

Intentionalism versus The New Conventionalism

DANIEL W. HARRIS*

Hunter College, The City University New York, New York, USA

Are the properties of communicative acts grounded in the intentions with which they are performed, or in the conventions that govern them? The latest round in this debate has been sparked by Ernie Lepore and Matthew Stone (2015), who argue that much more of communication is conventional than we thought, and that the rest isn't really communication after all, but merely the initiation of open-ended imaginative thought. I argue that although Lepore and Stone may be right about many of the specific cases they discuss, their big-picture, conventionalist conclusions don't follow. My argument focuses on four phenomena that present challenges to conventionalist accounts of communication: ambiguity, indirect communication, communication by wholly unconventional means, and convention acquisition.

Keywords: Communication, convention, intentionalism, Grice, indirect speech acts.

Introduction

To what extent is communication a matter of convention, and to what extent is it a matter of hearers recognizing the intentions of speakers? This has been one of the central questions in the contemporary philosophy of language since J. L. Austin and H. P. Grice began debating it in Oxford in the 1940's.

* My thanks to Elisabeth Camp, Michael Devitt, Ernie Lepore, Stephen Neale, Dan Sperber, Una Stojnic, Matthew Stone, Elmar Unnsteinsson, and Deirdre Wilson for helpful feedback on the ideas presented here. Thanks also to the audience at the Philosophy of Language and Linguistics conference at the Interuniversity Center in Dubrovnik in September of 2015, and in particular to Dunja Jutrović, for organizing a wonderful conference and for her editorial patience.

Recently, the conventionalist flame has been stoked with new fuel from linguistics, and the new conventionalism that has emerged as a result has found its most ambitious and philosophically sophisticated expression in Ernie Lepore and Matthew Stone's new book, *Imagination and Convention: Distinguishing Grammar and Inference in Language*. On the strength of innovations from dynamic semantics, coherence theory, and formal pragmatics, Lepore and Stone argue that many aspects of communication that have standardly been taken to involve Gricean mechanisms are actually governed by linguistic conventions that had previously escaped notice.

As for the rest of the supposedly inferential parts of communication—the parts that can't be given conventionalist treatments: Lepore and Stone argue that they shouldn't be understood as aspects of *communication* at all. Instead, they are ways in which we engage our hearers in imaginative thought that needn't have any single goal.

Lepore and Stone's squeeze play is the most comprehensive attack on Gricean approaches to language and communication in decades. They must be stopped! In this paper, I therefore won't focus on their detailed analyses of particular examples. Instead, I'll try to show that these analyses don't add up to convincing defense of their sweeping conclusions about the nature of communication.

Here's the plan. In §1, I formulate the disagreement between conventionalism and intentionalism as precisely as I can. In §2, I draw a distinction between two kinds of cognitive process involved in communication, and show how this distinction bears on the conventionalism–intentionalism debate in a way that can allow us to adjudicate it. Namely: although conventional properties of communicative acts can be interpreted by wholly algorithmic processes, intention-based properties have to be interpreted by inferential processes. In §§3–6, I discuss four phenomena, in order of increasing trickiness for Lepore and Stone, that force hearers to rely on inferential processes, and that therefore threaten conventionalism: ambiguity, indirect communication, wholly unconventional utterances, and the acquisition of new conventions. Although Lepore and Stone have things to say about how conventionalists can handle each of these phenomena, I'll conclude by arguing that it doesn't make sense to say all of these things at once.

1. *Conventionalism and Intentionalism*

Conventionalism is the idea that communicative acts are, essentially, things we do by behaving in conformity to conventions. Intentionalism is the view that communicative acts are, essentially, things we do by acting with certain intentions. By 'communicative acts,' I mean the things that speakers do, and that their addressees have to correctly interpret, in order for communication to take place. I am using the term 'speaker' quite loosely here, since communication needn't involve speech, or even language; it can also transpire by means of "utterances" that

are written, signed, gestured, displayed, or made available to the addressee in any number of other modalities. What is at issue between conventionalists and intentionalists is whether these utterances serve as the vehicles of communicative acts primarily in virtue of the conventions that govern them, or primarily in virtue of the intentions with which they're produced.¹

In order to formulate the terms of the debate with more precision, let's begin by distinguishing two kinds of questions that we can ask about the properties of communicative acts. Specifically, suppose that we are interested in some property ϕ of a communicative act a , such that the a 's addressee must interpret a as having ϕ in order for communication to succeed. Any given communicative act has many such properties. Suppose, for example, that a speaker performs a literal and direct speech act by uttering (1).

(1) He should be turning a profit by now.

In order to correctly interpret the communicative act performed with (1), let us suppose that an addressee will have to interpret the speaker as performing an assertion (rather than a question), as referring to Jeff Bezos (rather than Santa Claus) with 'he', and as making an epistemic claim (rather than a deontic claim) with the modal 'should'. Each of these properties of the speaker's communicative act, as well as various other properties, is a possible value for ϕ , in the way I have set things up.² We can now formulate the guiding questions of our inquiry as follows:

Guiding Questions

Where a is a communicative act with a property ϕ such that a 's addressee would have to interpret a as having ϕ in order for communication to succeed:

¹ It is worth noting that there are also positions about the nature of communicative acts that aren't species of either conventionalism or intentionalism. One example is Timothy Williamson's (2000) idea that an assertion is constituted by the fact that it is subject to the knowledge norm (it's unclear what Williamson thinks about other communicative acts). Another example is Wilfrid Sellars' (1954) idea that the properties of communicative acts, like the properties of intentional mental states, boil down to their conceptual roles. Another example is Robert Brandom's idea (1994), extended by Mark Lance and Rebecca Kukla (2009), that facts about communicative acts boil down to normative facts about how discourse may or ought to proceed. In order to focus on the debate between conventionalism and intentionalism, I will here pretend that these and other nonconventionalist, non-intentionalist alternatives are not live options.

² There is a further question about what interpreting a speech act as having a given property consists in. Interpreting an act as an assertion needn't consist in coming to believe that the speaker has performed an assertion, for example. The concept of assertion may best be understood as a theoretical concept that needn't be possessed by competent communicators, after all. On the simplest version of intentionalism, to recognize that someone has asserted p is just to recognize that they performed it with the intention of getting their addressee to believe p . Conventionalists are free to tell their own story about what interpreting a speech act as having a certain property consists in.

(MQ) THE METAPHYSICAL QUESTION
In virtue of what does a have ϕ ?

(EQ) THE EPISTEMIC QUESTION
How can a hearer come to recognize that a has ϕ ?

These two questions demand answers of different kinds. A theory that answers instances of (MQ) is a metaphysical theory of what constitutes or grounds the properties of communicative acts that play a role in the theory of communication. A theory that answers instances (EQ) is a psychological theory about the kinds of information and cognitive processes that hearers use to interpret speech acts.³

The disagreement between conventionalism and intentionalism is primarily about (MQ), but, as I'll argue below, how we answer (EQ) about particular values for ϕ can have implications for how we can answer (MQ) for those same values, and much of this essay will deal with the relationship between the two. But first let me focus on (MQ), which the most extreme versions of conventionalism and intentionalism will answer, respectively, as follows:

(TC) **Total Conventionalism**

For any property ϕ and communicative act a such that a 's addressee would have to interpret a as having ϕ in order for communication to succeed, a has ϕ in virtue of facts about the conventions that govern the utterance-type with which a was performed in the language in which a was performed.

(TI) **Total Intentionalism**

For any property ϕ and communicative act a such that a 's addressee would have to interpret a as having ϕ in order for communication to succeed, a has ϕ in virtue of facts about the intentions with which a was performed.

Since it's possible for the debate to play out in different ways with respect to different kinds of values for ϕ , there is really a spectrum of positions between (TC) and (TI). For example, a number of authors have defended versions of the view that the content of a context-sensitive expression on an occasion of use is fixed by facts about speakers' intentions on that occasion (Bach 1987, 1992; Kaplan; 1989; King, 2013, 2014), while leaving open the possibility that other properties of communicative acts are fixed by convention. Nonetheless, as I interpret Lepore and Stone, they occupy a position at or very near the total-conventionalist end of the spectrum. My own view lies at the total-intentionalist end of the spectrum, though I won't try to defend total intentionalism in this essay.

How can we adjudicate between conventionalism and intentionalism about particular properties of communicative acts? The answer

³ The importance of this distinction has been vigorously advocated by Stephen Neale (2004; 2005; 2016), who argues that the conflation of metaphysical and epistemic questions about communication is the source of much confusion in contemporary semantics and pragmatics.

depends on the nature of conventions—an issue on which Lepore and Stone defer to David Lewis (1969; 1975).

A convention, according to Lewis, is a kind of self-sustaining solution to a coordination problem. A coordination problem arises when it would benefit the agents in a group to perform actions of a certain kind in the same way, but there are various good candidate ways to choose from. In a situation like this, each of the agents' interests would be best served if the community arbitrarily settles on one way of performing actions of the kind in question, and sticks to it. A textbook example is the choice of which side of the road to drive on: each of us is better off if we all consistently drive only one side, but it doesn't really matter which. Different populations have adopted different conventions to solve the problem.⁴

Clearly, some aspects of how we use language are conventional in at least this sense. Although English-speakers use the word 'science' to talk about science, we could just as well have used 'Wissenschaft' instead. One of Lepore and Stone's main argumentative strategies is to point to a wide variety of other facts about the ways in which we speak that, contrary to what many have thought, are likewise arbitrary, and so conventional, in this sense. I will address these arguments in §4.

2. *Algorithm and Inference*

We also need to think about (EQ)—the epistemic question to be answered by a psychological theory of the processes by which hearers interpret the properties of communicative acts. Specifically, I want to focus on the distinction between cognitive processes that are *algorithmic* and those that are *inferential*.

By an algorithmic process, I mean one that draws on an encapsulated body of information and that can be finitely axiomatized. A paradigmatically algorithmic cognitive process is syntactic parsing—the task of assigning syntactic structures to linguistic perceptual inputs. Although much work remains to be done, phonologists, morphologists, and syntacticians have made impressive progress toward formulating

⁴ Lepore and Stone also go along with Lewis's particular account of the metaphysics of conventions, according to which the members of a group G participate in a convention of ϕ ing by ψ ing just in case it is common knowledge among members of G that they truly believe that they regularly ϕ by ψ ing, that they generally prefer to ϕ by ψ ing (as opposed to ϕ ing by other means), that their beliefs about this regularity gives them a reason to continue it, and that there is at least one other way of ϕ ing that would have served their purposes just as well (Lewis 1975). Lewis's theory of convention has been criticized along various lines, many of which I think are decisive (Burge 1975; Gilbert 1989; Harris 2014; Hawthorne 1990, 1993; Hawthorne and Magidor 2009; Millikan 1998, 2005; Petit 2002; Schiffer 1993; Skyrms 1996, 2010). Still, the major positive theories proposed as alternatives to Lewis's theory can still be thought of as theories of either what causes solutions to coordinations to arise and persist (Millikan 1998, 2005; Skyrms 1996, 2010), or as theories of what grounds solutions to coordination problems (Gilbert 1989; Harris 2014; Miller 2001; Schiffer 1993).

the body of information on which parsing relies, and psycholinguists have built impressive computational models of the processes by which these principles are acquired and deployed in parsing. All of this is possible because the parser is, at least for the most part, a discrete and limited component of the mind: it takes a certain class of inputs and algorithmically transforms them into outputs in a way that is largely insensitive to the agent's personal-level mental states.

By contrast, inferential processes are, to borrow terminology from Jerry Fodor, *isotropic*. In spelling out what this property amounts to, Fodor uses scientific confirmation as a case study, since he thinks that it is a good model for isotropic cognitive processes, such as analogical reasoning, belief fixation, and inference to the best explanation (a.k.a., abduction).

By saying that confirmation is isotropic, I mean that the facts relevant to the confirmation of a scientific hypothesis may be drawn from anywhere in the field of previously established empirical (or, of course, demonstrative) truths. Crudely: everything that the scientist knows is, in principle, relevant to determining what else he ought to believe. In principle, our botany constrains our astronomy, if only we could think of ways to make them connect. [...]

These generalizations about cognitive architecture connect up to the conventionalism–intentionalism debate by way of the following principle: wholly conventional properties of communicative acts can be interpreted by wholly algorithmic decoding processes, whereas intention-based properties of communicative acts can be interpreted only by processes that are, at least partly, inferential. If we can show that inferential processes are needed to interpret a communicative act as having a property of a certain kind, therefore, we will have shown that conventionalism is false about properties of that kind.

The argument for the convention–algorithm link goes as follows. Because conventions are arbitrary, it must be possible for agents to become competent participants in them, whatever that entails, in a finite amount of time. I will call the process of becoming a competent participant in a convention ‘convention acquisition’. It follows from their acquirability that conventions must be finitely axiomatizable: in order to become a competent participant in a convention in a finite amount of time, there must be a finite amount of convention to acquire.⁵ But if a convention is finitely axiomatizable, then it can be interpreted by an algorithmic process that draws on an encapsulated body of information—namely, the axioms, or some equivalent body of information.

This link holds even in the case of linguistic conventions, which are extremely complex. We can take linguistic conventions to be pairings of

⁵ This is a generalization of Davidson's argument for the conclusion that a semantic theory must be finitely axiomatizable (Davidson 1965). It is generalized in two ways: (i) it applies to conventions generally, and not just linguistic conventions, and (ii) I don't assume that acquiring a convention is a matter of coming to know anything, or of entering into any other personal-level mental states.

types of communicative acts with the types of observable utterances by means of which they can be literally performed.⁶ Because languages are productive—they contain an indefinitely large number of conventionally meaningful expressions, many of which have never before been uttered—linguistic conventions must be spelled out by means of a grammar, which is a set of principles that together entail a pairing of utterancetypes and communicative-act types. A grammar can thus be thought of as a finite specification of the conventions governing a language. Within generative linguistics, it is also common to use the word ‘grammar’ to refer to the body information a speaker must come to in order to be a competent speaker of a language, and on which the algorithmic component of their language production and interpretation mechanisms draw.

The fact that intention-based properties must be interpreted by inferential processes follows from the fact that recognizing someone’s intentions is an application of *mindreading*—the process of attributing propositional attitudes to others—together with the fact that mindreading is a paradigmatically isotropic process. There are no normative or practical limitations on the information that I might draw on in order to figure out your beliefs and intentions. If interpreting your communicative act as having a certain property is a matter of recognizing your intentions, it follows that doing so is an isotropic, and therefore inferential, process.

The connection between inference, mindreading, and interpreting communicative acts lies at the core of intentionalist ideas about communication. Here is Stephen Neale making a concise case for these connections, for example.

To interpret is to provide an *explanation*, and the concept of interpretation makes no sense in the absence of a *problem* to be solved. We reflexively generate *hypotheses* about the things we perceive. Nowhere is this more in evidence than when we perceive one another’s *actions*. We act out of *reasons*. To interpret an action is to form a hypothesis about the *intentions* behind it, the intentions that *explain* it. Interpreting a speech act is a special case of this. (Neale 2005: 179)

I agree with Neale that interpreting a communicative act is always and essentially a matter of explaining the speaker’s utterance by inferring their intentions. Although *some* of the processes involved in the interpretation of *some* communicative acts are algorithmic and grammar-driven, other communicative acts cannot be interpreted even in part by means of algorithmic processes. (I’ll make the case for this claim in §5.) Moreover, the algorithmic processes involved in interpreting even highly explicit, linguistic communicative acts serve only to narrow down the range of possible interpretations; the rest of interpretation

⁶ Total conventionalism can be understood as the view that the linguistic conventions governing an utterance-type fully specify the properties of a communicative act that a hearer would have to interpret it as having in order for communication to succeed.

must be inferential. (I won't make the case for this view here, though I will chip away at the alternative view.)

In the next four sections, I will consider four kinds of objection to conventionalism about various properties of communicative acts and, so, to total conventionalism of the sort advocated by Lepore and Stone. Each of these objections revolves around a task involved in interpreting run-of-the-mill communicative acts that looks to be inferential, and so seems to present Lepore and Stone with counterexamples.

3. *Ambiguity*

First, consider the problem raised by ambiguous sentences.

- (2) I forgot how good beer tastes.
- (3) I saw her duck under the table. (Perry 1997: 593)

Each of these sentences has more than one literal meaning, which is to say that it is governed by at least two distinct linguistic conventions. To interpret a literal and direct communicative act performed with one of the sentences requires figuring out which convention is operative.

It is easy to see that choosing between conflicting conventions is an inferential task. In order to understand someone who utters (2), for example, it might help to know whether they've spent the last several years going to AA meetings or merely a bar with nothing but Bud Light on tap. But to know that, it might help to know whether alcoholism runs in their family, whether they have ever used the phrase, 'fake it til you make it', and whether they belong to a fraternity. In short: who knows what sorts of information might be relevant to disambiguating (2)? Likewise for (3): was the utterance made by someone who hangs out with animal lovers, or by a school teacher in 1953 who has been running her students through nuclear drills? You get the point: disambiguation is inferential. In fact, Lepore and Stone accept this (2015: 12). How, then, can they be conventionalists?

The answer, I think, is that disambiguation is the only inferential process that Lepore and Stone take to be involved in interpreting genuine communicative acts. And the fact that interpretation often involves disambiguation does not threaten conventionalism, for the following reason: insofar as disambiguation involves inferring the speaker's intentions, the intention at issue is merely the intention to make use of a certain convention. It's this convention, chosen by the speaker and inferred by the hearer, that does the work of fixing the properties of the communicative act. Once the relevant convention is inferred, the rest of interpretation is just an algorithmic process, guided by grammar. As Lepore and Stone sometimes put it, "Pragmatics can be, at most, a theory of disambiguation; pragmatic reasoning never contributes content to utterances" (2015: 83).

Recategorizing many of the inferential processes involved in interpretation as varieties of disambiguation is one of the main strategies

that Lepore and Stone deploy to defend conventionalism throughout the book. Is it a legitimate maneuver? I confess that the difference between disambiguation and “pragmatic reasoning” that “contributes content to utterances” strikes me as a somewhat tenuous one, particularly as more and more inference gets labeled ‘disambiguation.’ I’ll return to this point.

I also worry that Lepore and Stone’s talk of “pragmatic inference contributing content to utterances” risks conflating the epistemic question (EQ) with the metaphysical question (MQ) in a way that obscures a problem with their view. Pragmatic reasoning can be part of a theory of how hearers interpret communicative acts, but it presumably can’t be a theory of what makes it the case that communicative acts have the properties they do. For that, we must look to that which pragmatic reasoning is aimed at recognizing—namely, the speaker’s intentions. In particular, I see no alternative to the view that, when a speaker uses an ambiguous expression, the fact that one convention rather than another is operative in their communicative act is grounded in their intentions. So even if the only role for inference in interpretation is disambiguation, that still points to a role for the speaker’s intentions in an answer to (MQ), since even if all of the relevant properties of a communicative act performed with an ambiguous sentence are fixed by the operative convention, the fact that a given convention is operative is fixed by the speaker’s intention. Lepore and Stone might concede this point, but I will argue in §4.3 and §6 that doing so creates serious difficulties for them.

4. *Indirect Communication*

Probably the most obvious objection to conventionalism revolves around *indirect* communicative acts, such as the ones that would naturally be performed by uttering the following sentences:⁷

- (4) Can I have the french toast?
 \rightsquigarrow Serve me the french toast.
- (5) Oil prices doubled and demand for consumer goods plunged.
 \rightsquigarrow The doubling preceded/caused the plunge.
- (6) Well, it looked red. [said of a magician’s handkerchief]
 \rightsquigarrow I doubt that it is actually red.

How do we communicate indirect meanings like those indicated under (4)–(6)? The traditional and still-dominant story draws on Grice’s theory of conversational implicature. According to this theory, speakers and interpreters operate under the assumption that, in communicating, they’re engaged in a cooperative endeavor—one that is governed by the Cooperative Principle and maxims of conversation. Given this assumption, if

⁷ I use the symbol \rightsquigarrow to indicate a paraphrase of the most natural indirect meaning.

a speaker produces an utterance that would constitute an uncooperative contribution to the conversation if taken literally, an interpreter will assume that the utterance meant to be taken non-literally, and will search for an alternative, indirect interpretation. The fact that the speaker has implicated thus-and-such—or, more broadly, that they have performed an indirect communicative act of thus-and-such kind⁸—is, according to Grice, “what has to be supposed in order to preserve the supposition that the Cooperative Principle is being observed” (1989: 39–40).

In the stereotypical circumstances in which a speaker would utter (4), for example, it would be an uncooperative waste of everyone’s time for the speaker to merely ask whether they can have the french toast. They already know that french toast is available, since they’ve presumably looked at the menu (their use of the definite article further signals this assumption), and, presumably, they are talking to a waiter, whose job it is to get them whatever food they choose. So, on the assumption that they’re being cooperative, asking about the availability of French toast can’t be *all* they’re doing by uttering (4). Given the circumstances, the obvious further explanation of why they uttered (4) would be that they are trying to get the waiter to bring them some french toast. This is the best explanation of why they produced the utterance that they did, given the circumstances.

Both coming to recognize that an utterance shouldn’t be taken at face value and working out the correct indirect interpretation are, according to Grice’s picture, inferential processes. Depending on factors that can’t be exhaustively predicted, the most natural interpretations of some utterances of (4)–(6) could, in some circumstances, turn out to be the literal ones after all. For example: sometimes it might be clear that a speaker who utters (4) just wants to know what is on the menu. In other cases (4)–(6) might be used to perform indirect communicative acts entirely different than those listed. In response to a confident assertion that the magician’s handkerchief was not red, for example, (6) could be used to implicate that the speaker thinks it *was* red, or might have been. These interpretations depend, in messy and unpredictable ways, on what hearers can work out about speakers’ intentions on particular occasions of utterance.

It follows from this theory that indirect communicative acts have several special properties that can be used as diagnostics. One is that they are *cancellable*, which is to say that a speaker can always follow

⁸ Although Grice makes it clear that he thinks that his theory of implicature can be extended to explain non-assertoric indirect communicative acts, and that, in particular, the maxims of quality and quantity can be generalized to apply acts that aren’t true and don’t aim to inform (1989: 28), he does not work out the details of such an account. Since it is awkward to talk about someone implicating that *q* by saying that *p* when what they’ve done is indirectly requested something by asking a question, I speak more broadly of indirect communicative acts, and assume that a generalized Gricean account could handle them just as well as it handles standard cases of implicature.

up an utterance by which they seem to have conversationally implicated p by clarifying that not- p . There's nothing incoherent about saying 'Can I have some toast? I don't want any, I'm just wondering if it's available', for example, or 'Oil prices doubled and demand for consumer goods plunged, but not in that order'. This distinguishes the content and force of indirect communicative acts from conventional properties of direct communicative acts, which can't be cancelled. (It is incoherent to follow (5) with 'but oil prices haven't doubled', for example.) The explanation is clear: the best explanation of an action, including an utterance, needn't remain the best explanation once new evidence becomes available. Suppose I get up and walk toward the kitchen, for example. You might conclude that my action is best explained by the hypothesis that I intend to get a snack. But suppose you then see me walk straight through the kitchen without stopping, and head toward the bathroom. In light of this new evidence, your best explanation of my original action will change. Cancellation works in the same way: even if your initial best explanation of my utterance was that I indirectly meant q , this explanation is defeasible in light of new evidence, and, unless you have reason to think that my cancellation is disingenuous, what better way to defeat the conclusion that I meant q than to say that I didn't, thereby forcing you to look for a new explanation?

A second property that Grice attributes to indirect communicative acts is that they are *calculable*.

...the presence of a conversational implicature must be capable of being worked out; for even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a conversational implicature; it will be a conventional implicature. (Grice 1989: 31)

Lepore and Stone say some strange things about Grice's calculability requirement. Here, for example, is a passage in which they criticize Gricean pragmatics, seemingly on the ground that it attributes insufficient isotropy to the processes by which hearers interpret indirect communicative acts.

The Cooperative Principle might be a useful way of thinking about the background constraints that inform all of this intention recognition. But Grice's theory of conversational implicature calls for more than just recognizing the speaker's intention according to the Cooperative Principle. It requires the content of the implicature to come directly from your social understanding of the speaker, and not from any of the rest of the rich background that informs intention recognition. (2015: 230)

But Grice places no restrictions on the information that hearers can draw on in interpreting either indirect *or* direct communicative acts. With regard to implicature, Grice makes this clear immediately after outlining the calculability requirement, in a passage that Lepore and Stone quote, in full, twice:

To work out that a particular conversational implicature is present, the hearer will rely on the following data; (1) the conventional meaning of the

words used together with the identity of any references that may be involved; (2) the cooperative principle and its maxims; (3) *the context, linguistic or otherwise, of the utterance*; (4) *other items of background knowledge* and (5) the fact (or supposed fact) that all relevant items falling under the previous headings are available to both participants and both participants know or assume this to be the case. (1989: 31; emphasis added)

Lepore and Stone go on to argue that “Most of Grice’s alleged derivations of conversational implicatures fall far short of establishing that the implicated meaning is calculated the way Grice’s theory calls for” (2015: 230). The implication here seems to be that calculability is a very strong requirement—that the inferences by which implicatures are worked out must be gapless and demonstrative—whereas Grice’s examples of calculations are sloppy and filled with gaps. Lepore and Stone seem to confirm this reading earlier in their book, where they admit that they “interpret these passages [on calculability] rather more stringently than other commentators sometimes do (Lepore and Stone 2015: 22). Lepore and Stone are of course right that Grice’s “derivations” of conversational implicatures are gappy. But I fail to see why this constitutes a problem for Grice’s theory. As Fodor has taught us, inferential psychological processes are not subject to precise computational modeling. And in saying that implicatures must be calculable—that they must, in his words, “be capable of being worked out”—Grice was insisting only that implicatures must be interpreted by means of an inferential process that draws on, along with whatever other information is available, the fact that the speaker said (or made as if to say) thus-and-such, together with the assumption that they were being cooperative.

A final quality of at least most indirect communicative acts is, according to Grice, that they are *nondetachable*, which is to say that two equally good ways of performing (or making as if to perform) a direct communicative act on a given occasion will also be equally good ways of performing the same indirect communicative act(s) on that occasion. If I would be understood as implicating *q* by saying *p* with one sentence, for example, then saying *p* with a different sentence should be an equally good way to implicate *q*.

Like cancellability and calculability, nondetachability follows from the fact that the content and force of an indirect communicative act aren’t matters of convention. If they were, then there should be no reason why two different sentences that could be used to say that *p* would also be equally good ways to implicate the same things, since the conventions governing the sentences’ potentials for indirect communication could differ in arbitrary ways.

All of these considerations have led intentionalists to hold up conversational implicature, and indirect communicative acts more generally, as the most serious hurdle for conventionalist theories of communicative acts (Bach and Harnish 1979: §7.2).⁹ A significant portion of

⁹ Schiffer apparently took this objection to conventionalism to be so obvious and crushing that, in his 1972 book, *Meaning*, he decided that it needn’t be stated

Imagination and Convention is devoted to three strategies for countering the argument from indirect communication, which I will now consider, in turn.

4.1. *The Conventional Defense*

Lepore and Stone's first defense against the argument from indirect communication is to argue that many putative examples of indirect speech are in fact conventional. One way that Lepore and Stone do this is by arguing that many standard examples of indirect communication fail Grice's own nondetachability test. They point out that it is much more natural to request french toast using (4) than with (7), for example.

- (4) Can I have the french toast?
 ~→ Serve me the french toast.
- (7) Am I able to have the French Toast?
 ~→ Serve me the french toast.

But, since (4) and (7) seem to be synonymous with respect to what we can literally and directly say with them, Lepore and Stone conclude that "English speakers somehow know that [(4)], rather than [(7)], is the ordinary formulation of an indirect request" (2015). Lepore and Stone argue that, together with other data, the difference between (4) and (7) supports the conclusion that the indirect-request reading of (4) is built into the conventions of English. In a forthcoming paper, they show how this idea can be worked into a formal-semantic theory (Lepore and Stone, 2017).

Lepore and Stone offer a similar argument about (5), comparing it to (8).

- (5) Oil prices doubled and demand for consumer goods plunged.
 ~→ The doubling preceded/caused the plunge.
- (8) Oil prices have doubled and demand for consumer goods have plunged.
 ~→ The doubling preceded/caused the plunge.

Again, it would be much more natural to implicate that the doubling of oil prices preceded or caused the plunge in demand for consumer goods by uttering (5) than by uttering (8). Grice's explanation for how speakers communicate facts about the order of events described in a conjunction is that they do so by exploiting the maxim of manner, which directs speakers to, among other things, "be orderly". Grice seemed to think that describing events in an orderly way normally requires presenting them in the order in which they occurred. A speaker who

explicitly, instead making (and thus illustrating) the objection indirectly by using the following example sentences to make an unrelated point (1972: 94):

- A: "A necessary condition of someone's meaning that *p* is that he utter a sentence which means '*p*'."
 S: "But then one could never mean that *p* by uttering a sentence metaphorically."

utters (5) can thus normally be understood as implicating something about the order of the events they're describing.¹⁰

But as Lepore and Stone point out, an utterance of (8) presents the same events in the same order, and yet the implicature that they occurred in that order is far less natural. Again, Lepore and Stone take this to be evidence that something grammatical is responsible for the order effect communicated by (5), and they defend a dynamic-semantic treatment of the simple past in order to implement this idea (2015: ch.7).

With nondetachability out of the way, what about calculability and cancellability? Lepore and Stone simply deny that indirect readings are calculable in anything like Grice's sense. They argue that we rely on language-specific information in order to work out what someone indirectly means by (4)–(6), for example. And they have a nice explanation for why the indirect reading of (4), for example, is cancellable, which is that (4) is semantically ambiguous between a pure question reading and an indirect-request reading. What looks like the cancellation of an implicature—'not that I want french toast, I'm just curious about the menu'—is actually a disambiguating clarification.

I find myself convinced by Lepore and Stone that the very natural indirect readings of (4) and (5) are more conventional than intentionalists have tended to think. But we must be careful about the big-picture lessons we draw from these arguments. And Lepore and Stone draw some very bold lessons from these and a few other case studies in conventionalizing indirect communication. They argue that indirect communicative acts, understood as communicative acts that must be interpreted by means of the kind of inferential processes posited by Grice, do not exist: "We have no use for a category of conversational implicatures, as traditionally and currently understood" (2015: 6).

Not so fast! A full defense of conventionalism against the argument from indirect communication would require giving a grammatical account of every purported example of indirect communication that intentionalists can cook up. In order to achieve their grand philosophical goal of expunging implicature from the theory of communication, Lepore and Stone find themselves in the always-tricky dialectical situation of having to defend a universal generalization one case at a time. Even if we accept their accounts of the examples they discuss in their book, there are always more examples.

In fact, I think we can use Lepore and Stone's own examples, (7) and (8), against them. All that's needed is to find situations in which, for whatever reason, one of these sentences can be used to indirectly communicate in the relevant way. By filling in the surrounding context in the following way, I think we can accomplish this for (8):

¹⁰ Some more recent intentionalists have argued that enriched uses of 'and' are not best understood in terms of implicature. See, for example, Carston (2002: ch.3). These views are nonetheless incompatible with Lepore and Stone's conventionalism, since they entail that the process by which 'and' is enriched is inferential.

Ernie is eating lunch with his friend Paul, who has a casual interest in economics. Ernie takes out his phone and glances at his stock app, looks distressed, and blurts out: “Darn! All of my stock in Samsung, Nike, and Proctor & Gamble has been tanking!” Paul looks thoughtful for a moment and responds: “[It’s because of Iraq.]

- (8) Oil prices have doubled and demand for consumer goods have plunged.
 ↳ The doubling caused the plunge.

In this case, I think the causal reading of ‘and’ is quite natural. But it is a premise in Lepore and Stone’s own argument for the conventional temporal reading of (5) that there is no such conventional reading of (8). The reading in this particular context must therefore be an intention-based property, to be interpreted inferentially.

I think this point exposes a serious flaw in Lepore and Stone’s methodology and choice of examples. In choosing (4)–(6), Lepore and Stone build their arguments around examples that admit of indirect readings very naturally, and without situating them in detailed extralinguistic contexts. The indirect readings of these examples are so natural, in fact, that it shouldn’t be too surprising that they are good candidates for conventional treatments. It is natural to think of the indirect readings of (4)–(6) along the lines of what Grice dubbed ‘generalized conversational implicatures’—cases in which “the use of a certain form of words in an utterance would normally (in the absence of special circumstances) carry such-and-such an implicature or type of implicature” (1989: 37). As Grice himself points out, “Noncontroversial examples are perhaps hard to find, since it is all too easy to treat a generalized conversational implicature as if it were a conventional implicature” (1989: 37).

The really tough nuts for Lepore and Stone to crack will be what Grice called ‘particularized conversational implicatures’ and other highly idiosyncratic indirect communicative acts that can be interpreted only due to highly specific features of the context in which they are performed. As the contextualized example of (8) that I’ve just given illustrates, indirect communicative acts of this kind can’t normally be detected simply by glancing at a numbered example and imagining up a stereotypical context in which they might have been uttered.

One way to see this would be to find examples of sentences that, in different real-world contexts, could be used to perform different indirect communicative acts. An example:

- (9) It’s 6AM.

First, imagine my wife uttering (9) as she wakes me up on a morning when I have to catch an 8:30 flight. In this context, it would be natural for me to interpret my wife as indirectly suggesting that I should get out of bed. Second, imagine my wife uttering (9) after emerging from

our bedroom on a morning when I've been up all night writing. In this context, it would be natural for me to interpret my wife as indirectly demanding an explanation for why I hadn't slept. Third, imagine my wife uttering (9) as we hike up a mountain on a morning when, we both know, the sun is scheduled to rise behind us at 6:02. In this context, it would be natural for me to interpret my wife as indirectly informing me that the sun is about to come up. Clearly, none of these communicative acts is linked to (9) by the grammar of English. I need all sorts of background information in order to understand my wife in any of these three cases.

These particularized indirect communicative acts, along with metaphorical utterances, flatly-intoned but contextually obvious sarcasm and irony, and the like, make for the strongest case against Lepore and Stone's big-picture conclusions. But this is where their next two defensive strategies come into play.

4.2. *The Imaginative Defense*

Lepore and Stone's second defense against the argument from indirect communication is to deny that many purported cases of indirect communicative acts are really communicative acts at all. This argument is perhaps most convincing as applied to certain metaphorical utterances, such as (10).

(10) Juliet is the sun.

What communicative act is Romeo performing when he utters (10)? As scholars of metaphor have long pointed out, it's hard to say, because metaphors are hard to paraphrase. But as Lepore and Stone point out, this presents a serious problem for intentionalists who wish to claim that Romeo is performing any communicative act at all by uttering (10). After all: in order to perform a communicative act, Romeo has to have a meaning intention of the kind posited by Grice. A meaning intention is an intention both (a) to produce a response (such as a belief in a certain proposition) in one's addressee, and (b) to get them to recognize that one intends to produce this response. But if Romeo is trying to produce a belief in his addressee by uttering (10), it's not going to be clear to them which belief this is. Communication is therefore not likely to succeed. Moreover, it should be obvious to Romeo that this is the case. But intending to do something that you don't think you can do is either impossible or just irrational. It follows that it is either impossible or irrational for Romeo to have any sort of meaning intention in uttering (10), and that it is therefore impossible or irrational for him to perform a communicative act.

By generalizing this version of what Schiffer has dubbed the 'meaning-intention problem' (Schiffer 1992, 1994), Lepore and Stone argue that a wide variety of purported indirect communicative acts, including those involving metaphor, sarcasm, irony, joking, hinting, insinuation, and various other apparent cases of implicature, cannot be communicative acts after all, since they can't be backed by genuine meaning

intentions. Lepore and Stone argue, instead, that the speakers who produce these utterances are doing something that is rather different than attempting to communicate in the usual, Gricean sense.

...we argue that such utterances are better characterized as invitations to the audience to follow a specific direction of thought in exploring the contributions of the utterance. This is the “Imagination” part of our title. The thinking involved is heterogeneous and diverse—it’s not just a circumscribed or uniform application of principles of rationality. And so, the insights interlocutors get by pursuing and appreciating this thinking also fall outside the scope of pragmatics as traditionally conceived. (Lepore and Stone 2015: 4–5)

In particular, Lepore and Stone argue that there’s no one right way to interpret metaphorical and other “imaginative” utterances, and so it doesn’t make sense to talk about communication either succeeding or failing when it comes to them. “Any conclusions the audience thereby discovers are implicit and tentative suggestions, rather than transparent and public contributions” (2015: 39).

Some of what they say suggests that Lepore and Stone would wish to apply this treatment even to my three examples of my wife uttering ‘It’s 6AM’ in different contexts. Would she really be indirectly communicating with me in those cases, or simply pointing out the time and thereby inviting me to draw my own conclusions?

These considerations can be resisted in various ways. One is to insist that genuine communication can, at least sometimes, happen via metaphor and other unconventional indirect speech. One way to do this is to point out that even nebulously unparaphrasable metaphors can be *misinterpreted*.¹¹ Suppose, for example, that Juliet takes Romeo to have uttered (10) in order to imply that she is huge, gaseous, and dangerous to get close to. Clearly, she will have misinterpreted Romeo’s utterance. Similarly, if my wife awakes to find me at the end of an all-night writing binge and utters (9), I will have misinterpreted her if I conclude that she was trying to congratulate me on my manic state. But the possibility of misinterpreting these acts can arise only if it also makes sense to talk about some interpretation (or a range of interpretations) as *correct*.

Another way to respond to Lepore and Stone’s argument is to point out that it overgenerates, since it consigns not only many purported indirect communicative acts to the imaginative waste bin, but also many seemingly perfectly good direct communicative acts as well. That the meaning-intention problem affects direct as well as indirect communicative acts has been pointed out by Ray Buchanan, who imagines Chet addressing (11) to Tim before their party:

(11) Every beer is in the bucket. (Buchanan 2010: 347)

Clearly, Chet doesn’t mean that every beer *in the universe* is in the bucket; the quantifier he expresses using the DP ‘every beer’ is restricted in

¹¹ This point is influenced by discussion with and unpublished work by Elisabeth Camp.

some way. But *how* is it restricted? Does he mean that every beer *in the house* is in the bucket, that every beer *that Chet and Tim bought* is in the bucket, that every beer *that they planned to drink at the party* is in the bucket, or something else? These are all equally good, non-equivalent ways of restricting Chet's quantifier. But that means that there is no single candidate that Tim can reasonably take Chet to intend, and that means that Chet can't reasonably intend for Tim to interpret him as meaning any of them.

This is the same problem raised by Lepore and Stone for metaphor, but now in the case of a much less mysterious-seeming utterance. Moreover, if the problem turns up here, we should expect to find it all over the place: quantificational expressions of various kinds, including DPs, modals, conditionals, generics, adverbs, tense morphemes, and so on, are ubiquitous in natural language, and they are always, or nearly always, implicitly restricted. By parity of reasoning with Lepore and Stone's conclusion that we don't ever communicate with metaphor, should we conclude that we don't ever genuinely communicate with quantified sentences either?

Similar concerns apply to many other common expressions, including possessives ('my horse'), gradable adjectives ('tall'), predicates of personal taste ('delicious'), plural pronouns ('we'), and so on. The leading semantic treatments of these expressions tell us that they install hidden variables in the LFs of sentences in which they appear, and that these variables must be "saturated by context" in order for those sentences to express propositions. But, as many others have argued (Bach 1987; King 2013, 2014; Neale 2004, 2005), it's hard to see how anything but the intentions of the speaker could, in general, do the job of fixing the values of these variables. The problem is that there will often be many non-equivalent candidate saturations that will seem equally natural to a hearer, and so none that a speaker can uniquely intend.

In short: if the meaning-intention problem establishes that we can't communicate indirectly, then it also seems to establish that we almost never communicate directly either. But this is a dark path down which, I suspect, Lepore and Stone do not wish to lead us.

Luckily, Buchanan (2010) has a positive suggestion about how to respond to the meaning-intention problem. He argues that we should loosen the conditions on successful communication to some extent. The contents of communicative acts aren't propositions, he suggests, but properties of propositions; for communication to succeed, a hearer need only come to believe that the speaker performed a communicative act whose content is some proposition or other with that property. Some aspects of this solution are underdeveloped—what is it about the speaker's intentions or other mental states that determines exactly which property they mean, for example?—but it seems to be a more promising first step toward a solution than giving up on whole categories of communication altogether. Buchanan's strategy also gives us an explanation of

how metaphorical communicative acts can be quite open-ended in the conditions of their correct interpretation, while at the same time having conditions of correct interpretation (and, so, misinterpretation).

Indeed, a careful reading of Grice reveals that he anticipated a version of this problem as it applies to conversational implicature, as well as something in the neighborhood of Buchanan's solution (if a bit more inchoate), and that he seemed untroubled by the issue.

Since, to calculate a conversational implicature is to calculate what has to be supposed in order to preserve the supposition that the Cooperative Principle is being and there may be various possible specific explanations, a list of which may be open, the conversational implicatum in such cases will be disjunction of such specific explanations; and if the list of these is open, the implicatum will have just the kind of indeterminacy that many actual implicata do in fact seem to possess. (1989: 39–40)

Like Buchanan, Grice seems to have thought communicative success does not require there to be a single, precise propositional content intended by the speaker and recognized by the addressee. The idea that communication is about coordination in this strict sense is, at best, an idealization that Grice knowingly built into his model. Of course, it is an open question how intentionalists should lift this idealization, but I can't think of a pressing reason to think that this can't be done.

4.3. *The Novelty Defense*

Suppose that we resist letting Lepore and Stone toss every example of indirect communication into either the convention box or the imagination bin. Still, they have one more line of defense, which is to posit new conventions that we make up and acquire on the fly.

Even in these cases [of non-conventional, indirect communication], however, we suspect that hearers must make an intuitive guess about a conventionalized indirect speech act, and so the listener's inference lacks the content of a Gricean calculation. On the ambiguity view, the question to ask

(106) What is a plausible convention that I could postulate to assign this utterance a likely intended interpretation?

It's not the Gricean question (107).

(107) How do principles of rationality and collaboration explain the creative use to which the speaker has put this utterance with its known interpretation? (2015: 105)

My first reaction to this defense of conventionalism is that it is ad hoc. To cover just those situations that can't be explained by appeal to pre-existing conventions, Lepore and Stone argue that we should simply posit new ones. Although I think this initial worry holds some force, I don't think that it is the best reason to be suspicious of the novelty defense.

The best reason, I think, is that the defense simply replaces one inferential task—that of answering (107)—with another, equivalent, inferential task—that of answering (106). Indeed, answering these two questions will require roughly the same information and the same sorts

of inferences. After all: a convention is an arbitrary pairing of utterance-types with communicative-act types. Learning a new convention on the fly is therefore a matter of being presented with an utterance and, without prior knowledge of the convention, figuring out what kind of communicative act is being performed with it. But this is indistinguishable from the task of simply interpreting the communicative act inferentially. In particular, since the fact that a given convention is operative is itself grounded in the speaker's intention to make it operative—a point I argued in §3—inferring that a given convention is operative is a kind of intention recognition. I will return to this criticism in §6.

In any case, it looks as though we've found another inferential task, aside from disambiguation, that regularly figures in the interpretation of communicative acts, thus contradicting Lepore and Stone's claim that "Pragmatics can be, at most, a theory of disambiguation; pragmatic reasoning never contributes content to utterances" (2015: 83). Acquiring a new convention on the fly is not the same thing as disambiguating between two conventions that one has already acquired. Right?

I am worried that Lepore and Stone might try to resist this point. They might reply that, after all, disambiguation and on-the-fly convention acquisition are both a matter of figuring out which convention is at work in a given communicative act. Convention acquisition, they might insist, is just *total disambiguation*—disambiguation from an indefinitely long list of meanings. I won't protest if Lepore and Stone want to use the word 'disambiguation' in this way,¹² but in that case I think that their distinction between disambiguation and "pragmatic reasoning" that "contributes content to utterances" is no longer an interesting one. As I have just pointed out, after all, interpreting a communicative convention by acquiring a new convention on the fly is approximately the same task, requiring approximately the same sorts of inferences, as simply interpreting the communicative act using intention recognition.

5. *Unconventional Communication*

A consequence of Lepore and Stone's position is that it entails the impossibility of communication by unconventional means. The only genuine communicative acts, they argue, are those performed with the aid of conventions. Everything else is an invitation to an imaginative jam session where there are no right or wrong answers.

The idea that there is no genuinely unconventional communication runs counter to the whole spirit of the intentionalist project as outlined in Grice's work. Grice used the labels 'utterer's meaning' and 'nonnatural meaning' for the central concept in his theory of communication—what I have been calling 'performing a communicative act'.¹³ This concept is

¹² Hey look: I've just acquired a new convention!

¹³ Strictly speaking, 'nonnatural meaning' is broader in scope than utterer's meaning, since it also includes utterance-type meaning and utterance-occasion

now usually called ‘speaker meaning’, but that label is misleading: Grice wanted his theory to apply to nonlinguistic and linguistic communication alike, and he took essentially the same psychological mechanisms to be at the core of both phenomena. Grice used ‘utterance’ in a technical sense, “as a neutral word to apply to any candidate for [non-natural meaning]” (1989: 216). In other words: any behavior that might be produced for communicative purposes could be an utterance—an act of pronouncing a sentence, sure, but also a gesture, a raised eyebrow, the act of passive-aggressively doing the dishes, or just the right dance move. Grice’s broad use of ‘utterance’ is illustrated by his choice of examples of non-natural meaning (“meaning_{NN}”) in his 1957 article, ‘Meaning’—the original articulation of the intentionalist project:

- (G1) Those three rings on the bell (of the bus) mean that the bus is full. (1989: 214)
- (G2) That remark, ‘Smith couldn’t get along without his trouble and strife’, meant that Smith found his wife indispensable. (1989: 214)
- (G3) I draw a picture of Mr. Y [displaying undue familiarity to Mrs. X] and show it to Mr. X. ...[T]he picture (or my drawing and showing it) meant_{NN} something (that Mr. Y had been unduly familiar), or at least that I had meant_{NN} by it that Mr. Y had been unduly familiar. (1989: 218)
- (G4) If I frown deliberately (to convey my displeasure), an onlooker may be expected, provided he recognizes my intention, ...to conclude that I am displeased. [Grice goes on to argue that this case counts as meaning_{NN} provided that the frowner intends the addressee to conclude that the frowner is displeased via the recognition of the frowner’s intention.] (1989: 219)
- (G5) If...I had pointed to the door or given him a little push, then my behavior might well be held to constitute a meaningful_{NN} utterance, just because the recognition of my intention would be intended by me to be effective in speeding his departure. (1989: 220)
- (G6) ...a policeman who stops a car by waving. (1989: 220)
- (G7) ...if I cut someone in the street, I do feel inclined to assimilate this to the cases of meaning_{NN}, and this inclination seems to me dependent on the fact that I would not reasonably expect him to be distressed (indignant, humiliated) unless he recognized my intention to affect him in this way. (1989: 220)
- (G8) If my college stopped my salary altogether, I should accuse them of ruining me; if they cut it by one pound, I might accuse them of insulting me [This example immediately follows the previous one, and the implication is that the latter case is an example of meaning_{NN}.] (1989: 220)

meaning. But Grice often speaks of ‘nonnatural meaning’ when he’s talking about utterer’s meaning, including in the examples discussed below.

Of these examples, only (G2) involves a linguistic utterance, and only (G1), (G4), and (G6) are plausibly conventional in nonlinguistic ways. Even (G4) and (G6) don't seem to be *essentially* conventional. If there were no convention of deliberately frowning in order to convey displeasure in my community, it seems likely that the non-conventional relationship between involuntary frowning and displeasure would allow one to communicate displeasure via an obviously deliberate frown, given the right context. Likewise, it is possible to imagine the policeman in example (6) getting his point across by waving in a hitherto-unconventional way, given the right circumstances.

The fact that Grice uses examples of communicative acts that (mostly) don't involve language or conventions is not an accident. In justifying his distinction between natural and nonnatural meaning, for example, Grice argues that it does a better job of capturing "what people are getting at when they display an interest in a distinction between 'natural' and 'conventional' signs", in part because "some things which can mean_{NN} something...are not conventional in any ordinary sense" (1989: 215). What Grice's examples share is just that they are all intelligible ways, in their respective contexts, of providing the addressee with evidence of the speaker's intentions. Although it may often be easiest to provide this evidence by exploiting conventions, there are sometimes other ways that will work well enough. Grice's theory posits the same underlying psychological mechanisms to explain both convention-aided and unconventional cases of communication.

What can Lepore and Stone say about Grice's examples—in particular, the clearly unconventional cases: (G3), (G5), (G7), and (G8)? It seems to me that they have two options: either these aren't genuine communicative acts, or they involve the on-the-fly creation and acquisition of novel conventions. I don't like either of these options—the latter for reasons that I touched on in §4.3 and that I will address in greater detail in §6, and the former for reasons that I will briefly take up here.

My main problem with the thesis that we don't communicate in unconventional ways is that this idea doesn't fit with what we know about the human appetite and aptitude for mindreading. Our drive to take the intentional stance toward the world around us—to interpret and predict others' actions in terms of what's going on in their minds—is so constant and routine that it may be easy not to notice the enormous role that it plays in our lives. But try reading a newspaper or a novel, or watching a movie, without attributing beliefs and intentions to the real or fictional agents in the stories. Try planning dinner, going shopping, caring for children, conducting market research, negotiating a cease-fire, or designing public policy without a constant and mostly reliable stream of information about the mental states of those around you. If our mindreading abilities were all suddenly switched off, things would get ugly fast.

This is because the ability to attribute mental states to others and to oneself—including higher-order mental states—is the basic capacity required for any sort of intelligent social interaction. Some have argued that the capacity for mindreading and the social intelligence that came with it are the central functions for which human brains evolved to be so large and costly in the first place (Dunbar 1998, 2003). Others have argued that the capacity for mindreading was an essential precondition for the evolution of language (Scott-Phillips 2014). Whether or not those theses are true, a great deal of evidence has recently emerged to support the view that mindreading is an innate capacity in humans—one that develops in all neurotypical humans in early infancy.¹⁴ The ability to detect others' beliefs and distinguish them from one's own has been observed in children as young as ten months (Luo 2011), for example, and the ability to attribute intentions to agents in order to explain their behavior has been observed in infants as young as three months—an age below which we don't have any viable experimental paradigms because infants' muscles and eyesight aren't sufficiently developed (Sommerville et al. 2005).¹⁵ Intention-recognition is thus among the earliest higher cognitive capacities to come online in babies—a fact that is perhaps unsurprising, given that getting over many other developmental hurdles requires infants to engage in socially intelligent ways with others.

I doubt that Lepore and Stone would find much to disagree with in this mixture of platitudes and well-supported cognitive science. But they seem not to appreciate the consequences of it. Interpretive inferences of the sort posited by Grice to explain both linguistic and nonlinguistic communication is just a special case of the same old mindreading in which we continuously and reflexively engage as we interact with other agents. Specifically, it is the special case that arises when our already-hyperactive drive to read agents' minds is intentionally initiated and guided by the very agents whose minds we're already trying to read. But, surely, an interpretive capacity that we engage in with such frequency and success does not become deeply mysterious just when the very people we're trying to interpret shape their own behavior with the goal of making it easier for us.

Lepore and Stone might try to respond by arguing that Grice posits a *special* kind of mindreading in his theory of implicature, so that we can't show that Gricean explanations make sense just by showing that there's nothing mysterious about intentionally getting someone to read one's mind. In their discussion of Grice, they tend to focus on the fact that implicatures are supposed to be interpreted by relying on the

¹⁴ For an excellent summary of this evidence, see Carey (2009).

¹⁵ Intentions are commonly referred to as 'goals' in the psychological literature. Some have argued that the states that infants attribute are not quite beliefs and intentions but simpler counterparts of those states (Butterfill and Apperly 2013), but Carey presents evidence that can't be accounted for by this sort of "minimal" model of infant mindreading (2009, 166–170).

Cooperative Principle and the maxims. Toward the end of the book, they summarize what they take themselves to have proven as follows, for example.

There are no special meanings, over and above the meanings of our utterances, that interlocutors infer by calculation from a Cooperative Principle, maxims of conversation, or other general principles for pragmatic enrichment and reinterpretation. (2015: 199)

But the idea that pragmatic inference is guided the Cooperative Principle does not make it a uniquely mysterious form of mindreading. Rather, the Cooperative Principle is just a natural assumption for me to make as an addressee, given the assumption that the speaker is actively trying, in good faith, to trigger and guide my instinct to read their mind. If I make this assumption about a speaker, and if the speaker makes the analogous assumption that I am attempting in good faith to recognize their thoughts, then we're acting in accordance with something like the Cooperative Principle. This isn't to say that we must always act in accordance with it; sometimes we know that a hearer is being willfully obtuse, for example. But of course, these are precisely the occasions on which indirect communication tends to fail.

This broad line of thought gives us a very powerful reason to reject Lepore and Stone's view that most apparent examples of indirect and unconventional communication are actually mere invitations to mere imaginative reflection. It also gives us a new reason to be suspicious of the argument that Lepore and Stone use to defend that view. In particular, if we think of communication as nothing more than deliberately triggered mindreading, we find further reason to accept the idea, which I advocated at the end of §4.2, that communicative success often doesn't require the the speaker and addressee to coordinate on a single, precise propositional content. After all: mindreading is hardly ever that precise, and yet we don't conclude that it is mysterious or impossible. But if communicating with someone is just a matter of intentionally triggering and guiding the same mindreading capacity they're using the rest of the time, it shouldn't be surprising if communication involves the same sort of slack as mindreading.

6. *Convention Acquisition*

Finally, I would like to consider the question of how we acquire linguistic conventions in the first place. Conventions are pairings of types of actions (in our case, types of communicative acts) with ways of performing them (in our case, types of utterances). The task of acquiring a new convention for how to ϕ is a matter of somehow coordinating on a new way, ψ , of ϕ ing. But every story about convention-acquisition that I know of involves the following steps.¹⁶

¹⁶ Cf. Schiffer 1972: §5.1; Lewis 1969; Skyrms 1996, 2010; Millikan 1998, 2005; Hume 1738: §3.2.2

- (i) Agents begin by ϕ ing in unconventional ways.
- (ii) Somehow, out of the chaos, a pattern of ϕ ing by ψ ing emerges, at first for arbitrary or idiosyncratic reasons.
- (iii) The pattern becomes increasingly standardized and self-reinforcing.
- (iv) Finally, a convention of ϕ ing by ψ ing has come into being.

The problem for Lepore and Stone, is that this abstract story entails that for any convention of ϕ ing by ψ ing to develop in the first place, it has to be possible to perform at least some rudimentary acts of ϕ ing in an unconventional way. It follows that no kind of action that has gradually come to be governed by conventions can be *essentially* conventional, and this includes communicative acts.

This reinforces a point that I made in §4.3. Lepore and Stone explicitly argue that convention-acquisition is an important step in many instances of communication. Although I disagree with their idea that we have anything to gain by holding that convention acquisition should be posited in place of all of the other inferential processes involved in interpretation (other than disambiguation), I do think that there are good reasons to think that linguistic conventions are constantly being renegotiated during conversations, and that being an effective communicator often involves acquiring new conventions on the fly. Indeed, much recent work in the philosophy of language has hammered home the degree to which linguistic conventions are constantly being renegotiated (Armstrong 2016; Barker 2002; Ludlow 2014; Plunkett and Sundell 2013).

But as I argued in §4.3, the process of acquiring a new linguistic convention is a mindreading task of just the kind that Lepore and Stone wish to expunge from their theory of communication. Figuring out that you are using an unfamiliar expression with the intention of engaging in a convention on which it means XYZ is no more straightforward than inferring, without the aid of any prior knowledge of the relevant conventions, that you mean XYZ in using the expression.

These points about convention acquisition are backed up by empirical evidence about language acquisition in children. For example, after surveying the literature on the role played by mindreading in word learning, Paul Bloom concludes that, provided we bracket the parallel task of acquiring concepts, “learning the meaning of a word just reduces to intentional inference; once we know how children divine the intentions of others, there is nothing left to explain” (2000). In a similar vein, Shevaun Lewis, Valentine Hacquard, Jeffrey Lidz, and their colleagues have developed a sophisticated and empirically supported model of the acquisition of the meanings of attitude verbs on which the model’s inputs are the child’s innate knowledge of syntax together with their ability to work out speaker meanings (Hacquard 2014; Lewis 2013; Lewis et al., MS).

These findings confirm that intention recognition is developmentally, and so explanatorily, prior to the use of linguistic conventions in

communication, and they give us strong reasons to conclude that the interpretation of communicative acts needn't be guided by knowledge of conventions.

Aside from pointing to confirmed cases where unadulterated intention recognition is involved in interpreting communicative acts, and thereby further bolstering my argument in §5, these considerations should take the wind out of what I have called Lepore and Stone's 'imaginative defense' of conventionalism (§4.2). If it is sometimes possible to successfully interpret communicative acts without the aid of convention—and it has to be, given that we acquire conventions in the first place—then the problem of inferring the properties of indirect communicative acts can't be as bad as Lepore and Stone's version of the meaning-intention problem might seem to suggest.

Conclusion

When presented with a property, ϕ , of a communicative act, α , that seems not to be conventional, Lepore and Stone reply in one of three ways:

THE CONVENTIONAL DEFENSE: ϕ is really conventional after all (§4.1).

THE IMAGINATIVE DEFENSE: α is not really a communicative act, but an invitation to engage in imaginative reflection with no right or wrong answers (§4.2).

THE NOVELTY DEFENSE: Hearers interpret α as having ϕ by positing a novel convention on the fly (§4.3).

I have tried to rebut these strategies for defending conventionalism in several ways.

First, I have argued that each of these defenses fails on its own terms. Although the examples that Lepore and Stone discuss in mounting the conventional defense may be susceptible to conventionalist treatment, other examples can't. Lepore and Stone's argument for the imaginative defense overgenerates, consigning numerous instances of perfectly good communication to the imaginative waste bin. Moreover, it rests on the dubious premises, rejected by Grice and at least some contemporary intentionalists, that communication has to be precise in order to succeed. And, in deploying the novelty defense, Lepore and Stone conclude, in effect, that many cases of communication involve an inferential process (convention acquisition) that isn't relevantly different from the kind of Gricean intention recognition that they wish to expunge from the theory of communication.

Second, I have argued that Lepore and Stone's three defensive strategies are in tension with one another. Since acquiring a new convention on the fly is essentially the same task as interpreting a communicative act without the aid of convention, and Lepore and Stone think that we regularly do the former, they are also committed to the possibility of doing the latter. But if this is so, then the imaginative defense can't work,

since it is designed to show that we can't interpret unconventional communicative acts.

I conclude that although Lepore and Stone's treatments of particular semantic and pragmatic phenomena are fascinating and possibly correct, their broader defense of conventionalism in the philosophy of language is not one about which we should be optimistic.

References

- Armstrong, J. 2016. "The problem of lexical innovation." *Linguistics and Philosophy*, 39 (2): 87–118.
- Bach, K. 1987. *Thought and Reference*. Oxford: Oxford University Press.
- Bach, K. 1992. "Intentions and demonstrations." *Analysis* 52: 140–146.
- Bach, K. and Harnish, R. M. 1979. *Linguistic Communication and Speech Acts*. Cambridge: MIT Press.
- Barker, C. 2002. "The dynamics of vagueness." *Linguistics and Philosophy* 25: 1–36.
- Bloom, P. 2000). *How Children Learn the Meanings of Words*. Cambridge: MIT Press.
- Brandom, R. 1994. *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.
- Buchanan, R. 2010. "A puzzle about meaning and communication." *Noûs*, 44 (2): 340–371.
- Burge, T. 1975. "On knowledge and convention." *Philosophical Review* 84 (2): 249–255.
- Butterfill, S. A. and Apperly, I. A. 2013. "How to construct a minimal theory of mind." *Mind and Language* 28 (5): 606–637.
- Carey, S. 2009. *The Origin of Concepts*. Oxford: Oxford University Press.
- Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Davidson, D. 1965. "Theories of meaning and learnable languages." In Bar-Hillel, Y. (ed.). *Proceedings of the International Congress for Logic, Methodology, and Philosophy of Science*. North-Holland: 3–17.
- Dunbar, R. I. M. 1998. "The social brain hypothesis." *Evolutionary Anthropology*, 6 (5): 178–190.
- Dunbar, R. I. M. 2003. "The social brain: Mind, language, and society in evolutionary perspective." *Annual Review of Anthropology* 32: 163–181.
- Gilbert, M. 1989. *On Social Facts*. London: Routledge.
- Grice, H. P. 1957. "Meaning." *The Philosophical Review* 66 (3): 377–388.
- Grice, P. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Hacquard, V. 2014. "Bootstrapping attitudes." In Snider, T. (ed.). *Proceedings of SALT* 24: 330–352.
- Harris, D. W. 2014. *Speech Act Theoretic Semantics*. PhD Dissertation. City University of New York Graduate Center.
- Hawthorne, J. 1990. "A note on 'languages and language.'" *Australasian Journal of Philosophy*, 68: 116–118.
- Hawthorne, J. 1993. "Meaning and evidence: A reply to Lewis." *Australasian Journal of Philosophy* 71: 206–211.

- Hawthorne, J. and Magidor, O. 2009. "Assertion, context, and epistemic accessibility." *Mind* 118 (470): 377–397.
- Hume, D. 1738. *A Treatise of Human Nature*. London.
- Kaplan, D. 1989. "Afterthoughts." In Almog, J., Perry, J., and Wettstein, H. (eds.). *Themes from Kaplan*. Oxford: Oxford University Press: 565–614.
- King, J. C. 2013. "Supplementives, the coordination account, and conflicting intentions." *Philosophical Perspectives* 27 (1): 288–311.
- King, J. C. 2014. "Speaker intentions in context." *Noûs* 48 (2): 219–237.
- Kukla, R. and Lance, M. 2009. "*'Yo!' and 'Lo!': The Pragmatic Topography of the Space of Reasons.*" Cambridge: Harvard University Press.
- Lepore, E. and Stone, M. 2015. *Imagination and Convention*. Oxford: Oxford University Press.
- Lepore, E. and Stone, M. 2017. "Explicit indirection." In Fogal, D., Harris, D., and Moss, M. (ed.). *New Work on Speech Acts*. Oxford: Oxford University Press.
- Lewis, D. K. 1969. *Convention: A Philosophical Study*. Cambridge: Harvard University Press.
- Lewis, D. K. 1975. "Languages and language." *Minnesota Studies in the Philosophy of Science* VII: 3–35.
- Lewis, S. 2013. *Pragmatic Enrichment in Language Processing and Development*. PhD thesis. College Park: University of Maryland.
- Lewis, S., Hacquard, V., and Lidz, J. (MS). "Think' pragmatically: Children's interpretation of belief reports." *Unpublished Manuscript*.
- Ludlow, P. 2014. *Living Words*. Oxford: Oxford University Press.
- Luo, Y. 2011. "Do 10-month-old infants understand others' false beliefs?" *Cognition* 121: 289–298.
- Miller, S. 2001. *Social Action: A Teleological Account*. Cambridge: Cambridge University Press.
- Millikan, R. G. 1998. "Language conventions made simple." *The Journal of Philosophy* 95 (4): 161–180.
- Millikan, R. G. 2005. *Language: A Biological Model*. Oxford: Oxford University Press.
- Neale, S. 2004. "This, that, and the other." In Bezuidenhout, A. and Rieger, M. (eds.). *Descriptions and Beyond*. Oxford: Oxford University Press: 68–182.
- Neale, S. 2016. "Pragmatism and binding." In Szabó, Z. G. (ed.). *Semantics versus Pragmatics*. Oxford: Oxford University Press: 165–285.
- Neale, S. 2016. "Implicit and aponic reference." In Ostertag, G. (ed.). *Meanings and Other Things: Essays in Honor of Stephen Schiffer*. Oxford: Oxford University Press.
- Perry, J. 1997. "Indexicals and demonstratives." In Hale, B. and Wright, C. (eds.). *A Companion to the Philosophy of Language*. Oxford: Blackwell: 586–612.
- Pettit, D. 2002. "Why knowledge is unnecessary for linguistic understanding." *Mind* 111 (443): 519–550.
- Plunkett, D. and Sundell, T. 2013. "Disagreement and the semantics of normative and evaluative terms." *Philosophers' Imprint* 13 (23).
- Schiffer, S. 1972. *Meaning*. Oxford: Oxford University Press.
- Schiffer, S. 1992. "Belief ascription." *Journal of Philosophy* 89 (10): 492–521.

- Schiffer, S. 1993. "Actual-language relations." *Philosophical Perspectives* 7: 231–258.
- Schiffer, S. 1994. "A paradox of meaning." *Noûs* 28: 279–324.
- Scott-Phillips, T. 2014. *Speaking Our Minds: Why Human Communication is Different, and How Language Evolved to Make it Special*. Basingstoke: Palgrave Macmillan.
- Sellars, W. 1954. "Some reflections on language games." *Philosophy of Science*, 21 (3): 204–228.
- Skyrms, B. 1996. *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
- Skyrms, B. 2010. *Signals*. Oxford: Oxford University Press.
- Sommerville, J. A., Woodward, A. L., and Needham, A. 2005. "Action experience alters 3-month old infants' perception of others' actions." *Cognition* 96 (1): B1– B1.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

