of mind in an addressee, together with a second intention to reveal the first intention to them. One kind of communicative success consists in the addressee correctly recognizing what the communicator intended. A further kind of communicative success consists in the addressee entering into the intended state of mind.

If this Gricean model is correct, then natural language functions as an add-on to our capacity for Gricean communication—as a system for offering our addressees particularly rich evidence of our intentions. It seems to follow that language use depends on the cognitive capacities needed to intelligently form and act on complex metarepresentations—for example, intentions about beliefs about intentions about states of mind. Doing this well requires sophisticated and integrated capacities for mindreading and planning that allow us to infer and predict others' states of mind and use the resulting information to design and interpret communicative acts.

The idea that language use depends on these sophisticated cognitive capacities exposes the Gricean model to a variety of challenges, arising from doubts about whether child language users, neuroatypical language users, and even neurotypical adults could really be doing all of this reasoning when they use language to communicate.¹

What I say here will have implications for these debates, but my focus will be on a related collection of related challenges related to human evolution. Although non-human great apes have limited capacities for planning and mindreading, our current evidence suggests that these capacities are not sufficiently powerful or integrated to support Gricean communication. The same presumably goes for our last common ancestors with chimpanzees and bonobos.² Griceans are therefore stuck with a particularly thorny version of the already difficult problem of explaining how our capacity for linguistic communication could have evolved. In addition to explaining how humans came to possess and integrate the articulatory, perceptual, phonological, syntactic, semantic, prosodic, and other cognitive and social capacities that make up our capacity to use natural language itself, Griceans also have

¹On child language users, see Breheny (2006); Moore (2017b). On language use in people with autism-spectrum disorder, see (Tager-Flusberg, 2000).

²Rather than review the evidence for these empirical claims, I will take them as a given in what follows. For recent surveys of the evidence on primate mindreading, see (Krupenye, 2020; Lurz, 2011). On practical reasoning in non-human animals, see Camp and Shupe (2018). For the claim that non-human primates lack the capacity for Gricean communication, see Armstrong (2021); Cheney and Seyfarth (2007); Scott-Phillips (2015); Tomasello (2006). A notable dissenter on this point is Richard Moore, whose views I will discuss in §2.

to explain how our ancestors evolved the cognitive capacities that are needed for Gricean communication.

This problem is particularly difficult if Gricean communication had to evolve *before* natural language. At least some Griceans as well as some of their opponents have taken this claim about temporal priority to follow from the Gricean model. For example, Origg and Sperber argue from a broadly Gricean perspective that “language as we know it developed as an adaptation in a species already involved in inferential communication, and therefore already capable of some serious degree of mind-reading” (Origg and Sperber, 2000, 165). They reach this conclusion by arguing that natural-language utterances only ever serve as partial and defeasible evidence of speakers’ intentions, and a system like that could only have evolved in the presence of the cognitive capacities needed for Gricean communication.³ Meanwhile, Dorit Bar-On bases her criticism of the Gricean model on her claim that, “on the Gricean view, speaker meaning is a necessary step toward structured linguistic meaning” (Bar-On, 2013b, 345).

The claim that the evolution of Gricean communication had to precede the evolution of language makes the evolutionary challenge more difficult in at least two ways. First, it precludes a timeline in which sophisticated practical reasoning, social cognition, and natural language evolved in parallel, instead forcing us to explain how sophisticated capacities for planning and mindreading could have sprung up prior to and independently from our capacity for language, while also leaving just the remaining fraction of the last 6–9M years for the evolution of language itself. This forces us to cram our stories about the evolution of both Gricean communication and language into even shorter, sequential timeframes. Second, this timeline forces on us the view that neither Gricean communication nor any of the underlying cognitive capacities that support it depend on our capacity for natural language. But some have argued that the sort of planning and mindreading involved in Gricean communication themselves involve mental representations with language-like, recursive and hierarchical structures, and so these capacities may depend on the capacities for syntax and semantics that underlie natural language.⁴

³For other arguments that language had to evolve after Gricean communication, see Fitch (2010, 135) and Scott-Phillips (Scott-Phillips, 2014, §2.1). Michael Tomasello defends a similar but weaker claim: “it is only within the context of collaborative activities in which participants share intentions and attention...that arbitrary linguistic conventions could have come into existence evolutionarily” (Tomasello, 2008, 9).

⁴E.g. Bar-On (1995, 2013b).

2 Alternatives to the Gricean model

These evolutionary considerations have moved some theorists to wholly reject the Gricean model, seeking radical alternatives that promise greater continuity with our ancestors (e.g. Geurts 2019). Others have seen Gricean communication as a recent addition to our communicative repertoires, which explains some of our advanced abilities but which isn't needed for basic language use. Ruth Garrett Millikan (1984, ch.3; 1998) argues that in most communication between adult humans, a signal's communicative function is fixed not by the speaker's intentions, but by its proper function, which it has by virtue of a history of differential reproduction of other signals of the same kind. Relatively unsophisticated mechanisms may be responsible for the encoding and decoding of information in signals, and it is these mechanisms that have been trained by past usage to be coordinated across speakers. In a similar vein, Dorit Bar-On (2013a; 2013b; 2021) identifies a category of "expressive communication," wherein a communicator produces a behavior that expresses an inner state, thereby "showing" it (and, if it is an intentional state, its content) to their audience. On Bar-On's view, the expressive model captures what is happening in many cases of non-human communication, in child language use, and in many instances of language use by adult humans (Bar-On, 2013b, 355). Similarly, Josh Armstrong (2021) posits a category of "minded communication," in which both the sender's strategy for producing a signal and the receiver's strategy for interpreting it have the function of achieving coordinated mental representations, but in which the mechanisms serving these functions needn't involve sophisticated planning or mindreading.

Each of these theories is compatible with the possibility that some amount of communication between adult humans takes the form of Gricean communication. For example, Millikan appeals to communicative intentions to determine the contents of indirect speech 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). She likewise says of demonstrative pronouns, like "this," that their proper function is to indicate that the hearer should consult the speaker's referential intentions (1984, 167).⁵ Gricean communication, for Millikan, is a trick that neither non-human animals nor human children can pull off, and that adult humans use only for advanced communicative tasks. Most

⁵As we will see in §4, I think that Millikan is onto something here, but vastly underestimates the ubiquity and importance of context sensitivity in natural language.

of the time they can rely on their words' proper functions to do the communicative work for them, and so they don't need to form complex intentions. Bar-On thinks of expressive communication as an evolutionary and developmental precursor to Gricean communication—one that still shows up in at least some of our own linguistic communication, and that explains how we can do it without complex planning and mindreading. Armstrong (2021, 31, n14) thinks of Gricean communication as a special, intellect-fueled case of the broader category of minded communication that he identifies, and leaves it open just how often humans rely on Gricean communication and how often we instead rely on less demanding mechanisms for coordinating our states of mind.

Finally, Richard Moore (2015; 2017a; 2017c; 2018), has argued that the cognitive requirements of Gricean communication are much less demanding than most have thought, and that many non-human animals are themselves Gricean communicators. On Moore's view, all that is required for Gricean communication is that a communicator "intentionally produce[s] an utterance in order to solicit a response from her interlocutor, and that she intentionally addresses that utterance to her interlocutor as a way of soliciting this response" (Moore, 2017c, 803). However, Moore argues that it is possible for a communicator to meet these conditions without possessing a concept of belief, without "the ability to make complex inferences about others' goal-directed behavior", and without the ability to engage in higher-order metarepresentations (Moore, 2017c, 803). The reason that this is possible, Moore thinks, is that a communicator can minimally satisfy Grice's definition of utterer's meaning by intentionally performing an act with the joint function of (a) producing a state of mind in their addressee and (b) soliciting the attention of the intended addressee and focusing it on the act itself. But intentionally performing an action with a function of producing a certain outcome, on Moore's view, does not entail having an intention or other mental representation of that outcome. Unlike most other Griceans, Moore thinks that our capacities for mindreading and planning aren't essential to Gricean communication, but are add-ons to an evolutionarily ancient mechanism that explain how the it can be augmented to achieve more cognitively sophisticated, cooperative, and efficient forms of communication.

All of the authors I have mentioned in this section deny that linguistic communication in general must involve complex metarepresentations and sophisticated mindreading and planning. They all suggest that communication in young children and our proto-human ancestors could have been supported by the non-Gricean or minimally Gricean mechanisms that they posit, and that at least some communi-

cation between adult humans is supported by the same mechanisms. How much? Each of the authors is vague on this point, but it is compatible with their views that cognitively undemanding mechanisms are the norm even when it comes to adult language use, and that full-blown Gricean communication is a relatively rare strategy that we deploy only when our cognitive resources are abundant and the occasion calls for something special.

3 Functional dependence without diachronic priority

In what follows I will argue that ordinary language use *does* functionally depend on Gricean communication, but it does not follow from this that the evolution of Gricean communication had to precede evolution of language, or that all linguistic communication lives up to the full Gricean ideal. Full linguistic competence depends on sophisticated planning and mindreading, and it is part of the proper function of natural language to play a role in the cognitively sophisticated form of Gricean communication, but we do not always use language competently or in keeping with its proper function, and not all homologues of natural language need appear in Gricean communicators.

The evidence for these claims, which I will lay out in §4, is that there are many communicatively valuable features of natural language that we can't use properly unless we engage in full-blown Gricean communication, with its sophisticated blend of planning and mindreading. Moreover, these aren't exotic features of natural language, but quite ordinary features—things like the noun phrases, gradable adjectives, and the kind of specialized vocabularies that each of us has for talking about the subject matters that are particularly relevant to us. To be a fully competent user of these features—to use them in accordance with their proper functions—requires forming communicative intentions as part of the process of planning our communicative acts and organizing conversations. This is the sense in which our capacity for natural language depends on our capacity for Gricean communication, and in this sense I agree with Origgi and Sperber (2000) about the functional dependence of natural language on mindreading (and, I would add, planning).

However, unlike Origgi and Sperber, I deny that it follows from this line of thought either that our proto-human ancestors had to be Gricean communicators before language could evolve. Nor does it follow that contemporary human language users are always ideal Gricean communicators. To see why these conclusions don't follow, it helps to recognize that neither Gricean communication nor language

is a monolithic trait. Both are integrated systems of more basic capacities. This line of thought builds on previous work in which I have argued that the Gricean program is best pursued by dissecting our capacity for communication into underlying capacities and ultimately cognitive mechanisms (Harris, 2019).⁶ In particular, I have focused on explaining the nuances of human communication by understanding the interactions of our capacities for planning, mindreading, verbal working memory, and language. This approach will likewise animate what follows.

In §5, I will argue that it is plausible that our pre-Gricean ancestors had homologues of our capacities for planning, mindreading, and at least many of the ingredients in our capacity for language. These ingredients may have evolved in parallel for millions of years, playing a variety of functional roles, before becoming further developed and integrated into a package capable of supporting Gricean communication. It was only once our ancestors assembled this package that a wide range of communicatively valuable features of natural language could have developed. (I will discuss these features in §4.) The communicative benefits of this integrated package may have been so great that its emergence generated new evolutionary pressure to further refine and integrate its underlying ingredients.

Although I will argue in §4 that sophisticated planning and mindreading are essential ingredients in fully competent linguistic communication, it needn't follow that all language use is fully competent in this way. Children and some neurodivergent adults may routinely fail to integrate mindreading and planning in the way that an ideal Gricean communicator would, and the same goes for neurotypical adults who are short on cognitive resources or when the stakes of miscommunication are low. For example, psycholinguistic research shows that children, as well as adults who are under cognitive load or who have selective working-memory deficits, find it more difficult to design and interpret speech in a way that accounts for available information about their interlocutors' beliefs and perspectives, instead falling back on their own beliefs and perspective—a kind of “egocentricity heuristic.”⁷ However, even when we rely on this heuristic, we are still attempting to reveal and recognize each other's intentions; it's just that we are doing so in a way that is more likely to lead to miscommunication. Taking this risk is often better than not trying to communi-

⁶This sort of explanatory approach, which could be fleshed out as a version of mechanistic explanation (Bechtel, 2008; Glennan, 2017; Godfrey-Smith, 2008; Machamer et al., 2000), functional analysis (Cummins, 1975), homuncular functionalism (Dennett, 1975), or cognitive-architectural decomposition (Carruthers, 2006; Fodor, 1983), stands in contrast to Grice's original methodology, which revolved around conceptual analysis. I discuss this methodological contrast in detail in Harris (2019).

⁷See, e.g. Keysar 2007; 2008; I will discuss this research in greater detail in §4.

cate at all, and so our reliance on easier substitutes for the full-blown mindreading and planning that would go into ideal Gricean communication may be a form of graceful degradation—a backup mechanism that offers a diminished but often still worthwhile results.⁸ This degraded form of Gricean communication and accompanying miscommunication may be ubiquitous for young children, who have acquired sophisticated grammatical tools but lack the broader cognitive capacities needed to use these tools in fully competent ways.

If the foregoing picture is right, then Millikan, Bar-On, and Armstrong are wrong to think that Gricean communication is reserved for special occasions. And although Moore could be right that the sort of “minimal” Gricean communication that he posits is widespread among animals, I am skeptical that this represents genuine continuity with humans. Rather, if I am right, then the sort of Gricean communication in which humans routinely engage, and that belongs at the center of an explanation of language use, depends in a deep and essential way on our sophisticated capacities for mindreading and planning, and involves genuine communicative intentions rather than merely intentional actions with communicative functions.

An interesting question raised by the foregoing claim is: How “regularly” and how ideally would we have to behave as Gricean communicators in order for natural languages to have evolved as they have? This is an interesting empirical question. Some systems (e.g. hearts) must continually and successfully perform their proper functions in order to continue existing, while others (e.g. sperm) need perform their proper function extremely rarely (Millikan, 1984, 34). My claim that it is part of the proper function of natural languages to be used for Gricean communication therefore entails little about the frequency with which this happens. If the benefits of the odd episode of Gricean communication were high enough, and the downsides of relying on degraded methods of communication low enough, then full-blown Gricean communication could turn out to be relatively rare. However, I think that is unlikely, and that lots of mindreading and planning goes into most ordinary communication. It should become clear why I think this by the end of §4.

4 Why Gricean communication?

The position I have outlined so far depends on the claim that Gricean communication, supported by sophisticated planning and mindreading, is common, required

⁸I have borrowed the term “graceful degradation” from web designers, who often design websites to take advantage of the features of newer web browsers while also remaining readable on older browsers.

for competent language use, and confers significant communicative benefits, despite being cognitively demanding. I turn now to a defense of these claims.

My case will have two parts, corresponding to what I will take to be the two parts of a communicative intention. On the view that I will articulate and defend below, a communicative intention is a complex plan consisting of two simpler intentions standing in a subplan relation. First there is the “effective intention,” which is an intention to produce a mental state (such as a belief) in an addressee. Second, there is the “revelatory intention,” which is an intention to reveal the effective intention to an addressee.⁹ A revelatory intention is a subplan of the effective intention in the sense that it is formed as part of a strategy for achieving the effective intention. (Revealing my intention for you to believe something is a way of getting you to believe it.)¹⁰ A communicative intention therefore has the structure illustrated in Figure 4.




Figure 1: The structure of a communicative intention, which is a complex plan consisting of an effective intention with a revelatory intention as its subplan.

In answering the question why we form communicative intentions, I will there-

⁹My terminology could confuse some readers. Sperber and Wilson (1995) use the term ‘informative intention’ to refer to what I am calling an ‘effective intention’ and ‘communicative intention’ to refer to what I am calling the ‘revelatory intention.’ I prefer ‘effective intention’ because not all effective intentions are intentions to inform. For example, I have elsewhere argued that directive acts (such as requests or commands) are aimed at producing intentions to act, rather than beliefs, in addressees (Harris, 2021). I use ‘communicative intention’ to refer to the entire complex plan that has an effective intention and a revelatory intention as parts, following another Gricean tradition whose usage diverges from that of Sperber and Wilson (e.g. Strawson 1969; Recanati 1986; and Neale 1992). My main reason for preferring this usage is that it will be important for me to have separate terminology for the complex plan and each of its components.

¹⁰The two intentions that I posit within communicative intentions correspond to the first two clauses of Strawson’s (1964) and Grice’s (1968; 1969) three-clause analysis of speaker meaning. My requirement that the revelatory intention be a subplan of the effective intention does the job of Grice’s third clause, but without itself positing further intention.

fore first answer the question why we form effective intentions, and second the question why we form revelatory intentions as subplans of our effective intentions. My claim will be that we form these intentions because doing so provides us with big communicative benefits and unlocks ubiquitous and valuable features of natural language. Since all of the non-Gricean accounts of communication that I have discussed claim that much ordinary communication does not involve communicative intentions, this section will serve as the foundation of an objection to those accounts, and a response to the argument that it is psychologically implausible to posit communicative intentions in most ordinary human communication.

4.1 Why intentions to communicate?

First, consider effective intentions. What advantages does a communicator get by forming an intention to produce a certain state of mind in their addressee? My answer, in brief, is that effective intentions are pivotal steps in the reasoning process that leads from our extra-communicative goals to the performance of specific communicative acts. They arise at the moment in this process at which we have chosen the message we want to communicate and the addressee to whom we want to communicate it, and they are a crucial input to the process of deciding how to get that message across to that addressee. If we didn't do the sort of planning that results in communicative intentions, we would be unable to customize our communicative acts for our addressees in the flexible ways that we do, and both our capacity for communication and our competence with basic features of natural language would be greatly diminished as a result.

To flesh out this answer, it will help to think about the nature of intentions in general. I will build on Michael Bratman's (1987) planning theory of intention. On Bratman's view, intentions typically arise mid-stream in the process of practical reasoning, as elements in the larger plans that we construct in order to turn our abstract goals into specific actions. For example, I currently have an intention to eat lunch in an hour, but I haven't yet worked out any of the details of how I will accomplish this. What will I eat, where, and with whom, for example? These questions will have to be answered before I eat, and answering them will require engaging in practical reasoning that will result in my adoption of further intentions—subplans of my original intention. A subplan of a prior intention is an additional intention about how to accomplish the prior intention. If I form an intention to have pizza, this will be a subplan of my intention to have lunch, for example, since it partially

settles the question of how I'll eat lunch. Normally, any subplan that I adopt will raise further questions of means and implementation: Where will I get pizza, how will I get there, and what kind should I have? Practical reasoning is thus an iterative process, repeatedly proceeding from intention to subplan, resulting in complex, hierarchical plans that are eventually fleshed out enough to act on with specific bodily movements. It is a central part of the functional role of intentions to serve as the waypoints in this process, and as the elements in these complex, hierarchical plans.

Communicative intentions are themselves complex, hierarchical plans with these qualities. A communicative intention is a complex plan consisting of an effective intention with a revelatory intention as a subplan. And this is because communicative intentions arise from practical reasoning of the kind that Bratman describes: We form intentions to reveal our effective intentions to our addressees because doing so is part of a strategy for having the effects that we intend to have on them.

But, of course, an effective intention doesn't arise from nowhere: We decide to change others' states of mind because doing so will help us to accomplish our broader goals. And a revelatory intention is not normally the final output of practical reasoning: After deciding to reveal an intention to someone, we still have to design an utterance that will provide them with the evidence of our effective intention. And so communicative intentions are typically embedded within larger hierarchical plans, and themselves arise midstream in practical reasoning, partway between our abstract, extra-communicative goals and the muscle movements by means of which we try to accomplish them.

If communicative intentions are just thin slices of larger plans, why would they play a privileged role in the theory of communication? There is a straightforward answer to this question. The formation of a communicative intention is the moment in the planning process when the terms of successful communication are set. Actually, there are two kinds of communicative success worth mentioning: It is one thing to be understood by someone, and it is another thing to convince. In the Gricean model, these two forms of success are the satisfaction conditions of the two component-intentions of a communicative intention. To be understood is to have your effective intention recognized (thereby satisfying your revelatory intention), and to convince someone is to produce the state of mind that you effectively intend (thereby satisfying the effective intention).¹¹

What makes all of this practical reasoning worth the effort? Could our interests

¹¹These aren't the only kinds of communicative success worth distinguishing. I discuss some others in Harris (2020).

be served just as well, and with fewer cognitive resources expended along the way, if we were to communicate in a less cognitively demanding way—say, in one of the non-Gricean ways described in Section 2?

There is a very good reason why we would bother to put so much practical reasoning into communicating. By reasoning about what to say to whom and then about how to convey it to them—particularly in light of our information about their beliefs, goals, linguistic capacities, and other states of mind—we are able to customize both *what* we communicate (our messages) and *how* we communicate it (our signals) for our addressees. This sort of customization hugely increases the efficiency, flexibility, and expressive power of human communication, and is a prerequisite for full competence with natural languages of the kinds that humans use.

A crucial property of practical reasoning, as theorists like Bratman have understood it, is that it provides an open-ended channel for our beliefs, including our beliefs about other agents and their states of mind, to intelligently influence our actions. When we choose subplans to flesh out our prior, partial plans, we are under rational pressure to make sure that this new intention fits with all of our other intentions and beliefs. Insofar as we repeatedly succeed at this step-by-step process, we construct complex, hierarchically organized plans made up of many intentions that cohere both with each other and with our best information about our circumstances. We need plans like this in order to intelligently act on abstract goals (whose pursuit would require different specific actions in different circumstances), social goals (whose success depends on what other agents will think or do, including how they will react to what we do), and goals whose fulfillment requires complex actions (which consist of many coordinated action-elements). Each of these features of practical reasoning is ubiquitous in human communication. We often communicate in the service of abstract, social goals, and we do so by means of complex actions with many intelligently organized components. In doing this, we bring to bear all sorts of information, including information about any subject matter about which we wish to communicate as well as information our addressees and their states of mind.

By customizing what we say with our addressees' beliefs in mind, we can steer between the twin hazards of redundancy and incomprehensibility, avoiding both telling people what they already know and saying things that they aren't in a position to understand. By conditioning what we say on our addressees' goals, we can be either cooperative or strategic, divulging or withholding information that would advance their ends.

As Armstrong (2018; 2021) and Bar-On (2013b) both emphasize, non-human animals engaging in non-Gricean communication have some mechanisms by which to avoid redundancy and incomprehensibility in addressee-specific ways. For example, primates of several species are more likely to produce alarm calls when in the presence of kin than when alone or in the presence of non-kin, and some primates continue to issue alarm calls only until the other members of their group have answered with alarm calls of their own (Cheney and Seyfarth, 1985, 2018; Coppinger et al., 2017; Crockford et al., 2017). Chimpanzees are also less likely to address alarm calls to those who, they can see, have direct visual access to the threat (Crockford et al., 2017). But although chimpanzees and other primates have a number of ways of conditioning the production of alarm calls, gestures, and other signals on their addressees' identities, attentional states, and other reliable indicators of their information states, this sort of audience design is not the product of online reasoning about addressees' beliefs, goals, and other mental representations. Rather, primate signals are "*functionally designed* to update audience members' states of mind in characteristic sorts of ways" (Armstrong, 2018, 12). This lack of online mindreading-driven planning severely limits message and signal customization in primates vis-a-vis humans, who can make much more flexible decisions about whether and to whom to issue a given signal, and who can also modify the details of both the contents and forms of our signals in response to subtle features of our addressees' thoughts. And, importantly for present purposes, we routinely make use of this ability, and need to do so in order to make competent use of basic and ubiquitous features of natural language.

For example, consider natural languages' rich noun-phrase systems, which supply us with many alternative ways to pick out a given referent.¹² Suppose that you ask me who I am meeting for lunch next Tuesday. As it happens, I will be dining with my colleague, Sandeep Prasada, and so I need to answer you with a noun phrase that will get you to think about him. Even holding the referent fixed, there are numerous options with which I could reply. Here are a few that could work, depending on the circumstances:

- (1) A psychologist friend.
- (2) The psychologist at Hunter College who works on concepts.

¹²For another example of how competent use of the noun-phrase system must be supported by nonlinguistic social cognition, see Rubio-Fernández (2022).

- (3) The guy whose paper I sent you last week.
- (4) Sandeep Prasada.
- (5) Sandeep.
- (6) Him.

Which should I choose? The answer should depend on specific features of your (i.e., my addressee's) state of mind, and specifically your cognitive and/or perceptual perspective on my intended referent. If I think you've never heard of him, then it might be best to utter an indefinite description, such as (1), because indefinites typically function to introduce novel referents.¹³ If I think that you already know something about Sandeep but don't know his name and aren't already thinking about him, then I should use a definite description, such as (2) or (3). Which one it would be better to use depends on what I take you to already believe about Sandeep, such that putting that information into the description's nominal would be a good way to get you to think about him. If you do know Sandeep's name, then I can utter (4), but I should consider whether you take yourself to be on a first-name basis with him (and perhaps also whether you take me to be on a first-name basis with him), in which case I should go with (5). Finally, Sandeep is currently highly salient to you, either because you're looking at him (or a picture of him), or because we've just been talking about him (and we haven't also been talking about another male), or because (I happen to know) you're actively thinking about him and can infer that I know this, I can just use a pronoun, such as (6). More generally, whenever I have some individual or plurality in mind and I have to choose a noun phrase with which to tell someone about it, I need to bring to bear information about my addressee's antecedent cognitive or perceptual perspective on that referent. If I don't do this, then miscommunication, or at least infelicity, will often result. But every sentence includes a noun phrase, and so this is an utterly ubiquitous task for natural-language users.

The benefits of this kind of design process should be obvious. Our complex noun-phrase system gives us an incredibly flexible system of tools for leveraging our addressee's perspective on an intended referent in order to get them to think something new about it. If we didn't have all of the different kinds of noun phrases that we do, our ability to convey information to addressees in varying states of mind

¹³It is a matter of debate whether indefinites introduce novel referents as a matter of their semantics (Heim, 1983; Kamp, 1981; Karttunen, 1976) or via pragmatic mechanisms (Lewis, 2021).

would be greatly diminished. But this benefit can be had only by on-line reasoning that rationally and flexibly integrates information about addressees' states of mind with our own information about whatever subject matter we want to talk about, all in the service of producing signals that are often syntactically and semantically novel. This requires far more sophisticated on-line reasoning than the sort of audience-conditioned "functionally designed signals" produced by other primates.

The noun-phrase system is just one example of a communicatively valuable feature of natural language that is made possible by our ability to bring the fruits of mindreading and practical reasoning to bear on the process of designing communicative acts. As the last several decades of research in semantics has taught us, language is bursting with context-sensitive vocabulary, including quantifier phrases (Barwise and Perry, 1983; von Stechow, 1994; Neale, 1990), modals (Kratzer, 1977, 1981), gradable adjectives (Kennedy, 2007), taste predicates (Lasnik, 2005), and, arguably, polysemous expressions, which include many (perhaps all) open-class vocabulary items (Carston, 2010; Travis, 2008). Because they can be used to say different things on different occasions, these expressions greatly expand what we can say while economizing on both syllables and lexical items. But this works only because we are the kinds of creatures who can somewhat reliably predict when our addressees are likely to arrive at a different interpretation of the expression than the one we intended, and so avoid using it—something that requires on-line mindreading and planning.

The reader might be skeptical about whether ordinary speakers and hearers really do all of this reasoning. We often don't notice ourselves doing it, and it sounds to some like an implausible amount of work to put in every time we use language. However, the fact that we are not aware of ourselves engaging in a psychological process is not, in general, good evidence that it isn't happening.

There is also a body of empirical research that supports what I have said so far.¹⁴ The picture that emerges from this work is that when we produce or interpret linguistic utterances, we normally do much more reasoning about our interlocutors' perspectives than we are conscious of doing. However, this reasoning is resource intensive, and our capacity to do it can fail in systematic and predictable ways, which leads us to disregard others' perspectives, sometimes resulting in miscommunication.

¹⁴What follows is a quick overview of just one line of relevant empirical work. With more space, I would also discuss work in the "Rational Speech Act" (RSA) framework, which uses game-theoretic Bayesian models to generate impressive predictions about how agents design and interpret utterances by reasoning about each other's states of mind. For an overview, see Degen (2023).

In variations on the “director task,” a subject must either give or follow instructions about how to move an object within a grid that sits between them and their interlocutor. However, the subject can see that some of the objects are visible only to them and not to the interlocutor. For example, consider a situation in which there are three candles of different sizes in the grid, the smallest of which is visible to the subject but not to their interlocutor (Keysar et al., 2000). What will they do if their interlocutor directs them to move “the small candle”? The correct, “allocentric” response would be to move the middle-sized candle, which is the smallest one that the director could have seen, and therefore what they most likely intended. On the other hand, if the subject moves the smallest candle from their perspective, which the director couldn’t see, then they behave “egocentrically.” In the same situation, what if the subject is tasked with directing their interlocutor to move the medium-sized candle?¹⁵ If they describe it as “the medium candle,” they behave egocentrically, and will likely confuse their interlocutor, whereas if they call it “the small candle,” they behave allocentrically. To behave allocentrically in situations like these requires reasoning about the interlocutor’s states of mind—in this case, about their visual perspective and present or future beliefs.

Experimenters have found that most subjects ultimately behave allocentrically in situations like these. However, they do sometimes fixate their gaze on or begin to reach for the egocentric options before correcting themselves, which suggests that they at least sometimes consider the egocentric option before landing on the allocentric one (Keysar et al., 2003; Nadig and Sedivy, 2002; Tanenhaus and Trueswell, 2003). This evidence of egocentric thinking, as well as outright egocentric behavior, is systematically more pronounced in children (Epley et al., 2007), adults with low verbal working-memory capacities or whose verbal working memories are kept busy during the experiment (Lin et al., 2010), and people who have been manipulated to be in a happy mood (Converse et al., 2008). These findings suggest that the perspectival reasoning involved in the task is costly, and that we get worse at doing this reasoning when our cognitive resources run short or when we judge the stakes to be low. On the other hand, when subjects are made to repeatedly interact with a consistently egocentric interlocutor, they compensate for this by becoming less egocentric—a finding that suggests that we reason not just about others’ perspectives but also about their capacity to adopt our own perspective (Hawkins et al.,

¹⁵The director’s task has more often been used to measure addressee responses than speaker behavior, but variations on this kind of elicited-speech task has been done, for example, by Nadig and Sedivy (2002) and by Hawkins et al. (2021).

2021). And when a speaker must direct an interlocutor to move an object in a condition where other objects are hidden from the speaker, they utter longer, more detailed descriptions of the target object, seemingly in order to rule out whatever possible distractor objects could be there (Hawkins et al., 2021). This suggests that speakers routinely reason not just about what they know about their addressee's perspectives, but also about what those perspectives *might* be like. Taken together, all of this research suggests that a good deal of resource-constrained mindreading goes into designing ordinary speech and interpretation—both in order to gather information about what our interlocutors' perspectives are or might be, and in order to make intelligent judgments about when and how to invest the resources needed to do this very reasoning.

4.2 Why do we reveal our intentions?

I have argued that we form effective intentions as part of the process of reasoning about what to say and how to say it to our addressees. But once we have an intention to produce a certain state of mind in an addressee, why do we try to accomplish this by revealing this intention to them? Why do we form revelatory intentions?

The best strategy isn't *always* to be transparent about what I intend. We sometimes use reverse psychology, hiding our real intentions, or provide evidence that will lead our addressees directly to the intended belief without taking a detour through our intentions, for example. But these strategies are often difficult to pull off, and tend to work only in special circumstances. Reverse psychology requires a contra-suggestible addressee, and showing instead of telling is possible only when you have the right kind of evidence available. In many cases, there is no good way to communicate without making your intentions understood. Imagine how much more difficult it would be to accomplish many ordinary communicative tasks while concealing your communicative aim from your addressee. How would you get your partner to have the right belief about when to expect you home from work? How would you get your friend to form an intention to buy you a certain specific drink? How would I get you to believe the sorts of things that I am trying to convey in this chapter?

This line of thought makes it plausible that making our intentions overt often goes hand in hand with easier, more efficient communication. But why, exactly? Why, in particular, would revealing our effective intentions be such a powerful means to their satisfaction?

In short, my answer is that human communication works best when it is a cooperative joint activity, governed by the participants' interlocking shared plans and expectations. Consider another joint activity, like moving a sofa across the room together. Imagine how much more difficult it would be to get someone to do this with you while concealing from them what you intended to accomplish. Perhaps not impossible, but it would normally require much more effort and plotting on your part. That situation is analogous to the problem faced by a communicator who wishes to change their addressees' beliefs without revealing their communicative aims.

Bratman (1992; 2014) has argued that the distinguishing feature of joint action is that it is done under the guidance of a shared intention—an intention on the part of two or more people about what they will do together. Once a group of people have a shared intention in this way, they are under rational pressure to adopt meshing subplans of their joint intention—further intentions that fit together into an inter-subjectively coherent plan about how to satisfy their shared intention. One very good way to avoid clashing subplans is for each agent to publicize their relevant intentions to the others, and to maintain accurate representations of the others' plans and expectations.¹⁶

For example, if you and I are moving a sofa together, things will go best if each of us intends that we move the sofa to the same place, if we form meshing subplans of this intention, and if each of us is aware of the other's relevant intentions and expectations. If I intend to pick up one end of the sofa at a certain time, you should intend to pick up the other end at the same time, and each of us should be aware of the other's intentions. None of this is *required* to get the sofa to its new spot, but we will achieve our goal more efficiently, directing our sofa-lifting energy in more complimentary and focused ways, insofar as we are able to conform to this ideal of joint planning.

Communicative exchanges are analogous. Imagine teaching a class or giving a talk, and think about the amount of cognitive effort that your audience must put into the process of understanding and evaluating what you say in order for the talk to be a success. Think of a time when *you* have been in the audience's position, and consider how much effort, attention, and thought went into understanding and

¹⁶Bratman makes it a condition on possession of a shared intention that the intenders have common knowledge of their shared intention as well as their intention to form meshing subplans of this intention (Bratman 1999, 143–144; 2014, 5, 57–60). I think this is too demanding. I doubt that we ever achieve full-blown common knowledge (Lederman, 2017), and I have elsewhere argued that the role of common knowledge and similar notions in the theory of communication has been greatly exaggerated (Harris, 2020); see also Jankovic (2014).

assessing what you were hearing. Consider the resources that you are expending now, as you read this chapter. A communicator whose addressee is investing this kind of effort in a way that meshes with their own intentions is at an extraordinary advantage, both when it comes to being understood and when it comes to convincing their addressee of whatever they are saying. This advantage will often make it rational to try to communicate more things in more ambitious ways.

By contrast, imagine talking with someone who is downright uncooperative—who won't attend to what you're saying, who constantly interrupts, who pedantically nitpicks everything you say, who makes no attempt to charitably interpret your utterances, pouncing on every ambiguity or unclarity in an attempt to misconstrue your point, who repeatedly claims that you are implying things that you didn't intend, who answers every question in the least informative way possible, who constantly says irrelevant things, or who simply says lots of things that are false, misleading, or nonsensical. These are some of the hallmarks of highly uncooperative communicators. They are analogous, I think, to the tactics that someone might use if they were pretending to help you to move a sofa but were actually trying to thwart your attempt to move it. And of course, if someone is uncooperative enough, they might just refuse to talk (or to move the sofa) at all. There are also milder failures of cooperation. We can be ambivalent about our interlocutors' plans, or lazily fail to track them, in ways that can lead to miscommunication or merely inefficient or unsatisfying conversations.

I have been highlighting some of the communicative benefits of transparently shared conversation plans. There are also valuable features of natural language that we could not competently use without them. Grice (1975) influentially argued that an important part of our strategy for interpreting indirect and non-literal speech is to consider what a speaker must have intended, given their commitment to a shared conversational goal. It's not clear that cooperativity or joint planning is an *essential* precondition of indirect communication.¹⁷ However, it is easy to see how interlocutors' awareness of their shared conversational plans would facilitate indirect communication. If I know that our conversational goal is to decide where to have dinner, and if I assume that your communicative intentions will be a subplan of this shared intention, this will make it much easier for me to recognize your utterances of "Veselka has good pierogis" as an indirect proposal about where we should eat. In general, because we are under rational pressure to make our communicative intentions cohere with and flesh out our shared plans, shared plans are a valuable

¹⁷See Buchanan and Schiller (2022) for some reasons why.

extralinguistic source of information about communicative intentions.

Shared plans can also play an important role in facilitating direct and literal language use. As one example, consider how we use gradable adjectives, such as ‘tall.’ Suppose that Mike and Reggie are looking at a group of basketball players, each of whom is over 6 feet, 5 inches tall, when they have the following exchange:

- (7) Mike: Now, *they* are tall.
Reggie: Is your cousin tall?
Mike: No.

The most influential theories of gradable adjectives tell us that the literal content of Mike’s first utterance is that the basketball players are taller than d , where d is the contextually relevant degree of tallness (Kennedy, 2007). In a situation in which Mike and Reggie already agreed on a standard for tallness, and in which Mike but not Reggie could see how tall the basketball players were, Mike could have used the same utterance *factually*, to inform Reggie of their height. That’s what’s happening in Mike’s second utterance, in which he informs Reggie that his cousin is not tall by the standard on which they have coordinated. But when Mike makes his first utterance, both he and Reggie can see for themselves how tall the basketball players are. In this situation, Mike’s utterance has a *metalinguistic* aim, which is to propose an approximate value for d —a contextual standard for what they will count as tall in this conversation. And it is their coordination on that new standard that allows Mike to then convey some information to Reggie about his cousin’s height.

A number of philosophers and linguists have argued that we often use context-sensitive expressions to make metalinguistic proposals of this kind (Barker, 2002; Ludlow, 2014; MacFarlane, 2016, 2020; Plunkett and Sundell, 2013). Barker (2002) thinks of these proposals as updating interlocutors’ shared beliefs about the conversation’s context. But MacFarlane (2016; 2020) points out that it makes more sense to think of utterances like Mike’s first as proposing shared *plans* about how to use context-sensitive expressions for the foreseeable future of the conversation. After all, there is no independent fact about which height counts as tall, over and above what the participants in a conversation decide, and so there is nothing for them to have metalinguistic beliefs about in this case other than the results of their own decisions—i.e., their plans.¹⁸ If this is right, then we should understand the conversation in (7) as one in which Mike and Reggie first adopt a shared metalinguistic

¹⁸Building on a formal model first developed by Gibbard (2003), MacFarlane goes on to develop a rich, “plan-expressivist” theory of gradable adjectives and other vague expressions, on which our uses

plan, and then rely on this shared plan to make their later utterances intelligible to each other. This gives them a mechanism for packing more information into fewer, shorter utterances, but it would seem to require that they actively track each other's intentions and expectations with respect to the conversation.

This sort of discourse planning appears to be baked into normal language use in many other analogous ways. For example, Karen Lewis (2021, §§4.1–4.2) has argued that it helps to explain our ability to use anaphoric noun phrases. What, she asks, is the difference between (8) and the truth-conditionally equivalent but infelicitous (9)?

- (8) (a) I dropped ten marbles and found all but one.
(b) It's probably under the couch.

- (9) (a) I dropped ten marbles and found nine of them.
(b) #It's probably under the couch.

Lewis argues that by asserting something about a missing marble in (8) (as opposed to saying something that merely entails that a marble is missing in (9)), the speaker gives evidence that they have a discourse plan to work with a new shared representation of an entity (a new “discourse referent”), which will be relevant at later stages in the conversation. The addressee's awareness of this plan helps them to interpret the second utterance by letting them infer which discourse referent to update with new information. And whereas this is all quite transparent in (8), the listener in (9) has to go through a slower repair process to infer the speaker's plan, making (9)(b) infelicitous.¹⁹ In this way, coordination on a shared plan about what or whom to talk about allows interlocutors to efficiently discuss the same topic over the course of a conversation. Again, this looks like a basic design feature of natural language, and one that is much more available to speakers who can organize their conversations around shared and mutually understood plans.

of them typically function as factual–metalinguistic hybrids, both conveying information about the world and proposing more specific plans about which standards to adopt.

¹⁹For Lewis, this account is part of a defense of a pragmatic alternative to dynamic semantics, which encodes updates to discourse referents in the semantics of noun phrases (e.g. Heim 1982; 1983 and Kamp 1981). But we needn't reject dynamic semantics to accept Lewis's idea that anaphora requires joint planning. After all: dynamic semantic theories typically don't tell us anything about what has to be going on in the minds of the participants in a conversation in order for a certain discourse referent to be active and available for update. Lewis offers a plausible answer to this question—namely, that the participants have a shared plan to update a shared representation of an individual, or can accommodate a speaker's presupposition of such a plan.

Finally, consider Craig Roberts' view that conversations are organized around discourse plans that can be represented as questions.²⁰ Roberts argues that we can represent the immediate goal of many conversations as the "question under discussion" (or "QUD")—roughly, the question that the interlocutors are currently attempting to answer with their conversation. By modeling the QUD as set of mutually exclusive and exhaustive propositions that would count as answers, Roberts and others have generated impressive empirical predictions about the semantics and pragmatics of prosodic focus (Roberts, 2012), projective content (Simons et al., 2017, 2010), loose talk and metaphor (Hoek, 2018), disjunction (Simons, 2001), epistemic modals (Beddor and Egan, 2018), attitude verbs (Schaffer, 2007; Yalcin, 2018), and our ability to interpret semantically underspecified expressions in general (Schoubye and Stokke, 2016), among a wide range of other phenomena.

At the center of Roberts' theory is the idea that the QUD determines which utterances are relevant. An assertion is relevant if it partially answers the QUD (meaning that its content is incompatible with at least one antecedently live answer), and asking a question is relevant if it is a subquestion of the current QUD (meaning that a full answer to this new question would be a partial answer to the old one). These ideas are direct applications of the idea that the intention behind a communicative act should be a subplan of the shared plan governing a conversation (Roberts, 2012, 3). Many shared plans can be represented as questions. For example, a plan to discuss where to eat can be represented as the question, "where should we eat?" And a plan to discuss last night's basketball game can be represented as the question, "what happened in the basketball game?" If the latter is our plan, then a reasonable subplan might be to discuss who won—a subplan that itself can be represented as subquestion of the QUD—"who won?"

This helps to explain why models featuring the QUD would make such useful predictions about a range of context-sensitive expressions. In effect, the QUD models interlocutors' information about their shared discourse plans. I have already argued that information about shared plans can be an important source of evidence about a speaker's intentions, particularly when the linguistic evidence that they offer doesn't tell the whole story. And so it's unsurprising that information about what question the speaker was trying to answer would help us to infer what they meant with a context-sensitive expression.

All of this adds up to an argument that a capacity for shared planning is an

²⁰Some important precursors to Roberts' view about the relationship between discourse and planning include Cohen and Perrault (1979); Grosz (1986); Thomason (1990).

important part of being a pragmatically competent language user, and that natural languages themselves are organized to be used by adept social planners.²¹ Of course, all of this works only in situations when the addressee is disposed to cooperate with the communicator and go along with their plans, at least to some extent. This represents a significant limitation of Gricean communication. Getting you to adopt a shared discourse plan will be a good way to get you to recognize what I intend, and getting you to recognize what I intend you to think will be a good way to get you to think it, only insofar as you are disposed to trust and cooperate with me, at least when it comes to this topic.

But the communicative benefits of this sort of cooperation are so great that it gives us very significant incentives to cooperate. These incentives operate on the individual timescale, incentivizing us to cultivate relationships of trust, cooperative reciprocity, and mutual scrutability. On the evolutionary timescale, the communicative benefits of cooperative shared planning give rise to selection pressure that favors a whole package of traits: the basic pro-social dispositions that allow us to coexist in large groups without (too much) conflict, the kind of memories and dispositions that allow us to engage in long-term reciprocity, mental representations of social structure and reputations that allow us to track the trustworthiness of other group members from our information about their past interactions, and—most importantly for my purposes—the kind of mindreading and planning capacities that allow us to coordinate and act on shared plans. Each of the traits in this package confers benefits outside of its role in communication, and could originally have had non-communicative functions. However, once we were able to put them to communicative use, it is plausible that a self-reinforcing evolutionary process resulted.

5 Evolutionary precursors of natural language?

I have argued that Gricean communication is resource intensive, but we routinely do it anyway because of its enormous communicative benefits. As evidence of this, I have pointed to a number of ways in which the design of natural languages presupposes that competent speakers will be Gricean communicators who can engage in sophisticated audience design and joint planning, both of which entail plenty of integrated mindreading and practical reasoning.

In this respect, I agree with those who have posited *some* kind of dependence of

²¹For related ideas of “pragmatic competence,” see Roberts (2017, 435; 2022), and Unnsteinsson (2022).

natural language use on Gricean communication. In particular, it follows from what I have said that our capacity for Gricean communication must have been present before at least some features of natural languages evolved. For example, our pre-Gricean ancestors would not have spoken languages with richly perspectival noun-phrase systems, ubiquitous context-sensitivity, or rich systems of anaphoric dependency and gradable adjectives that we can use to communicate only by keeping track of shared discourse plans. Likewise, we should not expect our pre-Gricean ancestors to have engaged in extended communicative exchanges with relevance norms organized around shared plans to answer a series of shared questions, and we should not expect them to have engaged in a lot highly flexible indirect and non-literal communication. To be clear, my point is not that there are no versions of these phenomena that are possible for non-Gricean communicators; it is merely that our advanced capacities for mindreading, practical reasoning, and cooperative shared planning make extremely powerful versions of these communicative strategies available to us, and that this is why we find them to be ubiquitous in present-day human communication.

Does it also follow from what I have said that our pre-Gricean ancestors could not have used *any* system of linguistic communication, or that they could not have had any traits that are homologous with those that make up the capacity for natural language that we have now? No! To draw that conclusion would be to falsely assume that the features of natural language on which I have focused are essential to any capacity that deserves to be called ‘language,’ or even any trait that would be homologous to such a capacity.

As I argued in §3, language itself is an emergent capacity, arising from the non-trivial interaction of a number of underlying cognitive and anatomical capacities. Even aside from planning and mindreading, natural language requires precise motor control over our mouths, vocal cords, and hands, perceptual capacities well suited to representing the phonetic and prosodic properties of speech, cognitive capacities for representing, encoding, and decoding the phonological, syntactic, and semantic properties of linguistic utterances, a capacity for open-ended conceptual thought that gives us plenty of things to say, and, very importantly, the right kinds of interfaces between all of these capacities to allow us to combine them for the purpose of interpersonal communication. Although there are no non-human animals who possess this complete package, there are plenty of creatures who possess analogues and homologues of at least some of the relevant underlying capacities and interfaces, and it is plausible that other great apes possess homologues of most or

all of them.

There should be little doubt that other great apes possess homologues of the articulatory and perceptual mechanisms with which we produce and consume linguistic utterances, although other apes have a smaller vocal range than humans and may have less voluntary control over vocalization (Fitch, 2010, 327–328). Likewise, we now have ample empirical reasons to posit rich forms of conceptual thought in animals, even if it is also true that humans possess more concepts, and can do more things with them.²² The most controversial questions deal with our phonological, syntactic, and semantic capacities, which allow us to create and interpret signals whose meanings systematically vary with changes to their structures.

If we restrict our attention to animal communication systems, we find evidence for analogues but probably not homologues of human syntax and semantics. Several kinds of non-human animals communicate with syntactically complex signals, including songbirds (Berwick et al., 2011) and humpback whales (Whitehead and Rendell, 2015, 76–84). However, as far as theorists can tell, their syntactic capacities are not paired with corresponding compositional-semantic capacities, and so the meanings of these animals' signals do not depend in systematic ways on their structures (Hurford, 2012, ch.1). On the other hand, Schlenker et al (2016) have recently argued that several species of monkeys communicate with signals that have simple syntactic structures and a corresponding compositional semantics. For example, Campbell's Monkeys sometimes add a suffix-like “-oo” sound to two of their alarm calls, altering their meanings in possibly related ways (Ouattara et al., 2009). But since this added structure only increases the monkeys' alarm-call repertoire from three to five total signal-types, some have doubted that semantic composition is the best explanation, as opposed to the memorization of two additional calls that sound a bit like the others (Hurford 2012, 16–17; Fitch 2016). More recently, evidence of similar call combinations has been found in Chimpanzees, raising the tantalizing (if far from confirmed) possibility of a rudimentary form of syntax and semantics that is homologous, and not merely analogous, to those found in human language (Leroux et al., 2023).

Another promising place to look for homologues of natural-language syntax and semantics is in the mental representations of other primates. Syntactic and semantic properties are not present on the surface of our speech. The fact that our utterances have these properties is grounded in the fact that speakers and hearers

²²For defenses of conceptual thought in non-human animals, see Camp (2009); Carey (2009); Galistel (2011); Quilty-Dunn et al. (2022); Spelke (2022).

represent them as such. The reason that ‘Fred loves Mary’ has a different meaning than ‘Mary loves Fred’ is that speakers represent them as having different syntactic structures that guide semantic composition in different ways. Syntactic and semantic theories are computational models of the systems in our minds that build and use these representations. Rather than insisting that anything homologous to human syntax must be found in an animal communication system, we might instead look for systems in animals’ minds that represent other subject matters in similar ways. Perhaps our ancestors started out using hierarchically structured, semantically interpreted mental representations to think about some specific domain, and then later gradually generalized the use of mental representations to other subject matters, including our own utterances.

Several more specific versions of this idea have been proposed. For example, Cheney and Seyfarth (2007) argue that baboons possess hierarchical mental representations of the social-dominance relations within their groups, and argue that similar mental representations of social hierarchy may have been the evolutionary precursors of human syntax and semantics. Their conclusion is that “the discrete, compositional structure we find in spoken language did not first appear there. It arose, instead, because understanding social life and predicting others’ behavior requires discrete, compositional thinking” (Cheney and Seyfarth, 2007, 272). In a similar vein, others have argued that the precursor to human syntax and semantics was our ancestors’ systems for building hierarchical representations of their plans for complex actions (Steedman, 2009). For example, Planer and Sterelny (2021) emphasize the plans that would have been needed to build the increasingly complex stone tools that we find in the hominin fossil record.

Of course, there is a huge gap between these mental representations and the syntactic and semantic properties of natural-language expressions. As Camp (2012) has pointed out, the baboons’ system for representing social hierarchy is apparently domain specific, in that its syntactic relations have a fixed semantic interpretation, always representing social-dominance relations. In natural language, by contrast, syntactic relations encodes highly abstract semantic relationships, such as predication, and this allows natural language to convey information about an open-ended range of subject matters. We could say something similar about the hierarchically structured plan representations that our ancestors may have used to build stone tools. But it seems at least conceivable that representational systems of these kinds could have gradually evolved to be more domain general—for example, by a process wherein a system of representational vehicles came to have more flexible semantic

interpretations.

A further apparent discontinuity between natural language and these would-be precursors is that natural-language sometimes has recursive syntax—a feature that is central to its productivity and expressive power. On one influential view, it is the capacity to build recursive mental representations that is the truly distinguishing feature of human language, and the evolution of this trait must have been a moment in human evolution that was as sudden as it was consequential (Berwick and Chomsky, 2015; Hauser et al., 2002). This may be so, but it would not follow that the kinds of representational systems that I have been discussing aren't homologues of natural language. Here it is worth considering what sort of cognitive foundation would have had to be in place in order for recursion to have evolved in the way that Chomsky and colleagues have suggested.²³ If no system of discretely recombinable representational units had already existed, how could organisms have suddenly gained an ability to organize those representations into recursive structures? This line of thought should lead us to think of recursion as an upgrade to a pre-existing representational system—albeit a very significant upgrade—rather than the beginning of a wholly new system.

All of this is compatible with the possibility that the origin story of human language involved a stage at which a spoken or signed proto-language was used by our pre-Gricean ancestors.²⁴ Such a proto-language could even emerge as a new function of a pre-existing cognitive capacity before the emergence of Gricean communication. My conclusions in §4 entail certain predictions about the limitations of any such proto-language, but do not rule out a system of communication that married complex syntax and compositional semantics to our articulatory and perceptual systems in a way that unlocked significant communicative benefits. However, if this proto-language did exist, then the emergence of Gricean communication eventually changed its function, turning it into a system that is now intricately designed for the role of revealing and recognizing intentions. There is nothing bizarre about the idea that the biological function of a trait can change over the course of its evolutionary history (Godfrey-Smith, 1994). If that is what happened in the case of natural language, then we should think of it as an exaptation, which was at some point cobbled together from parts with diverse origins.

What I have said is not even a sketch of a full answer to the evolutionary chal-

²³Thanks to Josh Armstrong (p.c.) for helping me to think of things in these terms.

²⁴For versions of this view, see Armstrong (2021); Bar-On (2013b); Fitch (2010); Planer and Sterelny (2021).

lenge. We still need an explanation of how the many cognitive underpinnings of linguistic communication originated, and how they could have reached the degree of advancement and interconnectedness that we now see in humans, all in just the last 6–9 million years. But I hope to have undermined the more drastic evolutionary challenge supposedly faced by Griceans, and that I articulated in §1. Although natural language in its present form is indeed the sort of system that can be competently used only by cognitively sophisticated Gricean communicators, this does not give us good reason to doubt that precursors of the various cognitive and anatomical ingredients that make up our capacity for language could have been evolving prior to or alongside the ingredients of our capacity for Gricean communication. Rather, we are forced to conclude only that once all of these capacities got hooked up together in the right ways in our minds, they gave rise to a new capacity for Gricean linguistic communication. And it is plausible, given what I have said here, that this capacity was so valuable that its presence gave rise to new selection pressure whose effect was to further enhance and connect its ingredient capacities, resulting in a capacity for natural language that has become intricately designed for its present functional role in Gricean communication.

6 Conclusions

Modern humans are Gricean communicators, and we could not competently use natural languages as we now find them if we weren't. This makes linguistic communication a resource-intensive cognitive endeavor, but it is worth it, because the features of natural language thereby made available are extremely valuable.

But it does not follow that the same goes for anything that we should call a language, much less for any trait homologous to natural language, whether the trait in question is a medium for thought or for communication. And so there is no reason why our ancestors must have evolved to be Gricean communicators first, before moving on to the development of natural language.

References

- Armstrong, J. (2018). Provincialism in pragmatics. *Philosophical Perspectives*, 32(5–40).
- Armstrong, J. (2021). Communication before communicative intentions. *Noûs*, early view.
- Bar-On, D. (1995). “meaning” reconstructed: Grice and the naturalization of semantics. *Pacific Philosophical Quarterly*, 76:83–116.
- 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 and Language*, 28(3):342–375.
- Bar-On, D. (2021). How to do things with nonwords: Pragmatics, biosemantics, and origins of language in animal communication. *Biology and Philosophy*, 36(50):1–25.
- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25:1–36.
- Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press.
- Bechtel, W. (2008). *Mental Mechanisms*. Routledge, New York and London.
- Beddor, R. and Egan, A. (2018). Might to better: Flexible relativism and the qud. *Semantics and Pragmatics*, 11(7):1–43.
- Berwick, R. C. and Chomsky, N. (2015). *Why Only Us*. MIT Press, Cambridge, MA.
- Berwick, R. C., Okanoya, K., Beckers, G. J. L., and Bolhuis, J. J. (2011). Songs to syntax: the linguistics of birdsong. *Trends in Cognitive Sciences*, 15(3):113–121.
- Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press, Cambridge, Mass.
- Bratman, M. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2):327–341.
- Bratman, M. (1999). *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge University Press, Cambridge.
- Bratman, M. (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford University Press, Oxford.
- Breheny, R. (2006). Communication and folk psychology. *Mind and Language*, 21(1):74–107.
- Buchanan, R. and Schiller, H. I. (2022). Pragmatic particularism. *Philosophy and Phenomenological Research*, 105(1):62–78.
- Camp, E. (2009). Putting thoughts to work: Concepts, systematicity, and stimulus-independence. *Philosophy and Phenomenological Research*, 78:275–311.
- Camp, E. (2012). A language of baboon thought? In Lurz, R. W., editor, *The Philosophy of Animal Minds*, chapter 6, pages 108–126. Cambridge University Press, Cambridge.
- Camp, E. and Shupe, E. (2018). Instrumental reasoning in non-human animals. In Andrews, K. and Beck, J., editors, *The Routledge Handbook of Philosophy of Animal Minds*, pages 100–108. Routledge, London.
- Carey, S. (2009). *The Origin of Concepts*. Oxford University Press, Oxford.
- Carruthers, P. (2006). *The Architecture of Mind: Massive Modularity and the Flexibility of Thought*. Oxford University Press, Oxford.
- Carston, R. (2010). Lexical pragmatics, ad hoc concepts and metaphor: from a relevance theory perspective. *Italian Journal of Linguistics*, 22(1):153–180.
- Cheney, D. L. and Seyfarth, R. M. (1985). Vervet monkey alarm calls: Manipulation through shared information? *Behaviour*, 94(1/2):150–166.
- Cheney, D. L. and Seyfarth, R. M. (2007). *Baboon Metaphysics: The Evolution of a Social Mind*. University of Chicago Press, Chicago and London.
- Cheney, D. L. and Seyfarth, R. M. (2018). Flexible usage and social function in primate vocalizations. *Proceedings of the National Academy of Sciences*, 115(9):1974–1979.
- Cohen, P. R. and Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212.
- Converse, B. A., Lin, S., Keysar, B., and Epley, N. (2008). In the mood to get over yourself: Mood affects theory-of-mind use. *Emotion*, 8(5):725–730.
- Coppinger, B., Cannistraci, R. A., Karaman, F., Kyle, S. C., Hobson, E. A., Freeberg, T. M., and Hay, J. F. (2017). Studying audience effects in animals: what we can learn from human language research. *Animal Behavior*, 124:161–165.
- Crockford, C., Wittig, R. M., and Zuberbühler, K. (2017). Vocalizing in chimpanzees is influenced by social-cognitive processes. *Science Advances*, 3(11):1–12.
- Cummins, R. (1975). Functional analysis. *The Journal of Philosophy*, 72(20):741–765.
- Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, 9:519–540.
- Dennett, D. (1975). Why the law of effect will not go away. *Journal for the theory of social behaviour*, 5(2):169–188.
- Epley, N., Waytz, A., and Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114(4):864–886.
- von Fintel, K. (1994). *Restrictions on Quantifier Domains*. PhD thesis, University of Massachusetts, Amherst, Amherst, MA.
- Fitch, T. (2010). *The Evolution of Language*. Cambridge University Press, Cambridge.
- Fitch, W. T. (2016). Why formal semantics and primate communication make strange bedfellows. *Theoretical Linguistics*, 42(1–2):97–109.
- Fodor, J. A. (1983). *The Modularity of Mind*. MIT Press, Cambridge, MA.
- Gallistel, C. R. (2011). Prelinguistic thought. *Language Learning and Development*, 7(4):253–262.

- Geurts, B. (2019). Communication as commitment sharing. *Theoretical Linguistics*, 45(1–2):1–30.
- Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press.
- Glennan, S. (2017). *The New Mechanical Philosophy*. Oxford University Press, Oxford.
- Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs*, 28:344–362.
- Godfrey-Smith, P. (2008). Reduction in real life. In Hohwy, J. and Kallestrup, J., editors, *Being Reduced: New Essays on Reduction, Explanation, and Causation*, pages 52–74. Oxford University Press.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3):377–388.
- Grice, H. P. (1968). Utterer's meaning, sentence-meaning, and word-meaning. *Foundations of Language*, 4(3):225–242.
- Grice, H. P. (1969). Utterer's meaning and intention. *The Philosophical Review*, 78(2):147–177.
- Grice, H. P. (1975). Logic and conversation. In Cole, P., editor, *Syntax and Semantics 3: Pragmatics*. Academic Press, New York.
- Grosz, B. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.
- Harris, D. W. (2019). Intention recognition as the mechanism of human communication. In Sullivan, A., editor, *Sensations, Thoughts, Language: Essays in Honor of Brian Loar*, pages 11–37. Routledge, London.
- Harris, D. W. (2020). We talk to people, not contexts. *Philosophical Studies*, 177(9):2713–2733.
- Harris, D. W. (2021). Imperative inference and practical rationality. *Philosophical Studies*.
- Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298:1569–1579.
- Hawkins, R. D., Gweon, H., and Goodman, N. D. (2021). The division of labor in communication: Speakers help listeners account for asymmetries in visual perspective. *Cognitive Science*, 45(e12926):1–47.
- 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.
- Hoek, D. (2018). Conversational exculpation. *The Philosophical Review*, 127(2):151–196.
- Hurford, J. R. (2012). *The Origins of Grammar*, volume 2 of *Language in the Light of Evolution*. Oxford University Press, Oxford.
- Jankovic, M. (2014). Communication and shared information. *Philosophical Studies*, 169(3):489–508.
- Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk, J., Janssen, T., and Stokhof, M., editors, *Formal Methods in the Study of Language, Part 1*, volume 135, pages 277–322. Mathematical Centre Tracts, Amsterdam.
- Karttunen, L. (1976). Discourse referents. In McCawley, J. D., editor, *Syntax and Semantics 7: Notes from the Linguistic Underground*, pages 363–385. Academic Press.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30:1–45.
- Keysar, B. (2007). Communication and miscommunication: The role of egocentric processes. *Intercultural Pragmatics*, 4(1):71–84.
- Keysar, B. (2008). Egocentric processes in communication and miscommunication. In Kecskes, I. and Mey, J., editors, *Intention, Common Ground and the Egocentric Speaker-Hearer*, pages 277–296. Mouton de Gruyter, Berlin and New York.
- Keysar, B., Barr, D. J., Balin, J. A., and Brauner, J. S. (2000). Taking perspective in conversation: The role of mutual knowledge in comprehension. *Psychological Science*, 11(1):32–38.
- Keysar, B., Lin, S., and Barr, D. J. (2003). Limits on theory of mind use in adults. *Cognition*, 89:25–41.
- 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.
- Krupenye, C. (2020). The evolution of mentalizing in humans and other primates. In Gilead, M. and Ochsner, K., editors, *The neural basis of mentalizing: A social-cognitive and affective neuroscience perspective*, pages 107–129. Springer, Cham.
- Lasersohn, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28(6):643–686.
- Lederman, H. (2017). Uncommon knowledge. *Mind*, forthcoming.
- Leroux, M., Schel, A. M., Wilke, C., Chandia, B., Zuberbühler, K., Slocombe, K. E., and Townsend, S. W.

- (2023). Call combinations and compositional processing in wild chimpanzees. *Nature Communications*, 14(2225):1–8.
- Lewis, K. (2021). Anaphora and negation. *Philosophical Studies*, 178(5):1403–1440.
- Lin, S., Keysar, B., and Epley, N. (2010). Reflexively mindblind: Using theory of mind to interpret behavior requires effortful attention. *Journal of Experimental Social Psychology*, 46(3):551–556.
- Ludlow, P. (2014). *Living Words*. Oxford University Press.
- Lurz, R. (2011). Animal mindreading: The problem and how it can be solved. In Andrews, K. and Beck, J., editors, *The Routledge Handbook of Philosophy of Animal Minds*, pages 229–237. Routledge.
- MacFarlane, J. (2016). Vagueness as indecision. *Proceedings of the Aristotelian Society, Supplementary Volume*, 90(1):255–283.
- MacFarlane, J. (2020). Indeterminacy as indecision, lecture 3: Indeterminacy as indecision. *The Journal of Philosophy*, 117(11/12):643–667.
- Machamer, P., Darden, L., and Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1):1–25.
- 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.
- Moore, R. (2015). Meaning and ostension in great ape gestural communication. *Animal Cognition*, 19(1):223–231.
- Moore, R. (2017a). Convergent minds: ostension, inference and grice’s third clause. *Interface Focus*, 7(3):1–9.
- Moore, R. (2017b). Gricean communication and cognitive development. *The Philosophical Quarterly*, 67(267):303–326.
- Moore, R. (2017c). Social cognition, stag hunts, and the evolution of language. *Biology and Philosophy*, 32:797–818.
- Moore, R. (2018). Gricean communication, language development, and animal minds. *Philosophy Compass*, 13(12):1–13.
- Nadig, A. S. and Sedivy, J. C. (2002). Evidence of perspective-taking constraints in children’s online reference resolution. *Psychological Science*, 13(4):329–336.
- Neale, S. (1990). *Descriptions*. MIT Press.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15:509–559.
- Origg, G. and Sperber, D. (2000). Evolution, communication, and the proper function of language. In Carruthers, P. and Chamberlain, A., editors, *Evolution and the Human Mind: Language, Modularity and Social Cognition*, chapter 7, pages 140–169. Cambridge University Press, Cambridge.
- Ouattara, K., Lemasson, A., and Zuberbühler, K. (2009). Campbell’s monkeys use affixation to alter call meaning. *PLoS ONE*, 4(11):e7808.
- Planer, R. J. and Sterelny, K. (2021). *From Signal to Symbol: The Evolution of Language*. MIT Press, Cambridge, MA.
- Plunkett, D. and Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers’ Imprint*, 13(23).
- Quilty-Dunn, J., Porot, N., and Mandelbaum, E. (2022). The best game in town: The re-emergence of the language of thought hypothesis across the cognitive sciences. *Behavioral and Brain Sciences*, forthcoming.
- Recanati, F. (1986). On defining communicative intentions. *Mind & Language*, 1(3):213–242.
- Roberts, C. (2012). Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–69.
- Roberts, C. (2017). Linguistic convention and the architecture of interpretation. *Analytic Philosophy*, 58(4):418–439.
- Roberts, C. (2022). The architecture of interpretation: Dynamic pragmatics and pragmatic competence. *Unpublished manuscript of the 2022 Beth Lecture*, University of Amsterdam.
- Rubio-Fernández, P. (2022). Demonstrative systems: From linguistic typology to social cognition. *Cognitive Psychology*, 139:101519.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75(2):383–403.
- Schlenker, P., Chemla, E., Schel, A. M., Fuller, J., Gauthier, J.-P., Kuhn, J., Veselinovi, D., Arnold, K., Cäsar, C., Keenan, S., Lemasson, A., Ouattara, K., Ryder, R., and Zuberbühler, K. (2016). Formal monkey linguistics. *Theoretical Linguistics*, 42(1–2):1–90.
- Schoubye, A. and Stokke, A. (2016). What is said? *Noûs*, Online First.
- Scott-Phillips, T. (2014). *Speaking Our Minds: Why Human Communication is Different, and How Language Evolved to Make it Special*. Palgrave Macmillan.

- Scott-Phillips, T. (2015). Meaning in animal and human communication. *Animal Cognition*, 18(3):801–805.
- Simons, M. (2001). Disjunction and alternativeness. *Linguistics and Philosophy*, 24:597–619.
- Simons, M., Beaver, D., Roberts, C., and Tonhauser, J. (2017). The best question: Explaining the projection behavior of factives. *Discourse Processes*, 54(3):187–206.
- Simons, M., Tonhauser, J., Beaver, D., and Roberts, C. (2010). What projects and why. *Semantics and Linguistic Theory*, 20:309–327.
- Spelke, E. (2022). *What Babies Know, Volume 1: Core Knowledge and Composition*. Oxford University Press, Oxford.
- Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell, Oxford.
- Steedman, M. (2009). Foundations of universal grammar in planned action. In Christiansen, M. H., Collins, C., and Elderman, S., editors, *Language Universals*, chapter 9, pages 174–199. Oxford University Press.
- 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.
- Tager-Flusberg, H. (2000). Language and understanding minds: Connections in autism. In Baron-Cohen, S., Tager-Flusberg, H., and Cohen, D. J., editors, *Understanding other minds: Perspectives from developmental cognitive neuroscience*, chapter 6, pages 124–149. Oxford University Press, Oxford.
- Tanenhaus, J. E. H. M. K. and Trueswell, J. C. (2003). The effects of common ground and perspective on domains of referential interpretation. *Journal of Memory and Language*, 49(1):43–61.
- Thomason, R. H. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Philip. R. Cohen, J. M. and Pollack, M. E., editors, *Intentions in Communication*, chapter 16, pages 325–364. MIT Press, Cambridge, Mass.
- Tomasello, M. (2006). Why don't apes point? In Levinson, S. C. and Enfield, N. J., editors, *Roots of human sociality: culture, cognition, and 42. interaction*, pages 506–524. Berg, Oxford.
- Tomasello, M. (2008). *Origins of Human Communication*. MIT Press.
- Travis, C. (2008). *Occasion-Sensitivity: Selected Essays*. Oxford University Press.
- Unnsteinsson, E. (2022). *Talking About: An Intentionalist Theory of Reference*. Oxford University Press, Oxford.
- Whitehead, H. and Rendell, L. (2015). *The Cultural Lives of Whales and Dolphins*. University of Chicago Press, Chicago and London.
- Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1):23–47.