

Socializing Pragmatics*

William B. Starr[†]

(Preprint from *Beyond Semantics and Pragmatics*, G. Preyer (ed.), 2018. New York: OUP.)

Abstract

Lepore & Stone (2014) focus on two theoretically useful notions of *meaning*: conventional meaning and speaker meaning. For Lepore & Stone (2014: Ch.14), the former consists of our mutual expectations about how language is used — *conventions* — to make ideas public. The latter consists in ideas that are made public in virtue of the speaker's basic intentions in speaking (Lepore & Stone 2014: Ch.13). This paper argues that there is a third, more basic notion of meaning I call *significance*. The significance of an utterance is not reducible to the content it makes mutual, because it is partly based on the private commitments speakers have when they make utterances and the private commitments hearers form on the basis of utterances. More specifically, significance is the private speaker commitments and hearer effects which explain why utterances of a given type are reproduced in a population of agents (Millikan 2005). This leads to an approach that differs from Lepore & Stone (2014) in the treatment of non-conventional interpretive effects, speech acts and deception.

1 Introduction

Lepore & Stone (2014) focus on two theoretically useful notions of *meaning*. There are interpretive effects generated by linguistic conventions, and there are those generated by the intentions directly supporting the application of those conventions on particular occasions of use. For example, a speaker may refer to a person with the name *Janis*, and the hearer may recognize this because there is a shared expectation between speaker and hearer about what person *Janis* refers to. This is **conventional linguistic meaning**. By contrast, a speaker may refer to a newly discovered stray dog with the name *Luna*, and the hearer may recognize this, because there is a shared expectation that speakers use novel names with a contextually salient referent in mind and other contextual clues make it clear that the speaker intends the stray dog to be this referent. This is an interpretive effect generated by the *intentions* supporting linguistic conventions, and is a species of what Grice called **speaker meaning**.

One ambition of Lepore & Stone (2014) is to significantly simplify the Gricean account of speaker meaning, and to expand the empirical reach of explanations which appeal to linguistic conventions. This point of contest between Lepore & Stone (2014)

*This paper stemmed from the joint influence of my ongoing collaboration with Sarah Murray on speech acts (Murray & Starr to-appear-b, to-appear-a), and my formative experience as a Ph.D. student with Ernie Lepore and Matthew Stone. I am lucky to have had all three as interlocutors on these issues, and their impact on my thinking is too extensive to document in detail.

[†]Cornell University, Department of Philosophy; will.starr@cornell.edu

and those that remain closer to Grice's vision, like Neo-Griceans and Relevance Theorists, will no doubt attract much critical attention. However, another ambition of Lepore & Stone (2014) is to place some interpretive effects like metaphor and imagery in a distinct third class, also contrary to Neo-Griceans who treat it as an implicature and Relevance Theorists who treat it as the non-conventional, context-specific loose-use of words that is a routine part of interpreting all utterances (Carston 2002: §5.3). Lepore & Stone (2014) explain these effects in terms of **imaginative engagement**, which are produced by the open-ended psychological processes that hearer's creatively recruit when prompted by the speaker's utterance. Lepore & Stone (2014) argue that metaphor and imagery do not involve 'metaphorical meanings' because imaginative engagement does not fit the mold of conventional meaning, or speaker meaning.

This paper argues that besides conventional meaning and speaker meaning, there is a third theoretically important notion of meaning which I call **significance**. The significance of an utterance is not reducible to the content it makes mutual, because it is partly based on the private commitments speakers have when they make utterances and the private commitments hearers form on the basis of utterances. More specifically, significance is the private speaker commitments and hearer effects which explain why utterances of a given type are reproduced in a population of agents (Millikan 2005). Indirectly, this challenges the second ambition of Lepore & Stone (2014). I will not contest their particular account of metaphor. Instead, I will challenge the assumption that any interpretive effect which does not count as conventional meaning or speaker meaning, does not involve a kind of meaning. I will argue that it often involves significance, and that significance is crucially involved in generating speaker meaning and conventional meaning.

My argument is foreshadowed by considering a question not discussed by Lepore & Stone (2014) or Lewis (1979): why do the agents in a given population engage in a particular practice of conversational scorekeeping in the first place? Consider the practice of uttering a declarative sentence to update the informational score of a conversation. Suppose it makes the conversationalists mutually aware that this information is being assumed in their conversation. Why is this mutual awareness valuable to those conversationalists? If the conversationalists are going to do more than coordinate their *conversational actions*, this mutual awareness is only instrumentally valuable. It is valuable insofar as it brings about beliefs and actions that influence the world in a beneficial way. So what makes a given practice of conversational scorekeeping *significant* is its connection to the actual commitments and subsequent actions of the agents involved. Or to put the point less theoretically: language is valuable not for mere conversational coordination, but for coordinating on the problems that face us in our actual lives. Without an answer to this more basic question, we have no way of explaining why certain linguistic practices are self-sustaining and how they evolve to meet the practical needs of a linguistic population.

By focusing on the way language effects the conversational record, Lepore & Stone (2014) are able to address two very important issues. First, it addresses the issue of deception and distrust that plague previous accounts of speaker meaning and convention like Grice (1957) and Lewis (1969). Since Grice and Lewis focus on the actual effects that a signal has on hearers, one must address uses of that signal that produce deceptive effects (see also Lepore & Stone 2015: §20.3.3). Second, Lepore & Stone (2014) are able to give a clearer account of which interpretive effects count as semantic, by saying that only effects on the conversational record are semantic. I will conclude the paper

by discussing how these two issues are addressed in Murray & Starr (to-appear-a). I will begin by considering a fictional signaling system, and detailing the explanatory role of conventional meaning and significance.

2 From Signaling to Meaning: two paths

I will introduce a hybrid of Wittgenstein’s (1953) ‘builders’ language game, Lewis’ (1969) signaling games and Tetris. It will first be used to illustrate the concepts of convention and linguistic meaning articulated by Lewis (1969) and Lepore & Stone (2014). With some sophistications, it will also be used to challenge those accounts §§4.

Agents inhabit a world that naturally furnishes blocks of two precise shapes. Only



Figure 1: Two Basic Shapes of Blocks

some agents — the *builders* — can stack the blocks together into habitable dwellings. Their task is complicated by the fact that only U and T-shaped blocks fit together. The *collectors* have the exclusive ability to find and transport blocks of any shape (no builders are collectors). All agents need homes, so coordination is necessary. As Lewis (1969: Ch.4) spells out, the goal of the coordination is not particular acts — issuing this particular signal now, bringing this particular block now — but of *contingency plans*. A contingency plan informs an agent what to do in each kind of circumstance relevant to their interactions with other agents. The contingency plans for builders and collectors are spelled out as below, and to achieve coordination the either the builders must adopt B_{12} and the collectors C_{12} or the builders must adopt B_{21} and the collectors C_{21} .

Blocks Coordination Problem

- *Builders’ Contingency Plans*
 - B_{12} :
 - When a builder needs a U-shaped block they produce a signal σ_1
 - When a builder needs a T-shaped block they produce a signal σ_2
 - B_{21} :
 - When a builder needs a U-shaped block they produce a signal σ_2
 - When a builder needs a T-shaped block they produce a signal σ_1
- *Collectors’ Contingency Plans*
 - C_{12} :
 - When a builder produces σ_1 , collector brings a U-shaped block
 - When a builder produces σ_2 , collector brings a T-shaped block
 - C_{21} :
 - When a builder produces σ_1 , collector brings a T-shaped block
 - When a builder produces σ_2 , collector brings a U-shaped block

Each combination of these strategies can be assigned a utility for each agent. A payoff matrix, like that in Figure 2, displays this information — where (n, m) means that the result of combining the two strategies provides utility n for the builder and utility m for the collector. While the selection of numbers is somewhat arbitrary, this particular

	C_{12}	C_{21}
B_{12}	(1,1)	(-9,-9)
B_{21}	(-9,-9)	(1,1)

Figure 2: Payoff Matrix for Basic Blocks Game

assignment of utilities indicates that the result of failed coordination is far worse than successful coordination is good.

Suppose now that the builders and collectors have somehow solved their coordination problem, routinely adopting the policies of B_{12} and C_{12} . What follows about the conventional meaning of σ_1 and σ_2 ? Lewis (1969: 143) proposes the following:

Lewis (1969) Theory

1. The builders and collectors acting in accord with B_{12} and C_{12} counts as a convention because it is an arbitrary, recurrent solution to a coordination problem for all members of that population.
2. The meaning of σ_1 ($\llbracket \sigma_1 \rrbracket$) can be identified with the ways the world is when B_{12} is followed in accord with the convention, i.e. worlds where the builder needs a U-shaped block — similarly for σ_2 .

Lewis (1969:Ch.5) and Lewis (1975) goes on to characterize conventions of the kind outlined in 1 more generally as conventions of truthfulness and trust in a language, as used by a population. The convention of truthfulness governs the speaker’s production of a signal in the right situations, and the convention of trust governs the hearer’s response to the signal in the right way. Under this description, it is even more clear that the approach has explained how signals get their conventional meanings, at least when conventional meanings are taken to be truth-conditions.¹

In the context of Lewis (1969), the definition of meaning above is a huge success, insofar as it allows one to say how conventional meanings emerge from signaling:

I have been trying to demonstrate that an adequate account of signaling need not mention the meanings of signals — at least, not by name. But of course signals do have meanings.

This is because Lewis (1969) is attempting to give a naturalistic and non-circular definition of meaning. However, in the context of giving an empirically adequate account of natural language, it is less clear that Lewis (1969)/Lewis (1975) succeeds. Lepore & Stone (2014:Ch.14) detail two very important limitations on this front, and Lepore & Stone (2015) adds a third:

Empirical Limitations of Lewis (1969) Theory

1. The difference between imperatives and declaratives is not resolved at the level of conventional meanings or conventional contingency plans.

¹Of course, it’s not clear that σ_1 and σ_2 really have truth-conditions — perhaps they are better viewed as imperatives. This issue will come into focus shortly.

-
2. The difference between semantics and pragmatics, e.g. even if σ_1 is always used in situations where a builder is speaking, that does not count as part of its conventional meaning.
 3. Speakers can be deceptive and hearers can be distrustful — so there is no general regularity of most members of the population being truthful and trusting.

These limitations highlight three explanatory goals of a theory of meaning, and how that constrains the concepts we are willing to call meaning.

Explanatory Constraints

1. A theory of conventional meaning must make certain distinctions that are required by goals internal to semantic theory.
2. A theory of meaning in general must make certain distinctions that are required by general goals in the explanation of language use.
3. A theory of meaning must explain how conventions arise despite conflicting interests.

Each of these limitations and constraints deserve further discussion, including an articulation of how the account of conventional meaning in Lepore & Stone (2014) addresses them. The second limitation and constraint are really the central focus of Lepore & Stone (2014) explored thoroughly there. Accordingly, I will focus instead on the first and third.

Lewis (1969) was aware that his theory of conventional meaning did not adequately distinguish declarative and imperative meaning. While Lewis (1969:144-7) makes a relatively implausible attempt to distinguish them, Lewis (1975:172) offers an alternative account.² In natural languages, imperatives and declaratives are formally distinguished, and so there must actually be two different signals $\sigma_{1.1}$ and $\sigma_{1.2}$, the former declarative and the later imperative. Lewis (1975:172) proposes that the two kinds of signals are subject to different conventions of truthfulness and trust. While truthfulness of a declarative signal like $\sigma_{1.1}$ remains the same, the truthfulness of an imperative signal $\sigma_{1.2}$ is identified as taking the action required of the addressee by the convention. Accordingly, $[\sigma_{1.2}]$ will be worlds in which the collector brings a U-shaped block, while $[\sigma_{1.1}]$ will be worlds where the builder needs a U-shaped block. This approach, however, is not adequate when applied to languages with the compositional complexity of natural language. As discussed and argued in Murray & Starr (to-appear-b) and Starr (forthcoming), natural languages allow imperatives and declaratives to embed across connectives, making it impossible to partition sentences of the language into declaratives and imperatives. This means that conventions that assign contents differently depending on the type — like Lewis (1975:172) — or pragmatic treatments of the difference fall short.

This limitation is nicely overcome by a theory within the broad approach advocated by Lepore & Stone (2014). On such an approach, the conventional meaning of a sentence is characterized in terms of rules for updating the conversational record. Just

²See Zollman (2011) for discussion of the Lewis (1969) idea, and a new proposal based on the idea that imperatives should convey no information about the state of the world but perfect information about the response to an arbitrary receiver, and that declaratives are the reverse. I am ultimately sympathetic to this idea and think it consistent with the position outlined in Murray & Starr (to-appear-a). Indeed, its focus on informational dependencies also bears on the points below about mental imagery.

as the record can reflect both individuals mentioned, and standards of precision, it can reflect the difference between directive information contributed by imperatives, and representational information contributed by declaratives. Pragmatic theories of this kind already existed with Stalnaker (1978), Roberts (2004) and Portner (2004), but Murray & Starr (to-appear-b), Murray (2014) and Starr (forthcoming) show how to translate this into an empirically adequate semantic theory. So by treating semantic conventions as constituted by our rules for updating the conversational record, the broad approach outlined by Lepore & Stone (2014) makes for an empirically superior alternative to Lewis (1975:172). While their limited discussion of *There is a bug on your back* and *Swat the bug on your back* undersells this point, I think it is fair to say that they have made an important contribution here.

When it comes to deception, Lepore & Stone (2015) also invoke the framework developed in Lepore & Stone (2014:Ch.14). Lewis (1969) requires all speakers to use a signal truthfully and interpret it trustingly to sustain a meaningful language. And yet, we don't. If one adds to the builders/collectors setup that builders occasionally issue σ_1 not when they need a U-shaped block, but when they want collectors to busy themselves while the builder relaxes, then this kind of world will become part of $\llbracket \sigma_1 \rrbracket$. This highlights that Lewis (1969) has no way of saying which uses of σ_1 are constitutive of its meaning, and which are deviant. This is simply because he identifies the meaning of σ_1 with the way speakers actually use and hearers actually interpret it.

Lepore & Stone (2015) propose that this problem can be solved by drawing on Lepore & Stone (2014:Ch.14). The conventions of a language coordinate speaker and hearer, but they do so by providing rules for changing the conversational record. The record is distinct from the speaker and hearer's private beliefs. Even lies are used in accord with this convention: they put a proposition on the scoreboard. So by re-characterizing linguistic conventions in terms of the conversational record, Lepore & Stone (2015) narrow the scope of what it counts to use a signal in a way that sets aside the details of whether speaker is being sincere or whether the hearer trust them. This is good, because we live in a world full of insincere speakers and distrustful hearers. So too, we may assume, do the builders and collectors.

3 Why are Linguistic Conventions Self-Sustaining?

The approach of Lepore & Stone (2014) has fared well, so far. But there is a lingering issue that can be brought out by investigating the question of how conventions are self-sustaining on their approach. Lewis (1969) clearly recognized that it was crucial to explain how conventions were self-sustaining. His explanation was individual rationality: each agent has reason to conform if others do, and agents are rational. In short, everyone is rational so they know they want to act so as to end up in a (1,1) cell in Figure 2, and think that everyone else is going to do their part to end up in the B_{12} , C_{12} cell. This is why Lewis (1969) must require that all, or almost all, of the agents in the population conform. Lepore & Stone (2014) are able to maintain the universal quantification by narrowly characterizing language use so that even liars and the distrustful conform. But this raises the question of whether they can still adopt Lewis' (1969) explanation of how linguistic conventions are self-perpetuating.

Consider again the builders and collectors, and replace our earlier characterizations of their contingency plans with the new ones entailed by Lepore & Stone (2014):

Blocks Coordination w/Conversational Records

- *Builders' Contingency Plan*
 - B_{12}^* :
 - When a builder wants it to be on the **conversational record** that they need a U-shaped block, they produce a signal σ_1
 - When a builder wants it to be on the **conversational record** that they need a T-shaped block, they produce a signal σ_2
 - B_{21}^* :
 - When a builder wants it to be on the **conversational record** that they need a U-shaped block, they produce a signal σ_2
 - When a builder wants it to be on the **conversational record** that they need a T-shaped block, they produce a signal σ_1
- *Collectors' Contingency Plan*
 - C_{12}^* :
 - When a builder produces σ_1 , it goes on the **record** that the collector is to bring a U-shaped block
 - When a builder produces σ_2 , it goes on the **record** that the collector is to bring a T-shaped block
 - C_{21}^* :
 - When a builder produces σ_2 , it goes on the **record** that the collector is to bring a U-shaped block
 - When a builder produces σ_1 , it goes on the **record** that the collector is to bring a T-shaped block

Now suppose we offer the analogous payoff matrix: The real question is if we can

	C_{12}^*	C_{21}^*
B_{12}^*	(1,1)	(-9,-9)
B_{21}^*	(-9,-9)	(1,1)

Figure 3: Purported Payoff Matrix for Blocks Game w/Conversational Records

justify assigning a payoff of (1,1) for the new outcome described by the agents' contingency plans. I think it is relatively clear that we cannot, and this poses a serious dilemma for how Lepore & Stone (2015) propose to handle deception. Merely getting some information and expectations on the record does not guarantee that the builders and collectors will have a home to sleep in — so assigning (1,1) to the outcome of B_{12}^* and C_{12}^* is unjustified. It does not guarantee that the collector brings a U-shaped block, only that it's on the record that they are to do that. At most, keeping score with the conversational record is of instrumental value. It is valuable only insofar as it leads to changes in the actual world. But if that is right, then conventions cannot be sustained by the kind of practical rationality that justifies adopting certain contingency plans on the basis of their valuable consequences. This makes it hard to see how Lepore & Stone (2015) can both adopt the Lewis (1969) account of how linguistic conventions are self-sustaining and their approach to the problem of deception.

The problem facing Lepore & Stone (2015) generalizes to many related accounts that attempt to characterize language use primarily in terms of 'conversational score-keeping'. They all leave open the question of why that scorekeeping practice is self-sustaining in the first place. Leaving this question un-asked and un-answered leaves

open questions of fundamental interest, such as how the practice can be self-sustaining in the face of conflicting interests.

This issue is related to earlier discussion of force in an interesting way. While it was clear enough how to capture sentence force (declaratives, imperatives, interrogatives) within a scorekeeping approach, it is far less clear how to pull that off for what you might call utterance force. A single imperative *Do community service!* can be used as a command when issued by a judge in court, as advice when issued by a friend, or as an exhortation by a crowd of gathered community members already engaged in community service. While some theorists have proposed to treat this as semantic, and many more theorists have proposed to treat this in terms of speaker meaning, Murray & Starr (to-appear-a) argue that these accounts are philosophically and empirically inadequate. Semantic accounts fail to capture generalizations about the relationship between utterance force, context and linguistic form, while speaker meaning approaches face other challenges. For example, a sign reading *Take off your shoes and have a beer!* at the entryway to an apartment can naturally be read as commanding guests to take off their shoes while merely suggesting that they have a beer. However, speaker meaning only enriches complete sentences at the level of an utterance, and so cannot apply differentially to sub-sentential units. Murray & Starr (to-appear-a) argue that the phenomenon of utterance force can only be captured by talking directly about the appropriate individual (off-the-record) responses to the utterance. However, no account of this is possible with addressing head-on the issue of deception — §4 will present the solution to this problem proposed by Murray & Starr (to-appear-a).

Focusing on the self-sustaining nature of conventions raises another difficulty for Lepore & Stone (2014), one which they inherit from Lewis (1969). Some percentage of the population is always deviating from the signaling convention, either to deceive or to explore other solutions. In fact, this is good. This makes individual rationality a poor explanation of how conventions are sustained. This point can be illustrated with yet another twist on the builders/collectors game.

Suppose, like us, builders and collectors are capable of generating private mental imagery while communicating. Suppose further that they, like some of us, greatly enjoy imagining neon colored squares (the builders) and triangles (the collectors). There may then be two distinguishable conventions in competition for the use of σ_1 and σ_2 . B_{12}^+ and C_{12}^+ differ from B_{12} and C_{12} only in that the contingency plans are enriched with mental imagery.

Blocks & Imagery Coordination Problem

- *Builders' Contingency Plans*

- B_{12}^+ :
 - When a builder wants to **imagine a neon pink square** and needs a U-shaped block they produce a signal σ_1
 - When a builder wants to **imagine a neon green square** and needs a T-shaped block they produce a signal σ_2

- B_{21} :
 - When a builder needs a U-shaped block they produce a signal σ_2
 - When a builder needs a T-shaped block they produce a signal σ_1

- *Collectors' Contingency Plans*

- C_{12}^+ :
 - When a builder produces σ_1 , collector **imagines a neon blue triangle** and brings a U-shaped block

-
- When a builder produces σ_2 , collector **imagines a neon yellow triangle** and brings a T-shaped block
 - C_{21} : ◦ When a builder produces σ_1 , collector brings a T-shaped block
 - When a builder produces σ_2 , collector brings a U-shaped block

Given that the added imagery adds a boost in enjoyment, it makes sense to assign the combination of B_{12}^+ and C_{12}^+ a higher utility than B_{21} and C_{21} . Since B_{12}^+ and C_{12}^+

	C_{12}^+	C_{21}
B_{12}^+	(3,3)	(-9,-9)
B_{21}	(-9,-9)	(1,1)

Figure 4: Payoff Matrix for Imagery & Blocks Game

might emerge as options only in the presence of B_{21} and C_{21} , it is clear that it would be suboptimal to have a population without defectors. Without some segment of the population to show the rest the benefits of B_{21} and C_{21} , there would be no way for the population to reap the benefits of changing from the equilibrium of B_{21} and C_{21} . But a population without defectors is a population who's conventions are not sustained by individual rationality keeping every single individual acting in accord with a single regularity. This means that both Lepore & Stone (2014) and Lewis (1969) fail to explain a dynamic feature of signaling systems with this structure: why do systems like B_{12}^+ and C_{12}^+ emerge when competing with B_{21} and C_{21} , provided that there are some defectors?

This limitation ramifies in different ways for theories of conventional meaning like Lepore & Stone (2014) and Lewis (1969). Lewis (1969) would say that the mental imagery is part of the meaning of σ_1 in the signaling convention consisting of B_{12}^+ and C_{12}^+ — though only the builder's mental imagery. This is some small consolation, since the imagery does end up being relevant to the explanation of how signaling systems like this emerge. But it also seems deeply odd to build it into the *truth-conditions* of σ_1 since it is irrelevant to the action being chosen by the collectors. Lepore & Stone (2014) are in an even more awkward situation, since they would presumably exclude mental imagery from the way σ_1 influences the conversational score. Such a claim about the conversational score seems right. But, is that all that there is to say about the meaning of σ_1 ? The clear answer seems to be *no*: this does not explain why one convention governing σ_1 wins out in relevant circumstances, and surely this is a fact about σ_1 's meaning that needs explaining. After all, B_{12}^+ and C_{12}^+ wins out over B_{12} and C_{12} because on B_{12}^+ and C_{12}^+ allow σ_1 to do something that B_{12} and C_{12} do not.

This point can be made even more precise by a slight variant on the Imagery & Blocks game. Suppose instead of two competing conventions, there were two competing signals σ_1 and σ_1^+ . The former is governed by B_{12} and C_{12} while the latter by B_{12}^+ and C_{12}^+ . In such a setting, use of σ_1^+ will proliferate at the expense of σ_1 which will fall into obsolescence. Now the question to be explained is: what difference between σ_1^+ and σ_1 explains why this happened? It's hard to see how any answer could fail to mention the mental imagery associated with σ_1^+ and, in some sense, include this as part of σ_1^+ 's meaning.

The relationship of mental imagery and conventional meaning is a central topic of Lepore & Stone (2014). On a first pass, the Imagery & Blocks game beautifully illustrates why they want to exclude imagery from the conventional meaning of a sentence.

There is no common feature of the mental imagery which builders have and that collectors generate which counts as being ‘communicated’ when σ_1 is issued. Indeed, the builders could just as well have been clueless authors, experiencing no phenomenology themselves, but eliciting mesmeric patterns of neon shapes in the collectors — indeed it is no slight to the great authors of literature to assume that their skill is not in communicating a particular experience, but in generating diverse valuable experiences interactively generated by their audience. Further, the collectors need not have uniform mental imagery: they could all imagine rather different shades of neon blue — or different colors altogether — when builders issue σ_1 . And, yet, these effects of σ_1 on mental imagery are still a crucial feature of explaining why this signaling system works. Without referencing them, we would not know why B_{12}^+ and C_{12}^+ proliferates when competing with B_{21} and C_{21} . Similarly with the variant where σ_1^+ and σ_1 are in competition. If we want to explain why σ_1^+ proliferates at the expense of σ_1 , one must capture its capacity to generate mental imagery. This presents a clear challenge:

The Challenge of Mental Imagery The mental imagery that accompanies language use can be an ineliminable part of explaining why a particular signaling system is self-sustaining. If mental imagery are not part of the meaning of signals, one must explain why meanings and psychological effects like mental imagery inhabit distinct explanatory roles.

Nothing so far entails that this challenge cannot be met, but nothing Lepore & Stone (2014) say entails that it can be.

Let me take stock of the issues raised for Lepore & Stone (2014). By divorcing conventions of language from effects on agents’ individual psychological states, they undercut Lewis’ explanation of how the conventions of language are self-sustaining. Changes to the conversational record are not the intrinsically valuable outcomes that makes language self-sustaining. It is changes in the private beliefs and actions that are valuable. Makes it impossible to exclude deceptive uses from bearing on a convention by simply limiting the conventional regularity to the conversational record. It also makes it difficult to explain, in a principled way, why mental imagery should be excluded from the meaning of a signal. After all, that imagery can play a crucial role in making a particular practice of signal use, or a particular signal, self-sustaining. But if all of this is right, then we are bereft of some key benefits of the approach articulated by Lepore & Stone (2014). Is there another approach that secures those benefits? If so, where does it differ from Lepore & Stone (2014)?

4 Signaling, Meaning and Significance: the middle way

Millikan (2005: Ch.3) offers a very different way than Lewis (1969) of thinking about linguistic conventions and how they relate to theoretically important concepts of meaning. There are really two key ideas in Millikan (2005). One is that Millikan (2005: Ch.1) bases her theory of convention on the *reproduction* of signals and the speaker-hearer commitments they invoke. The other is that signals are thereby endowed with a *stabilizing function* (Millikan 2005: Ch.3):

Stabilizing Function The speaker and hearer commitments a signal generates which explains why that signal is reproduced (or reproducible) in a population.

Millikan’s approach to reproduction is complex and nuanced in ways that I cannot engage with here. The crucial feature of this account that will be important below is this. The speaker and hearer commitments that actually result from signal use do, sometimes, have to constitute coordination in the game-theoretic sense. However, there is no need to assume that speakers and hearers always or usually coordinate when they use a signal — just that future uses of the signal are explained by historically successful coordination. Crucially, a convention is self-sustaining not by reasoning about whether others conform, but just by observing cases where signals work to coordinate speaker and hearer, and repeating those behaviors.

Millikan’s (2005) approach clearly holds promise for solving the deception problem. It can distinguish all uses of a signal from the convention constituting ones: namely those successful coordinations of speaker hearer that keep the signal in use. However, Millikan (2005) does not present this account with the level of precision needed to speak directly to the issues raised in §3. Murray & Starr (to-appear-a) provides some of the needed tools to do so. It will take a bit of setup to present those tools. The first step is to present the account of communicative acts in Murray & Starr (to-appear-a).

Murray & Starr (to-appear-a) model communicative acts using **conversational states**. A conversational state is centered on a body of mutual assumptions in a conversational context A_C — this is roughly analogous to the conversational record of Lepore & Stone (2014). But, crucially, it also tracks the private assumptions of speaker A_S and hearer A_H — these are not generally known to other conversationalists. A communicative act is then modeled as a function from one conversational state $c = \langle A_S, A_C, A_H \rangle$ to another $c' = \langle A'_S, A'_C, A'_H \rangle$. A speech act type (or *utterance force*) has the stabilizing function of coordinating speaker and hearer. When that function is fulfilled c' is a Nash equilibrium — just as in Lewis (1969). An example helps illustrate this basic idea.

Consider a simple assertion like (1). It’s assertive force can be modeled using conversational states, as in (2).

- (1) The bug is on your back. (Asserted by S to H)
- (2) *Assertive Effect on* $c = \langle A_S, A_C, A_H \rangle$:
 - a. *Effect on* A_C : information that the bug is on H ’s back is made mutual.
 - b. *Effect on* A_S : S is privately committed to that mutual information.
 - c. *Effect on* A_H : S is privately committed to that mutual information.

Assigning the utterance of (1) an assertive force is just to say that the utterance’s stabilizing function is to produce the effects in (2a)-(2c). Of course, it can still count as an assertion even if it does not, in fact, produce those effects. When the utterance of (1) does achieve these effects, Millikan (2005:Ch.1) would describe it as part of a convention of assertion — after all the utterance and response are reproduced³ and the fact that this pattern of activity is reproduced is, in part, due to precedent: it’s a reproduction of an utterance that coordinated the speaker and hearer. It is important to clarify that this convention concerns the utterance of a sentence, rather than the sentence itself — i.e. the sentence being tokened by a speaker, at a time, world, etc. in the presence of a particular hearer. This raises the question, however, of what the

³It is important to note that Millikan (1984, 2005) offers a sophisticated theory of reproduction whereby the original does not completely determine the reproduction. This is crucial for language where *A and B* could be a reproduction of *A, B and C and D*, and inherit its function from *and, A and B*.

convention governing the *sentence* in (1) is. This is a crucial issue if we want to answer the question of how a convention endows a *sentence* with a meaning. In fact, this issue highlights one of the problematic features of many approaches to speech acts like Searle (1969): it does not adequately distinguish the linguistic contribution of the signal from the force of a particular utterance of that sentence (Murray & Starr to-appear-b). Making this distinction is essential for adequately explaining how linguistic and non-linguistic mechanisms interact to produce utterance force (Murray & Starr to-appear-a). Millikan (1984, 2005) does not speak directly to this issue, but Murray & Starr (to-appear-a) show how to unify her general approach with a dynamic semantics for sentential mood (declarative, imperative and interrogative).

Drawing on dynamic semantics, Murray & Starr (to-appear-a) propose that the linguistic conventions governing sentences are procedures for updating the mutual assumptions A_C . More concretely, the linguistic meaning of a declarative sentence like (1) is a function mapping one A_C to another which contains the information that the bug is on the hearer's back. This effect is *part of* many different functions the utterance of (1) could serve. For example, making this information mutual might not serve to transmit information from speaker to hearer, as in assertion, but instead to strengthen the social bond between them. Perhaps the speaker once truly told the hearer this and something hilarious happened. This second utterance could initiate a short line of pretense, where the humorous event is reenacted. On such a use, the sentence still serves to update the mutual assumptions — which facilitates the pretense — but utterances like this do not serve the function of transmitting information. Their function is social bonding, a commonly neglected function in research on human communication, but a central idea in work with primates (Tomasello 2008:Ch.6). So, on the approach developed in Murray & Starr (to-appear-a) the conventional meaning of a sentence constrains the force of an utterance by encoding a procedure for updating the mutual assumptions. But the particular force of an utterance concerns how that utterance type fits into the agents' social lives — more on what this amounts to shortly.

On this approach, linguistic conventions only concern what a sentence makes mutual. This characterizes them more narrowly than Millikan (1984, 2005), who characterizes the conventions directly in terms of the speaker and hearer commitments they generate.⁴ This compels her to posit polysemy via overlapping conventions to capture the diversity of uses to which a sentence type can be put. This is less attractive for the purposes of linguistic analysis than Murray & Starr's (to-appear-a) account for two reasons. First, it does not cohere well with linguistic typology, where each of these variants are realized in some language or other with an overt form. Second, it makes it impossible to formulate a general, non-linguistic theory of how mutual contributions relate to the private commitments of speakers and hearers. Murray & Starr (to-appear-a) propose just such a general theory.

Language provides a way of coordinating mutual assumptions, but each linguistic act also involves a more fundamental form of coordination via three general mechanisms that are not at all specific to language. The first two have been well-studied, and

⁴Millikan (1984) is deeply skeptical of higher-order intentions and mutual knowledge, as it's been used in the Neo-Gricean literature. However, we envision a characterization of the mutual assumptions in a way that avoids these concerns. Following Lewis (1969) instead of Stalnaker (1978), one can think of the mutual assumptions as what each agent is assuming for the purposes of their exchange, what each would be justified in assuming everyone is assuming, and so on. By talking in terms of justification, it is possible to grant that no agent actually need to maintain higher-order mental states while communicating.

some have even attempted to found a complete pragmatic theory on one or both of them.

Mechanisms of Coordination (Murray & Starr to-appear-a)

1. *Social Conventions* (Austin 1962)
 - An act by an agent with a given social position, with a particular audience, counts as a socially recognized event type.
2. *Communicative Intentions* (Grice 1957)
 - One agent X intends to influence another Y 's state of mind, and intends Y 's state of mind to change, in part, by recognizing that X intended to change Y 's state of mind in that way.
3. *Social Norms* (Bicchieri 2005)
 - Self-fulfilling expectations of what agents like us do in situations like these, enforced by informal sanctions (shame, disgust, exclusion, etc.)

However, neither social conventions nor communicative intentions are generally viable mechanisms for getting things done together. The reason is simple: our social interactions are not, under their most fundamental description, coordination games at all. They are 'mixed-motive' games, like the Prisoner's Dilemma. In the Prisoner's Dilemma, two criminals are separately given a choice between informing on the other and getting released, or staying quiet and serving a short sentence. If the two prisoners were allowed to exchange messages neither social conventions, nor communicative intentions would be effective means for enabling communication. Social norms, according to Bicchieri (2005), exist to solve problems like these. If others expect the prisoners to stay quiet, and will sanction defecting heavily enough to outweigh the costs of taking a light sentence, then this mixed-motive game is transformed into a coordination game. Similarly, if others expect the prisoners to be truthful in their message-exchanges, and will sanction lies heavily enough, then these two prisoners will at least be able to communicate.

Upon reflection, our everyday communication is no different than the Prisoner's Dilemma. It may be better for both of us to communicate than not, but I might want more information than you have time to articulate, and you may not be sure I'll do the same for you when I need it. And, obviously, you might have an interest in my being misinformed. In these cases, there is no Nash equilibrium that allows neither of us to end up a bit better off. So, if viewed in terms of narrow self-interest, and in isolation of our social organization, everyday communication seems as unsustainable as it was for the prisoners. One way to resolve this tension is to assume that our everyday communicative interactions are infused with social norms that transform our mixed-motive games into coordination games.

What exactly are social norms, and which particular ones are at play in human communication? Full answers are beyond the scope of this paper, but some simplified ones may make my proposal here more concrete. According to Bicchieri (2005:11): some social practice is a social norm just in case each agent A prefers to conform to the practice given that conditions (i) and (ii) obtain, and those conditions do obtain (Bicchieri 2005:11):⁵

⁵More precisely: a behavioral rule R is a social norm just in case almost everybody knows that R exists

-
- (i) *A* expects others to conform and
 - (ii) *A* either believes that others expect *A* to conform or that others prefer *A* to conform and will informally sanction non-conformity (shame, disgust, etc.).

As Bicchieri (2005:3) clarifies, this is a rational reconstruction of what a social norm is, and is consistent with a psychological implementation that is sub-personal, unconscious and economically approximates the concept defined by the rational reconstruction. Applying this idea to speech act types, assertions, and several other speech act types, involve a social norm that relates changes to the mutual assumptions to changes in private commitments. For example, the analysis of assertion from (2) can be understood as involving a norm which says that the speaker is expected to be privately committed to the information they made mutual, and that the hearer is expected to become privately committed to information made mutual by an authority. The relevant forms of sanction exploit humans' impressive social memory and intricate systems of reputation and authority (Scott-Phillips 2011, 2015). As long as these projections of authority and reputation somewhat reliably track a speaker's trustworthiness and competence, trusting their contributions and expecting them to be sincere will keep deception at bay, and facilitate communication. These systems of norms and reputation are an evolved cultural solution to the problem of living together, and each new generation inherits them when they learn how to live in a society.

I can now return to the issues of §3. Like Lepore & Stone (2014), deceptive uses of a sentence do not bear on the relevant linguistic convention: declaratives update the mutual assumptions appropriately even when they are used to lie. However, unlike Lepore & Stone (2014), there is a clear answer to how this linguistic convention is self-sustaining. It is embedded in behavioral regularities governed by social norms. These norms make it possible to coordinate despite our conflicting interests, and regulate the private commitments conversationalists form. These patterns of coordination are reproduced, and include the relevant linguistic practice of updating our mutual assumptions. This theory of language use that relates linguistic conventions to the actual psychological effects certain uses have, even though I am not identifying linguistic conventions with those particular effects. This also offers a principled reason for excluding mental imagery from the linguistic meaning or convention from the Blocks & Imagery Coordination Problem. While those psychological states are essential to explaining how utterances and their interpretations are reproduced, it is not plausible to include them in the distinctively linguistic effect of changing the mutual assumptions. In order to justify including the information that the builder is imagining a neon pink square in the mutual assumptions, something the builder and collector are doing together would have to depend on being able to take that information for granted.⁶ Without a far more nuanced signaling system, including perhaps signals about mental states, this is not possible.

and prefers to conform to *R* on the condition that (a) almost everybody believes that almost everybody conforms to *R* and either (b) almost everybody believes that almost everybody expects almost everybody to conform to *R* or (b') almost everybody believes that almost everybody expects almost everybody to conform to *R*, prefers them to conform to *R* and may sanction those that don't (Bicchieri 2005: 11).

⁶This raises a question: is it that the builder's signal doesn't provide the collector with any information at all about the builder's mental imagery, or just that it doesn't make this information *mutual*. The former idea is closely related to the central proposal of Skyrms (2010: Ch.3): σ provides information about some state of the world s only if $P(s | \sigma) \neq P(s)$. I leave this question for further research.

Both of these moves seem broadly consistent with the approach in Lepore & Stone (2014). But, the general framework presented above also presents a challenge to a key assumption in Lepore & Stone (2014): that all meaning is either speaker meaning or conventional meaning. Above, we saw a genuine sense of meaning that cannot be reduced to either: the stabilizing function of an utterance. In an homage to Welby (1896a,b) who's central focus seemed to be this sense of *meaning*, I call it **significance**:

Significance What certain kinds of utterances of σ do which explains why that kind of utterance of σ is reproduced. (I.e. the stabilizing function of an utterance.)

This is to be contrasted, on my view with:

Conventional Meaning What all utterances of σ do to A_C which explains why σ is reproduced. (I.e. the stabilizing function of a signal.)

While significance is clearly a concept of pragmatics, it should also be contrasted with speaker meaning. Some utterances, particularly novel ones, will inherit their significance from the speaker's communicative intentions, but it is perfectly possible for this to not be the case. As a purely theoretical example, a poem that generates some particular imagery or associations in a sizable chunk of the population may justify including that effect in the poem's significance even if the author did not intend it, and it is not part of the information which the poem makes mutual by linguistic conventions. Linda Lovelace's *Ordeal*, a much-discussed example in the feminist literature on pornography (e.g. Langton 1993: 321-2), may best be understood in this way. The book was intended to protest and undermine the pornography industry by providing details and images about the brutality and coercion women in the industry endured. However, it became a sensation among pornography consumers that were aroused by such depictions. It soon came to be sold at adult books stores and widely consumed as pornography. Perhaps it is right to say that its significance was pornographic, even though its conventional and speaker meanings were anything but.

Ordeal is not an isolated example in the feminist literature on language. There, one finds many phenomena which do not fit nicely into the theoretical concepts afforded by the Fregean and Gricean tradition. The particular implementation of utterance significance here suggests that the significant effects of an utterance are often mediated by social norms. As emphasized by Bicchieri (2005), many social norms are deeply oppressive. They transform a mixed-motive game into a coordination game by unjustly endowing a group with the power of sanctions that effectively erases the self-interests of others — foot-binding in China is just one such example. Social norms that impact our linguistic practices are not immune from this injustice. The inability of the socially oppressed to have their conversational contributions believed, or even engaged with, is a central issue discussed in work on illocutionary disablement (Langton 1993) and testimonial injustice (Fricker 2007). As proposed in Murray & Starr (to-appear-a), these phenomena are easier to understand in terms of social norms and utterance significance. Thinking of significance as drawing on social norms, and emerging from complex interactions in a population of agents amounts to a fundamental shift in the orientation of pragmatics. Socializing pragmatics will require different tools and methodologies, and far more substantive engagement with social philosophy, sociology and socio-linguistics.

It is no surprise, then, that feminist work on language suggests immediately promising frontiers to explore. Consider McConnell-Ginet (2012: 747):

“[C]hildren’s books, syntax texts, newspapers, and many other media include many more references to men and boys than to women and girls. Notice that no particular utterer or utterance need have ‘meant’ that male human beings are more important than female or even more interesting or less problematic to discuss nor does anyone have to embrace such beliefs explicitly. Indeed, many people who themselves contribute to these patterns might be dismayed to realize that they have done so.”

This kind of emergent social meaning fits well with the concept of utterance significance defined here. It is a process by which agents’ utterances replicate ideology not in virtue of the linguistic conventions, or speaker’s intentions, but by the background psychological processes by which we engage with language. Much more work is needed to deliver on this promise, but I hope to have made the case that this kind of pragmatic social meaning can be integrated within a precise formal approach to linguistic meaning. There is no reason to exclude it as Lepore & Stone (2014) do, and, in fact, I have argued that any adequate account of how language is self-sustaining must include it. Perhaps it will also answer the call of philosophers like Hornsby (2000) who have highlighted the limiting individualistic assumptions of the Gricean tradition.

5 Conclusion

Any account of conventional linguistic meaning must explain how particular patterns of use and interpretation are self-sustaining. There are two main tasks in doing so: explain how those practices are valuable to a given population, *and* specify a mechanism by which that value suffices to perpetuate the practice. Lewis and Grice appealed to high-order rationality for the second task, but that proposal is seeming increasingly inadequate. As anticipated by Millikan (1984), there are a wealth of tools from evolutionary biology that can provide more psychologically realistic explanations (Skrms 2010). Lewis and Grice address the first task by tying use and interpretation to the actual private psychological states of language users. Our linguistic tools are meaningful because they are intertwined with our actual beliefs and desires, intentions and goals. Lepore & Stone (2014) departed from this tradition, instead adopting an account where language is meaningful insofar as it constitutes a rule-governed practice of conversational scorekeeping. However, this does not explain why speaking a language is not a purely recreational game, nor how it is self-sustaining. Filling this gap in Lepore & Stone (2014) requires investigating the *significance* of language. I have said here how that can be made precise and tried to suggest that it opens a new, more social, perspective on pragmatic theory. It may yet provide a way of articulating insights about non-literal language and speech acts that do not fit into the orthodox categories of linguistic meaning or speaker meaning.

References

- AUSTIN, JL (1962). *How to Do Things With Words*. Oxford: Oxford University Press.
- BICCHIERI, C (2005). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.

-
- CARSTON, R (2002). *Thoughts and Utterances: the pragmatics of explicit communication*. Malden, MA: Wiley-Blackwell.
- FRICKER, M (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. New York: Oxford University Press.
- GRICE, HP (1957). 'Meaning.' *Philosophical Review*, **66**(3): 377–388. URL <http://www.jstor.org/stable/2182440>.
- HORNSBY, J (2000). 'Feminism in Philosophy of Language: Communicative Speech Acts.' In M FRICKER & J HORNSBY (eds.), *Cambridge Companion to Feminism in Philosophy*, 87–106. Cambridge: Cambridge University Press.
- LANGTON, R (1993). 'Speech Acts and Unspeakable Acts.' *Philosophy & Public Affairs*, **22**(4): 293–330. URL <http://www.jstor.org/stable/2265469>.
- LEPORE, E & STONE, M (2014). *Imagination and Convention: distinguishing grammar and inference in language*. New York: Oxford University Press.
- LEPORE, E & STONE, M (2015). 'David Lewis on Convention.' In B LOEWER & J SCHAFER (eds.), *A Companion to David Lewis*, 313–327. Oxford: Wiley. URL <http://dx.doi.org/10.1002/9781118398593.ch20>.
- LEWIS, DK (1969). *Convention: A Philosophical Study*. Cambridge, Massachusetts: Harvard University Press.
- LEWIS, DK (1975). 'Languages and Language.' In K GUNDERSON (ed.), *Language, Mind, and Knowledge. Minnesota Studies in the Philosophy of Science, Vol. 7*, 3–35. Minneapolis, Minnesota: University of Minnesota Press. Page references to reprint Lewis (1983).
- LEWIS, DK (1979). 'Scorekeeping in a Language Game.' *Journal of Philosophical Logic*, **8**(1): 339–359. URL <http://www.jstor.org/stable/30227173>.
- LEWIS, DK (1983). 'Languages and Language.' In *Philosophical Papers*, vol. 1, chap. 11, 163–188. Minneapolis, Minnesota: Oxford University Press.
- MCCONNELL-GINET, S (2012). 'Language, Gender and Sexuality.' In DG FARA & G RUSSELL (eds.), *Routledge Companion to the Philosophy of Language*, 741–752. New York: Routledge.
- MILLIKAN, RG (1984). *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- MILLIKAN, RG (2005). *Language: A Biological Model*. New York: Oxford University Press.
- MURRAY, SE (2014). 'Varieties of Update.' *Semantics and Pragmatics*, **7**(2): 1–53. URL <http://dx.doi.org/10.3765/sp.7.2>.
- MURRAY, SE & STARR, WB (to-appear-a). 'Force and Conversational States.' In D FOGAL, D HARRIS & M MOSS (eds.), *New Work on Speech Acts*. New York: Oxford University Press.

-
- MURRAY, SE & STARR, WB (to-appear-b). 'The Structure of Communicative Acts.' *Linguistics & Philosophy*.
- PORTNER, P (2004). 'The Semantics of Imperatives within a Theory of Clause Types.' In RB YOUNG (ed.), *Proceedings from Semantics and Linguistic Theory 14*, 235-252. Ithaca, NY: CLC Publications. URL <http://semanticsarchive.net/Archive/mJ1ZGQ4N/PortnerSALT04.pdf>.
- ROBERTS, C (2004). 'Context in Dynamic Interpretation.' In L HORN & G WARD (eds.), *The Handbook of Pragmatics*, 197-220. Malden, MA: Blackwell. URL <http://www.ling.ohio-state.edu/~croberts/context.pdf>.
- SCOTT-PHILLIPS, T (2015). *Speaking Our Minds: why human communication is different, and how language evolved to make it special*. New York: Palgrave Macmillan.
- SCOTT-PHILLIPS, TC (2011). 'Evolutionarily Stable Communication and Pragmatics.' In A BENZ, C EBERT, G JÄGER & R VAN ROOIJ (eds.), *Language, Games, and Evolution*, vol. 6207 of *Lecture Notes in Computer Science*, 117-133. Springer Berlin Heidelberg. URL http://dx.doi.org/10.1007/978-3-642-18006-4_6.
- SEARLE, JR (1969). *Speech Acts*. Cambridge, England: Cambridge University Press.
- SKYRMS, B (2010). *Signals: Evolution, Learning, and Information*. New York: Oxford University Press.
- STALNAKER, RC (1978). 'Assertion.' In P COLE (ed.), *Syntax and Semantics 9: Pragmatics*, 315-332. New York: Academic Press. References to Stalnaker 1999.
- STALNAKER, RC (1999). *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press.
- STARR, WB (forthcoming). 'A Preference Semantics for Imperatives.' *Semantics & Pragmatics*.
- TOMASELLO, M (2008). *Origins of Human Communication*. The Jean Nicod Lectures. Cambridge, MA: MIT Press.
- WELBY, V (1896a). 'Sense, Meaning and Interpretation I.' *Mind*, 5(17): 24-37. URL <http://www.jstor.org/stable/2247630>.
- WELBY, V (1896b). 'Sense, Meaning and Interpretation II.' *Mind*, 5(18): pp. 186-202. URL <http://www.jstor.org/stable/2247581>.
- WITTGENSTEIN, L (1953). *Philosophical Investigations*. Oxford: Basil Blackwell.
- ZOLLMAN, KJS (2011). 'Separating Directives and Assertions Using Simple Signaling Games.' *Journal of Philosophy*, 108(3): 158-169.