CROATIAN
JOURNAL
OF PHILOSOPHY

Vol. XVI · No. 47 · 2016

Articles

Book Reviews

Introduction

The papers in this volume were presented at the conference Philosophy of Linguistics and Language *held at the IUC in Dubrovnik, September 2015. A day of the conference was dedicated to the discussion of* Ernie Lepore and Matthew Stone's *book* Convention and Imagination *thus the first three contributions directly address different points in the mentioned book.*

First Ernie Lepore and Matthew Stone *give their "Précis of* Imagination and Convention*" where they give an overview of the arguments presented in their book and explain how ideas from the book continue to inform their ongoing work. One theme that they stress is the challenge of fully accounting for the linguistic rules that guide interpretation. They do this by attending to principles of discourse coherence and the many aspects of meaning that are linguistically encoded but are not truth-conditional in nature. Thus they argue that they get a much more constrained picture of context sensitivity in language than philosophers have typically assumed. The other theme is the heterogeneous nature of interpretive processes where they propose that the connotations of an utterance are often best explained in terms of the hearer's experiential engagement with language, without appeal to propositional content that the speaker somehow signals either semantically or pragmatically.*

In their article "Against Lepore and Stone's Sceptic Account of Metaphorical Meaning" Esther Romero and Belén Soria *discuss and critically assess Lepore and Stone's account of metaphor. They claim that this account is based on three of Davidson's proposals: (i) the rejection of metaphorical meanings; (ii) the rejection of metaphors as conveying metaphorical propositional contents; and (iii) the defence of analogy as the key mechanism for understanding metaphors. Lepore and Stone defend these proposals because of the non-sceptic strategy on metaphorical meanings while Romero and Soria show not only how their non-sceptic account of metaphorical meaning as a variety of ad hoc concept eliminates difficulties but also how it can solve related difficulties in Lepore and Stone's approach.*

Daniel Harris *in his paper "Intentionalism versus The New Conventionalism" asks the question: Are the properties of communicative acts grounded in the intentions with which they are performed, or in the conventions that govern them? Ernie Lepore and Matthew Stone argue that much more of communication is conventional than we thought, and that the rest is not really communication, but merely the initiation of open-ended imaginative thought. Harris argues that although Lepore*

and Stone may be right about many of the specific cases they discuss, conventionalist conclusions do not necessary follow.

Marilynn Johnson's *paper "Cooperation With Multiple Audiences" is not the direct discussion of Lepore and Stone's book. She critically approaches Steven Pinker's proposal of a game-theoretic framework to help explain the use of veiled speech in contexts where the ultimate aims of the speaker and hearer may diverge—such as a case of bribing a police officer to get out of a ticket. Pinker's proposal is a seeming failure in H. P. Grice's influential theory of meaning to recognize that speakers and hearers are not always cooperating. Johnson argues that Pinker mischaracterizes Grice's views on cooperation and then argues that the cases Pinker presents are best treated by recognizing that in each instance the utterance is formulated with two intentions towards two different audiences. Johnson then goes on to detail a resulting revision to Pinker's game-theoretic framework that reflects this proposal.*

Jessica Keiser *in her paper "Coordinating with Language" looks into the idea that linguistic meaning is determined by use pointing to the fact that this claim marks the point where metasemantic inquiry begins rather than where it ends. It sets an agenda for the metasemantic project: to distinguish, in a principled and explanatory way, those uses that determine linguistic meaning from those that do not. The prevailing view (along with its various refinements), which privileges assertion, suffers from being at once overly liberal and overly idealized. By parsing the most prominent aims we use language to achieve, noting their relations of dependence and the specific type of uses they involve, she arrives at a novel metasemantic account: facts of linguistic meaning are determined by locutionary action.*

Marco Ruffino's *contribution "Superficially and Deeply Contingent A Priori Truths" reviews some standard approaches to the cases of contingent a priori truths that emerge from Kripke's (1980) discussion of proper names and Kaplan's (1989) theory of indexicals. In particular, he discusses Evans' (1979) distinction between superficially and deeply contingent truths. He raises doubts about Evans' strategy in general, and also about the roots and meaningfulness of the distinction.*

DUNJA JUTRONIĆ

Croatian Journal of Philosophy
Vol. XVI, No. 47, 2016

Précis of Imagination and Convention

ERNIE LEPORE
Department of Philosophy, Center for Cognitive Science, Rutgers, the State University of New Jersey, USA

MATTHEW STONE
Department of Computer Science, Center for Cognitive Science, Rutgers, the State University of New Jersey, USA

We give an overview of the arguments of our book Imagination and Convention*, and explain how ideas from the book continue to inform our ongoing work. One theme is the challenge of fully accounting for the linguistic rules that guide interpretation. By attending to principles of discourse coherence and the many aspects of meaning that are linguistically encoded but are not truth conditional in nature, we get a much more constrained picture of context sensitivity in language than philosophers have typically assumed. Another theme is the heterogeneous nature of interpretive processes, as illustrated by the distinctive interpretive profile of metaphorical and poetic language. Such effects remind us that the suggestions and connotations of an utterance are often best explained in terms of the hearer's experiential engagement with language, without appeal to propositional content that the speaker somehow signals either semantically or pragmatically.*

Keywords: Semantics, pragmatics, discourse coherence, context dependence, metaphor.

Imagination and Convention is a response to recent work in the cognitive science of language—work which has deepened philosophers' understanding both of the rules of language and of the processes of interpretation by exploring in new detail the fine-grained distinctions that characterize the interpretation of utterances in context.

One tradition we engage with is that of formal semantics. While this research once focused on the truth-conditional meanings delivered by sentence-level grammar, in the tradition inaugurated by Montague (1974), recent work is much broader in scope. There are now a variety of formal theories of presupposition (e.g., van der Sandt 1992, Beaver 2001), expressive meaning (e.g., Potts 2005), projective and not-at-issue

meaning (e.g., Tonhauser et al. 2013), and the interpretive links that connect multi-sentence discourse (e.g., Asher and Lascarides 2003). If such developments in formal semantics pan out, philosophers will need new conceptual tools to get clear on the relationship between grammar, meaning, interpretation and communication. The received constructs that philosophers have used to frame intuitions about meaning, like Grice's 'what is said', don't capture what language encodes or how grammar shapes interpretation.

The second tradition we engage with is a psychological one, which explains how language users make sense of utterances and their speakers. A common suggestion is that interpretations are often constructed creatively (Atlas 1989, 2005), for example, by taking words to signal new 'ad hoc' concepts (Carston 2002), by understanding phrases to be implicitly 'enriched' to more specific interpretations (Recanati 2004), and by loosening and transferring literal interpretations in light of inferences that matter in context (Sperber and Wilson 1986, 2008). Other researchers attribute interpretive effects to our empirical understanding of others' choices (e.g., Pinker, Novak and Lee 2008), or to open-ended processes of imaginative engagement (e.g., Camp 2008). Again, these diverse models require us to refine philosophers' received constructs for characterizing pragmatic inference, notably, of course, Grice's notion of 'conversational implicature'.

Synthesizing the perspectives of current research in semantics and pragmatics brings further challenges. Pragmatic theories have not yet come to grips with the heterogeneous nature of linguistic meaning as hypothesized in current formal semantics (see Simons et al. 2016 for some of the challenges involved). Conversely, the interpretive variability exposed by pragmatic research is often understood to undermine the assumptions and framework of formal semantics (Atlas 2005, Travis 1997).

Imagination and Convention offers our take on this new intellectual landscape. In this précis, we give a brief overview of the philosophical positions that make our view distinctive and highlight some of the research directions that our new view affords. First, in §1, we draw some lessons about the linguistic rules that guide interpretation. Our contention is that context sensitivity is much more closely governed by linguistic rules than is often appreciated. However, these rules appeal to more diverse principles than figure in traditional conceptions of semantics. Most importantly, we argue that the rules are sensitive to principles of *discourse coherence*, which we think of as linguistic conventions that connect and structure sequences of linguistic expressions within, and across, sentences, and encode implicit inferential relationships among their contents. In particular, the rules that link context-sensitive expressions to their semantic values can only be stated in terms of the overall organization of coherent discourse.

Moreover, we believe that a broad characterization of linguistic structure and meaning is crucial for philosophers to correctly diagnose

the interplay between semantics and pragmatics. A particularly fruitful but neglected case is intonation, which linguists model as a level of grammar that helps to signal the information structure of sentences in context, via the abstract meanings it encodes. As an illustration of the untapped implications of these principles, we close §1 with a brief survey of some developments since *Imagination and Convention* that we have pursued partly in collaboration with Una Stojnic (Lepore and Stone 2017a, Stojnic, Stone and Lepore 2013, Stojnic, Stone and Lepore 2017), on semantic models of discourse coherence and context-dependent meaning.

Next, in §2, we draw some lessons about the interpretive mechanisms involved in appreciating the points that speakers have in using utterances. On our view, the insights that we gain from an utterance often come from thinking about it in specific, creative ways. We briefly sketch our account of metaphor as a quintessential example of such imaginative engagement. We see metaphor as a distinctive way of thinking of one thing as another—one whose effects can differ from person to person and from occasion to occasion, and cannot be fully characterized just in terms of propositional content. This broadly Davidsonian view, which we elaborated already in Lepore and Stone (2010), was in many ways the impetus for our critical take on implicature in *Imagination and Convention*. Interpreting a metaphorical utterance, on our view, requires the hearer to engage in this process of metaphorical thinking, and to appreciate the insights this thinking engenders. In some cases, on our account, listeners can perhaps gain a deeper understanding into a speaker's intentions in using a metaphor, as a side effect of their own metaphorical thinking. Note that this explanation flips the direction of explanation often suggested in pragmatic accounts of metaphor, such as Searle's (1979) Gricean account, or Sperber and Wilson's (2008) in terms of Relevance Theory, which attempt to show how general reasoning about a speaker might prompt a listener to pursue associated or enriched interpretations which theorists might characterize retrospectively as metaphorical. Poetry is another case that we have begun to explore (Lepore and Stone 2016) but which did not make it into *Imagination and Convention*. We close this section with a brief overview of our approach to poetic interpretation: on our view, it involves exploring the articulation of a linguistic expression for added insight into its meaning.

We close in §3 with some reflections on the limits of knowledge of language. Our view invites theorists to capture a wide range of conventional information within a broad overarching framework for linguistic meaning: this includes not only the truth conditional content that is at issue in the use of a sentence, but also content that is encoded yet not at issue, for example, because it is marked as presupposed background, or because it is attached to a form as a matter of conventional implicature. However, a key part of our view is that the insights that the imagination prompts don't have the status of linguistic meanings. Again, the consequences

of this suggestion are largely unexplored. For example, in Lepore and Stone (2017b), we explore the idea that such insights are nevertheless an integral part of speakers' ear for the tonality of language—following up the influential suggestion of Frege that words can carry tone that does not contribute to the thoughts that sentences express.

1. *The Interpretive Effects of Linguistic Rules*

It often seems, intuitively, as though the interpretation of utterances is much stronger than the linguistic meanings of the expressions we use. In the book, we give (1–3) as illustrations of these effects.

(1) Can I have the French toast please?
(2) Oil prices doubled and demand for consumer goods plunged.
(3) Well, it looked red.

Example (1) is a question about ability that's easily interpreted as making a request. Example (2) offers a pair of event descriptions; the speaker's point seems to be that the events happened in succession and were perhaps even causally related. Example (3) describes the appearance of an object—but we take the speaker also to suggest that things were not as they appeared.

On our view, these interpretations will normally count as cases of successful communication. The speaker has a particular interpretation in mind, and the listener succeeds in recovering that interpretation. The question is what knowledge enables the coordination between speaker and listener: it can be hard to see intuitively how closely, if at all, the interpretations we derive in these cases are associated with the linguistic forms the speaker uses. Part of the traditional appeal of pragmatic explanations, we think, is that they promise to explain how linguistic meaning gets enriched or amplified here, without reference to unsuspected encoded meaning.

In *Imagination and Convention,* however, we argue that these interpretations *are* encoded—despite appearances. To do so, we argue for a richer conception both of linguistic form and of linguistic meaning than philosophers have traditionally entertained.

To start, we suggest—following work of Asher and Lascarides (2003), Grosz and Sidner (1986), Hobbs (1979), Kehler (2001), and Webber et al. (2003)—that linguistic form crucially includes a level of discourse structure that gives an organization to linguistic expressions that can extend beyond an individual clause. In particular, discourse structure groups sentences together hierarchically into segments that are interpreted as a coherent whole, much as sentence syntax unites constituents together by structural and interpretive connections (Grosz and Sidner 1986; Webber et al. 2003). Elements in discourse structure play specific roles in underwriting hearers' interpretive inference. These interpretive connections are known as *coherence relations,* and they have a range of interpretive effects. For one thing, coherence rela-

tions mark the commitments that arise when speakers use utterances indirectly to signal reactions to previous discourse; this means they make many kinds of alleged implicatures explicit. For another thing, coherence relations put particular discourse entities at the center of attention, making them the most prominent values for resolving context-dependent expressions. Because coherence relations give qualitatively different structures to discourse and dictate the formal dynamics of context, we argue that they must be represented in the logical form of discourse. Thus, we arrive at a picture where many alleged implicatures are actually a consequence of logical form. Although these features of interpretation are still derived by the operation of abstract principles, the principles in question are linguistic rules rather than rational or psychological generalizations; they are ultimately continuous with the abstract operations of formal compositional semantics.

We take (1) as a characteristic example of the role of coherence relations in underpinning apparently indirect interpretations. *Imagination and Convention* makes the case informally, but we offer an extended, formal treatment of cases like (1) in Lepore and Stone (2017a). This was the subject of our presentation at the 2015 workshop on linguistics and philosophy in Dubrovnik.

The challenge of (1) is to formalize the differences among declarative, interrogative and imperative meanings. We model these differences, following Starr (2010), in terms of different roles information can play in moving conversation forward. Declaratives convey information; interrogatives raise questions; imperatives express preferences. Starr's formalism gives a dynamic model of the state of a conversation that can distinguish among contributions of each of these kinds, and can also predict certain inferential relationships among them—thus, for example, conveying the right information can settle an open question. We also need to be able to combine different moves compositionally. Starr lets us combine two contributions into a single overarching move that starts by making the first contribution and proceeds by making the second; he lets us make a contribution conditionally, depending on the results of some other one.

With these tools, we can represent (1) as ambiguous between two logical forms at the level of discourse. One, the simple question interpretation, just raises a question: here, the question whether it is possible for the speaker to have the French toast. The second, the "indirect" interpretation, raises that same question, then further expresses a conditional preference: here, the indirect interpretation raises the question whether it's possible for the speaker to have the French toast, then expresses the preference that the speaker should have the French toast, assuming the answer is yes and it is possible. We show that such alternations in meaning are characteristic of a kind of polysemy Horn (1984) calls 'autohyponymy'—often found in verb meanings—where words carry overlapping specific and general senses. We offer some

suggestions about capturing this polysemy, at an appropriate level of granularity, by a suitable linguistic rule.[1]

We think that representing these two interpretations of (1) in logical form shows the advantages of conventionalized coherence relations in giving a theory of interpretation. In particular, as *Imagination and Convention* considers in detail, there is ample evidence that the two interpretations are separately specified by speakers' knowledge of language, and, moreover, that the conventional indirect interpretation is visible to other grammatical rules.

We explain (2), meanwhile, by a different set of resources in linguistic meaning—the grammar of *discourse reference*. When we produce extended descriptions, narratives and explanations—including the one in (2)—grammar allows the interpretation of later elements to co-vary with the interpretations of earlier ones. Formally, this can be modeled by representing both elements with a common variable in logical form; however, to implement it correctly, we also need to set up an appropriate logical system so that we can assign values to variables across an entire discourse (this suggestion goes back to Heim 1982 and Kamp 1981—see Cumming 2008 for a broader defense and philosophical explanation of the idea).

We can think of a grammar of discourse reference in terms of two components. One set of grammatical rules determines where variables occur (we call these 'rules for anaphora' in *Imagination and Convention*); the second set of grammatical rules says how the selection of a suitable variable is determined in context (we call these 'rules for presupposition' in *Imagination and Convention*). The idea of capturing dependent interpretations via variables that are subject to constraints is common to diverse approaches to formal semantics and pragmatics, including not only van der Sandt (1992) but also the very different Hobbs et al. (1993).

We can illustrate this idea through an explanation of the understood temporal relationship in (2). Following Lascarides and Asher (1993), Partee (1973) and Webber (1988), we assume that the tense of past tense English verbs can trigger a dependent temporal interpretation. That is, a past tense verb describes an event or state as located within a specific temporal interval, its reference time. In (2), then, 'plunged' has a meaning similar in content to 'then plunged'. The reference interval is taken from context in a way that gives it an interpretation that can depend on previous discourse. In (2), the reference time for 'plunged' is derived from the event time for 'doubled'. Importantly, these intervals progress in a discourse as a function of the coherence relations that

[1] The idea is to use lexical rules—defaults that apply across general classes of words, with exceptions—to transform basic meanings into related, derived meanings. Such rules are needed quite independently, for example, to stipulate that the names of animals are also used as the names of meat, with a few marked exceptions including 'beef' and 'pork'. This strategy for capturing conventional indirect speech acts was originally proposed by Asher and Lascarides (2001).

implicitly connect the discourse together: in *Narrative* discourse, for example, event verbs update the most prominent reference interval to a period immediately after the event took place, when its consequences continued to hold. That's what happens in (2). Thus, overall, we explain the interpretation of (2), that the plunge follows the doubling, because we represent the meaning of the form 'plunged' (in particular, its tense and aspect) as locating that event within a reference interval *after* the doubling, a reference interval that is made prominent by the preceding use of the form 'doubled' as part of an extended discourse organized by the *Narration* relation.

Researchers have developed a range of different formal models of presupposition and anaphora. Since completing the book, we have been exploring a particularly strict conception of the rules for context dependence, in our collaborative work with Una Stojnic (Stojnic, Stone and Lepore 2017). Our proposal is that the state of the discourse completely determines which variable should be used to interpret a dependent, context-sensitive element. For example, just as 'I' picks out a distinguished semantic value in any context—namely, the speaker of the utterance—just so, 'he' picks out a distinguished semantic value—a variable that has been established by discourse coherence as the representation of the most prominent male with respect to the place of the current clause within the organization of discourse.

This approach depends on a synthesis of our approaches to discourse coherence and discourse reference. Many researchers have noticed that when the interpretation of pronouns and other anaphoric elements seems to be ambiguous, there are also corresponding ambiguities in the overall coherence of discourse (see Kehler et al. 2008 for review). Take (4), originally studied by Smyth (1994) and discussed extensively by Kehler et al. (2008).

(4) Phil tickled Stanley and Liz poked him.

The speaker here might mean Phil by 'him', but in this case the speaker is describing Liz's action as a sequel to, and perhaps even as a retaliation for, Phil's tickling. Alternatively, the speaker might mean Stanley by 'him', but in this case the speaker's point is to draw an analogy between Phil's and Liz's attacks on Stanley; this interpretation doesn't seem to involve any commitments about whether the poking preceded, followed or was simultaneous with the tickling. In short, the discourse in (4) is organized either via a kind of *Narration* or via a kind of *Resemblance* relation, and this relationship gives us the value of the pronoun.

In Stojnic, Stone and Lepore (2017), we offer a formalization of this idea that makes precise the effects that coherence relations have on the prominence of candidate interpretations, and makes good on the intuitive idea that pronouns are interpreted simply by retrieving the most prominent candidate interpretation in context. In her (2016a), Stojnic develops an analogous approach to the context dependence of modal vocabulary, and in her (2016b) she even proposes to handle quanti-

fier domain restriction and incomplete definite descriptions with these techniques.

The last of our motivating examples, (3), is a reminder that natural language utterances generally have a more complex linguistic structure than orthography alone captures. Nevertheless, all of the grammatical components of an utterance can carry encoded meanings. What matters for (3) is *intonation*. When we imagine (3) used, as Kripke (1978) does, to challenge a previous speaker's contention that the handkerchief in a magic act was not red, we tend to imagine the utterance delivered in a particular way. The speaker will emphasize 'looked' rather than 'it' or 'red'; the speaker will perform the utterance with a particular tune (or 'pitch contour') with a rise on 'looked', followed by a fall, so that 'red' comes with a rise of its own at the end. These aspects of the performance of the utterance are meaningful—they figure in the English grammar of information structure, which characterizes the different roles of linguistic material in making contributions to discourse.

We give a comprehensive survey of information structure in the book, focusing on intonation and drawing particularly on the work of Steedman (2000). But you can already explain the distinctive interpretation of (3) with reference to Ward and Hirschberg's (1985) theory of the rise-fall-rise contour. They suggest that this tune is associated with limited agreement in discourse, while the placement of accents signals a point of contrast relevant to that limited agreement. In other words, the intonation of (3) encodes the fact that the speaker cannot completely agree with the prior suggestion that the handkerchief was red: although the handkerchief did look red, there can be a contrast between how something looks and how something actually is. This meaning is signaled by the grammar of (3), not derived by implicature. Information structure in its full generality, we suggest, has far-reaching consequences for many other cases of alleged implicature as well.

In hindsight, we would draw a broader message from the discussion of information structure and intonation in *Imagination and Convention*. A full treatment of the logical form of utterances may have to incorporate the contributions of a wider range of communicative actions than philosophers of language typically consider. Take deictic gestures, for example, which normally accompany demonstrative noun phrases (like 'this' or 'that' in English). The received view from Kaplan (1989) is that gestures are nonlinguistic cues that let a speaker provide evidence about the referent they intend. However, many cognitive scientists—including McNeill (1992) and Kendon (2004)—see gesture and language as part of a single, integrated system for making our ideas public. That suggests that we can and should represent the interpretation of speech and gesture in a single formalism (Lascarides and Stone 2009), and even derive the interpretation of speech and gesture compositionally (Alahverdzhieva and Lascarides 2011, Giorgolo 2010).

When we adopt such theories, we may be led to significant departures from traditional views of context sensitivity in philosophy. For

example, in Stojnic, Stone and Lepore (2013), we provide a formalism where even the demonstratives 'this' and 'that' turn out to get their values directly as a function of the context—just like pure indexicals. The tools we use are parallel to those in Stojnic, Stone and Lepore (2017). We give a grammatical analysis of pointing gestures; their meanings update the context in which a subsequent demonstrative is interpreted by putting particular entities and situations at the center of attention. In the resulting context, the demonstrative automatically gets its correct, context-dependent semantic value.

In sum, we encourage readers to regard the case studies of discourse coherence anaphora and presupposition, and information structure that we consider in *Imagination and Convention*—as exhibited in the interpretations of (1–3)—merely as an indication of the diversity and importance of linguistic semantics in guiding utterance understanding. We think many more rules remain to be uncovered, and correspondingly, that there is much more to say about the linguistic knowledge that underpins interpretation. We encourage students of language to appreciate the ways in which appeals to Gricean reasoning—the idea that the audience simply constructs whatever interpretation makes sense, purely by intuition and common sense—forecloses inquiry that could expose and characterize such knowledge. The effect, we think, is both tempting and insidious.

We know that utterances make sense, and accordingly, that we can resolve ambiguities in part by considering what we know about the speaker. As Blackstone (1765) says, in describing the considerations of jurisprudence that should go into the interpretation of the language of a statute:

> The fairest and most rational method to interpret the will of the legislator, is by exploring his intentions at the time when the law was made, by signs the most natural and probable. And these signs are either the words, the context, the subject-matter, the effects and consequence, or the spirit and reason of the law. (Blackstone 1765: I, Introduction, §2)

When we read the alleged 'derivations' of Gricean pragmatics, we find they often hint, retrospectively, at plausible reasons why interpreters might prefer an attested reading from other candidate interpretations. In these informal accounts of disambiguation, however, Grice tends simply to proffer the correct interpretation without explanation, so it's easy to lose track of the principles that derive and license the possible interpretations in the first place. Our experience is that—at least when utterances have a specific, clear interpretation—the relevant principles are always principles of grammar.

2. *Varieties of Interpretive Reasoning*

A different line of argument in *Imagination and Convention,* meanwhile, makes the case that researchers have often been too quick to distill the *points* that speakers make with utterances, and the *insights*

that hearers derive from them, in terms of propositional content. On our view, audiences approach utterances through diverse kinds of imaginative engagement, which, we think, philosophers and cognitive scientists must describe in diverse and generally non-propositional terms.

We give a number of examples of such effects in the book. Perhaps the deepest and most persuasive is the case of metaphor. The discourses that best illustrate our thesis are extended, novel metaphors that call for active engagement on the part of the audience. The example we like to cite is (5).

(5) Love is a snowmobile, racing across the tundra. It flips over, pinning you underneath. At night the ice weasels come. (Matt Groening, given as (1) in Lepore and Stone (2010) and as (177) in *Imagination and Convention*.)

Many discussions—particularly those in the tradition of Lakoff and Johnson (1980)—claim that semantics is rife with metaphor. They posit active spatial metaphors at work in the grammar of change (e.g., 'the light *went from* green *to* yellow'), the grammar of mental states (e.g., 'I can't *get* that idea *out of* my mind'), etc. This contrasts with a more conservative view in formal semantics: that such locutions really involve abstract meanings that apply across semantic domains (see Hobbs 2011). Focusing on examples like (5) enables us to sidestep this controversy about the pervasiveness of metaphor.

Similarly, it's clear that many word senses have metaphorical origins that audiences need not activate—and usually do not activate—as part of understanding them (Glucksberg 2001). For example, experienced speakers probably understand 'family tree' directly in terms of their concepts of genealogy, 'syntax tree' directly in terms of their concepts of grammatical derivation, or 'binary tree' directly in terms of their knowledge of algorithms and data structures, and not by entertaining a biological metaphor. Such cases, of course, are generally described as 'dead metaphors'. There may be philosophically interesting things to say about dead metaphors.[2] However, our arguments in *Imagination and Convention* focus just on (5) and other cases where metaphor is used productively and creatively. We think there's no way to account for their interpretation by postulating conventional semantics or by Gricean pragmatics that somehow delivers a 'metaphorical meaning'. We describe what's happening in these cases in a different way.

We start from the suggestion—originally due to Black (1955) but recently defended in detail by Camp (2003, 2008, 2009)—that appreciating a metaphor involves a distinctive kind of imaginative effort. One

[2] For example, in Lepore and Stone (2017b) we consider the possibility that one aspect of the connotation and Fregean tonality of words arises from the ability of a perceptive reader or listener to redeploy and draw insights from metaphors that other speakers might ignore as dead or dying.

key component of this effort is perspective taking: thinking of one thing as another. More precisely, metaphorical perspective taking requires an audience to construct a wide-ranging correspondence between entities in a source domain and entities in a target domain. Thus, for 'love is a snowmobile', we must not only imagine love as a snowmobile, but see the lovers as passengers, see the motion and mishaps of the vehicle as placeholders for the events that unfold in the course of a typical relationship, and—most importantly—appreciate the similarity in feeling between the lovers' experience of their affair and the snowmobile passengers' experience of their ride.

Moreover, as Camp (2008) describes it, metaphorical perspective taking involves a specific direction of fit. The point of a metaphor is not just to blend ideas together, or to remap the world in pretense: we don't interpret (5) just by imagining or pretending that lovers are riding a snowmobile. Instead, we draw on our knowledge of snowmobiles, selectively and judiciously, to find features that help us to appreciate corresponding aspects of the experience of love.

Once we characterize the workings of metaphor this way, we are led to a position similar to Davidson's (1978) about the philosophical status of metaphorical interpretation. Metaphor is not a case of Gricean reasoning: the insights of metaphor come not from working out what the speaker must have intended, but from engaging metaphorically with the text itself. Metaphorical insights are a product of the audience's private psychology in confrontation with challenging imagery. (Davidson memorably—though in our view not entirely accurately—compares the effects of a metaphor to those of a dream or even those of a bump on the head!)

The insights of metaphor need not even be propositional in nature. As with Wittgenstein's duck-rabbit figure, which viewers can see as one thing or as another, what matters in metaphor is the dynamics of experience—such factors as attention, memory and inference, as they are deployed in real time to organize and explain things around us. To try to boil this active process down to some specific information that the speaker of a metaphor intends to convey and that the audience aims to reconstruct is to miss what's really going on in this kind of language use.

In the book, we describe other literary effects in similar terms. We describe sarcasm as an invitation to appreciate an utterance framed in familiar terms as an inversion of what circumstances actually demand. We describe irony as an invitation to derive insights from engaging with the imagined speaker of an utterance exhibited in pretense. We describe humor as the appreciation of the potential of an utterance simultaneously to sustain two perspectives with opposed affective import. We describe hinting as an invitation for the hearer to formulate her own reactions to the theme of the utterance—perhaps guided by the associations of apparently incongruous or irrelevant detail.

We emphasize that these forms of imaginative engagement are neither mutually exclusive nor exhaustive. In fact, we have followed up this part of the book (Lepore and Stone 2016) by using the framework to explicate one particular ingredient in the interpretation of poetry.

Poets, of course, recruit all the expressive resources at their disposal. We often find that metaphorical language, in particular, can be particularly poetic. But, we suggest, something special happens when interpreters approach the articulation of language as poetry—regardless of the semantic content it has or the other components of their imaginative engagement with it.

In a poem, articulation itself is a meaningful part of the experience of understanding. The reader of poetry can attend to the sound, rhythm, lineation and even typography of the work, as a prompt to better understand it and to draw richer insights into the experience it affords. These cues can add sensual qualities that heighten the imagery of a poem, call attention to formal relationships within the poem that take on a corresponding importance, or help to attune the reader to the dynamic consciousness behind the poem, giving voice to a distinctive flow of perception, emotion and judgment. In describing poetic interpretation this way, we draw on a range of antecedents—from Pope (1711) through the New Critics of the mid 1900s (e.g., Brooks 1947) up to present-day scholars such as Longenbach (2008). All of these authors highlight similar aspects of what can happen when an audience appreciates a poem. However, critics have often been too ready to blur together the appreciation of a poem and its meaning. We maintain that poetic language retains its ordinary meanings—we think that this conclusion is philosophically inescapable. Nevertheless, we can appeal to the particular imaginative practice poetic language recruits in order to talk about the special effects of poetic language. In doing so, we discover that poetry involves a form of engagement which is distinct not only from the forensic project of reconstructing a speaker's communicative intentions, but also from the many other imaginative practices that speakers can invite that we already survey in *Imagination and Convention.*

Here, then, as in our account of the conventions of grammar, we hope that research is just beginning: there's room to address new kinds of data and to develop correspondingly refined accounts of the strategies that our psychology and culture gives us for enriching the experience of making sense of language.

3. *Theorizing Semantics and Pragmatics*

Theories of meaning have represented one of the most vital contributions of philosophy to cognitive science. Clear thinking and good examples have been instrumental in helping researchers to get clear on the ways that speakers exploit knowledge of language, knowledge of other people, and common sense in order to get their ideas across to

one another and to carry out joint projects together. The phenomena are so challenging precisely because such wide-ranging knowledge is implicated in every episode of communication.

We think the contemporary debate about the meanings of slurs is indicative of the difficulty that is involved (Lepore and Stone 2017b). As we noted in our introduction here, and survey in *Imagination and Convention,* the results of formal semantics offer many ways to associate words with evaluative content, including at-issue content, presupposed content, projective and not-at issue content, and expressive contents. These resources can be crucial for getting clear on the meanings of some words. But there are also good reasons not to regard all differences of nuance and tone as differences in conventional meaning. Words can still invite metaphorical thinking and other kinds of perspective taking, they can still prompt us to think of others who have used them, and they are subject to prohibitions and other social constraints that—regardless of how the taboos arise—endow their use with a special charge. Especially for problematic terms such as slurs, only a close look at their interpretive profile—in light of the full range of theoretical possibilities—can reveal what amounts to encoded meaning, as distinct from mere suggestions and connotations.

Encoded meaning itself is just the starting point for much of our interactions with one another. People use language intentionally, and people work hard to understand one another as people—not just to understand the language they use. Scientifically, there's ample evidence that intention recognition shapes the choices that interlocutors make in dialogue (Grosz and Sidner 1986), that it helps them avoid and recover from failures in communication (Brennan 2005), and that it's crucial for enabling infants to learn language in the first place (Bloom 2000).

Nothing we say undermines these obvious realities—nor do they undermine anything we say. The book offers an extended discussion of the cognitive architecture of collaborative language use in support of this claim. But the general lines of argument will be familiar; other semantic minimalists, such as Borg (2004), have presented similar lines of thinking.

This précis has emphasized the polemical side of our book and its implications. But we also wrote the book as a survey of the important theories, data and arguments that are relevant to mapping the relationships between semantics and pragmatics. We expect that few, if any, readers of the book will agree with us in all its particulars. But we are optimistic that readers will come away from our book with a feeling for how broad the phenomena are that bear on the interface of semantics and pragmatics, a deeper appreciation for the kinds of facts that seem to distinguish most strongly among the theoretical alternatives, and a road map to the open problems where future work is likely to bring challenges that all current theories must respond to. We hope you agree with us at least that the papers in this special issue illustrate these possibilities!

Acknowledgments

Preparation of this article was supported by NSF grant IIS 1526723 and a sabbatical leave from Rutgers to Stone.

References

Alahverdzhieva, K. and A. Lascarides. 2011. "An HPSG Approach to Synchronous Deixis and Speech." In S. Muller (ed.). *Proceedings of the 18th International Conference on Head-Driven Phase Structure Grammar (HPSG)*. Seattle: 6–24.

Atlas, J. D. 1989. *Philosophy without Ambiguity*. Oxford: Clarendon Press.

Atlas, J. D. 2005. *Logic, Meaning, and Conversation Semantical Underdeterminacy, Implicature, and Their Interface*. Oxford: Oxford University Press.

Asher, N. and A. Lascarides. 2001. "Indirect speech acts." *Synthese* 128 (1–2): 183–228.

Asher, N. and A. Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.

Beaver, D. 2001. *Presupposition and Assertion in Dynamic Semantics*. Stanford: CSLI Publications.

Black, M. 1955. "Metaphor." *Proceedings of the Aristotelian Society*, ns 55: 273–94.

Blackstone, W. 1765. *Commentaries on the Laws of England*. Oxford: Clarendon Press.

Bloom, P. 2000. *How Children Learn the Meanings of Words*. Cambridge: MIT Press.

Borg, E. 2004. *Minimal Semantics*. Oxford: Oxford University Press.

Brennan, S. E., 2005. "How conversation is shaped by visual and spoken evidence." In J. C. Trueswell and M. K. Tanenhaus (eds.). *Approaches to studying world-situated language use: Bridging the language-as-product and language-as-action traditions*. Cambridge: MIT Press: 95–129.

Brooks, C. 1947. *The Well Wrought Urn*. New York: Harcourt Brace.

Camp, E. 2003. *Saying and Seeing-as: The Linguistic Uses and Cognitive Effects of Metaphor*. University of California at Berkeley (dissertation).

Camp, E. 2008. "Showing, Telling and Seeing: Metaphor and 'Poetic' Language." In Elisabeth Camp (ed.). *The Baltic International Yearbook of Cognition, Logic and Communication*, *III. A Figure of Speech*. Riga: University of Latvia: 1–24.

Camp, E. 2009. "Two Varieties of Literary Imagination: Metaphor, Fiction and Thought." In Howard Wettstein (Ed.), *Midwest Studies in Philosophy*, XXXIII. *Poetry and Philosophy*. Oxford: Wiley-Blackwell: 107–30.

Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Davidson, D. 1978. "What Metaphors Mean." *Critical Inquiry*, 5 (1): 41–7.

Giorgolo, G. 2010. "A Formal Semantics for Iconic Spatial Gestures." In M. Aloni, B. Harald, T. de Jager and K. Schulz (eds.). *Logic, Language and Meaning*. Springer: 305–314.

Glucksberg, S. 2001. *Understanding Figurative Language: From Metaphors to Idioms*. Oxford: Oxford University Press.

Grosz, B. J., and C. L. Sidner. 1986. "Attention, Intentions, and the Structure of Discourse." *Computational Linguistics* 12 (3): 175–204.

Hobbs, J. R. 1979. "Coherence and Coreference." *Cognitive Science* 3 (1): 67–90.

Hobbs, J. R. 2011. "Word Meaning and World Knowledge." In C. Maienborn, K. von Heusinger, P. Portner, and N. van Leusen (eds.). *Semantics: An International Handbook of Natural Language Meaning*, Volume 1. Mouton de Gruyter: 740–761.

Hobbs, J. R., M. Stickel, D. Appelt, and P. Martin. 1993. "Interpretation as Abduction." *Artificial Intelligence* 63 (1–2): 69–142.

Horn, L. R. 1984. "Toward a New Taxonomy for Pragmatic Inference: Q-Based and R-Based Implicature." In D. Schiffrin (ed.). *Meaning, Form, and Use in Context: Linguistic Applications*. Washington, DC: Georgetown University Press: 11–42.

Kaplan, D. 1989. "Demonstratives." In J. Almog, J. Perry, and H. Wettstein (eds). *Themes from Kaplan*. Oxford: Oxford University Press: 481–563.

Kehler, A. 2001. *Coherence, Reference and the Theory of Grammar*. Stanford: CSLI.

Kehler, A., L. Kertz, H. Rohde, and J. L. Elman. 2008. "Coherence and coreference revisited." *Journal of Semantics* 25 (1): 1–44.

Kendon, A. 2004. *Gesture: Visible Action as Utterance*. Cambridge: Cambridge University Press.

Kripke, S. 1978. "Speaker's Reference and Semantic Reference." In P. A. French, T. E. Uehling, Jr., and H. K. Wettstein (eds.). *Contemporary Perspectives in the Philosophy of Language*. Minneapolis: University of Minnesota Press: 6–27.

Lakoff, G. and M. Johnson. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.

Lascarides, A. and N. Asher. 1993. "Temporal Interpretation, Discourse Relations and Commonsense Entailment." *Linguistics and Philosophy* 16 (5): 437–93.

Lascarides, A. and M. Stone. 2009. "A Formal Semantic Analysis of Gesture." *Journal of Semantics* 26 (3): 393–449.

Lepore, E. and M. Stone. 2010. "Against Metaphorical Meaning." *Topoi* 29 (2): 165–80.

Lepore, E. and M. Stone. 2016. "The Poetic Imagination." In N. Carroll and J. Gibson (eds.). *The Routledge Companion to the Philosophy of Literature*. Routledge: 323–333.

Lepore, E. and M. Stone. 2017a. "Explicit Indirection." To appear in Fogal, Harris and Moss (eds.). *New Work on Speech Acts*. Oxford: Oxford University Press.

Lepore, E. and M. Stone. 2017b. "Pejorative Tone." To appear in D. Sosa (ed.). *Bad Words*. Oxford: Oxford University Press.

Longenbach, J. 2008. *The Art of the Poetic Line*. Saint Paul: Graywolf Press.

McNeill, D. 1992. *Hand and Mind: What Gestures Reveal about Thought*. Chicago: University of Chicago Press.

Montague, R. 1974, *Formal Philosophy. Selected Papers of Richard Montague*, Thomason, R. H. (ed.), New Haven: Yale University Press.

Partee, B. H. 1973. "Some Structural Analogies between Tenses and Pronouns in English." *Journal of Philosophy* 70 (18): 601–9.

Pinker, S., M. A. Nowak, and J. J. Lee. 2008. "The Logic of Indirect Speech." *Proceedings of the National Academy of Sciences of the United States of America* 105 (3): 833–8.

Pope, A. 1711. *An Essay on Criticsm*. London: W. Lewis.

Potts, C. 2005. "The Logic of Conventional Implicatures." *Oxford Studies in Theoretical Linguistics*. Oxford: Oxford University Press.

Recanati, F. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.

van der Sandt, R. A. 1992. "Presupposition Projection as Anaphora Resolution." *Journal of Semantics* 9 (4): 333–77.

Searle, J. R. 1979. "Metaphor." In A. Ortnony (ed.). *Metaphor and Thought*. Cambridge: Cambridge University Press: 265–77.

Sperber, D. and D Wilson. 1986. *Relevance: Communication and Cognition*. Cambridge: Harvard University Press.

Sperber, D. and D. Wilson. 2008. "A Deflationary Account of Metaphor." In R. W. Gibbs (ed.). *The Cambridge Handbook of Metaphor and Thought*. Cambridge: Cambridge University Press: 84–105.

Simons, M., D. Beaver, C. Roberts, and J. Tonhauser. 2016. "The Best Question: Explaining the Projection Behavior of Factives." *Discourse Processes*. To appear.

Smyth, R. 1994. "Grammatical determinants of ambiguous pronoun resolution." *Journal of Psycholinguistic Research* 23 (3): 197–229.

Starr, W. B. 2010. *Conditionals, Meaning and Mood*. Rutgers University (PhD Dissertation).

Steedman, M. 2000. "Information Structure and the Syntax–Phonology Interface." *Linguistic Inquiry* 31 (4): 649–89.

Stojnić, U. 2016a. "Content in a Dynamic Context." To appear in *Philosophy and Phenomenological Research*.

Stojnić, U. 2016b. *Context Sensitivity in a Coherent Discourse*, Rutgers University (PhD Dissertation).

Stojnić, U., M. Stone, and E. Lepore. 2013. "Deixis (Even without Pointing)." *Philosophical Perspectives* 26 (1): 502–25.

Stojnić, U., M. Stone, and E. Lepore, 2017. "Discourse and Logical Form: Pronouns, Attention and Coherence." To appear in *Linguistics and Philosophy*.

Tonhauser, J., D. Beaver, C. Roberts, and M. Simons. 2013. "Toward a Taxonomy of Projective Content." *Language* 89 (1): 66–109.

Travis, C. 1997. "Pragmatics." In Hale, B. and Wright, C. (eds.). *A Companion to the Philosophy of Language*. Oxford: Blackwell: 97–107.

Ward, G. and J. Hirschberg. 1985. "Implicating Uncertainty: The Pragmatics of Fall-Rise Intonation." *Language* 61 (4): 747–76.

Webber, B. L. 1988. "Tense as Discourse Anaphor." *Computational Linguistics* 14 (2): 61–73.

Webber, B. L., M. Stone, A. Joshi, and A. Knott. 2003. "Anaphora and Discourse Structure." *Computational Linguistics* 29 (4): 545–88.

Croatian Journal of Philosophy
Vol. XVI, No. 47, 2016

Against Lepore and Stone's Sceptic Account of Metaphorical Meaning

ESTHER ROMERO, BELÉN SORIA*
Department of Philosophy I
and Department of English and German Philology
University of Granada, Spain

In this paper, we discuss Lepore and Stone's account of metaphor which is based on three of Davidson's proposals: (i) the rejection of metaphorical meanings; (ii) the rejection of metaphors as conveying metaphorical propositional contents; and (iii) the defence of analogy as the key mechanism for understanding metaphors. Lepore and Stone defend these proposals because the non-sceptic strategy on metaphorical meanings, characterized in general by the negation of (i) and (ii), fails to come to grips with neither the power of metaphor nor the explanatory resources of traditional pragmatic theories. In this paper we show not only how our non-sceptic account of metaphorical meaning as a variety of ad hoc concept eliminates these difficulties but also how it can solve two related difficulties that appear in Lepore and Stone's account. One of them is that Lepore and Stone's account involves the possibility of interpreting all metaphorical utterances literally (metaphors only have one meaning, the ordinary meaning) as a criterion of metaphorical identification; the other is that their proposal is not suited for explaining how speakers can agree or disagree when they use metaphorical utterances.

Keywords: Metaphorical meaning, analogy, pragmatics, metaphorical identification, metaphorical interpretation.

* A version of this article was given at the Philosophy of Linguistics and Language Course (Dubrovnik 2015) where Ernie Lepore and Mathew Stone's latest book was discussed. We wish to express our gratitude to the members of that audience, especially to Lepore and Stone, for various discussions and insightful comments. Thanks also to John Keating for suggestions on an earlier draft. This research was supported by the Spanish 'Ministerio de Economía y Competitividad" through the project FFI2011–26418.

1. *Introduction*

In this paper we discuss the account of metaphor proposed by Ernie Lepore and Matthew Stone in chapter ten of their latest book *Imagination and Convention* (2015) titled "Perspective Taking: Metaphor" and in two of their previous publications, "Against Metaphorical Meaning" (2010) and "Philosophical Investigations into Figurative Speech Metaphor and Irony" (2014). Following the lead of Davidson (1978), Lepore and Stone do not recognize any role for metaphorical meanings in a theory of meaning and reject that the interpretive effects of metaphors are propositional in nature. Metaphorical usage does not carry a meaning its speaker is trying to communicate: there is no metaphorical communication. This is not our position and we sometimes feel the titanic challenge to convince the sceptic that metaphorical communication is possible and can be explained if we take into account that metaphorical utterances convey propositional contents, what is metaphorically said, that are explained appealing to metaphorical provisional meanings, a variety of *ad hoc* concepts. These propositional contents are directly communicated, not implicated, much less intimated, suggested, or merely caused (Romero and Soria 1997/98).

At first sight, our approach to metaphor is at the other theoretical end of Lepore and Stone's scepticism towards metaphorical meanings but, far from what might be expected, we do not disagree on everything when dealing with their approach to metaphor. Thus, in the next section of this paper we focus on the points of agreement and on the differences related to our respective conceptions of metaphor. We agree on the crucial point about metaphor interpretation since we all defend the distinctiveness of metaphor and consider analogy as its key mechanism. Nevertheless, our main point of disagreement is on their claim that the result of metaphorical analogical thinking is not part of speaker meaning. We present this proposal and the way in which Lepore and Stone articulate it in the third section. Our disagreement with it leads us, in the fourth part of this paper, to provide the arguments for our defence of metaphorical meaning. To explain our proposal, which we argue is more explanatory, we first specify two conditions to identify a use of language as metaphorical: contextual abnormality and conceptual contrast. Second, we explain how the interpretation is achieved by means of a pragmatic process of context-shifting that affects the language parameter. The analogical mapping from source domain to target domain, which results in a metaphorical restructuring of a concept, permits us to construct provisional meanings for the words used metaphorically and thus the language parameter in that context changes. In this way, our pragmatic theory of metaphorical provisional meaning does not fail to come to grips with the power of metaphor since this meaning is conceivable only from the metaphorically restructured concept.

2. *Points of Agreement and disagreement*

We agree with Lepore and Stone on several points. To begin with, we all defend the position that a theory of metaphor must specify the peculiar characteristics of metaphor[1] and what is distinctive about metaphor is that it is related to a distinctive process of perspective taking. Perspective taking in metaphor, or as we usually call it "analogical reasoning" (Romero and Soria 2014), consists in using information about a domain to organize the information about another domain with which it is not previously related and this is done through an analogical correspondence (mapping). We then also agree that metaphor can issue in distinctive cognitive and discourse effects. As Lepore and Stone claim

> metaphorical interpretation involves a distinctive process of PERSPECTIVE TAKING. Metaphor invites us to organize our thinking about something through an analogical correspondence with something it is not. (2015: 162)[2]

Secondly, we also agree that metaphor recognition (identification) is essential to metaphorical interpretation. An account of metaphorical utterances must include the features of metaphorical identification. If (1)

(1) Love is a snowmobile racing across the tundra and then suddenly it flips over, pinning you underneath. At night, the ice weasels come. (Example taken from the comedian Matt Groening as quoted by Lepore and Stone 2010: 165)

is a metaphor, (1) must be identified as such. This identification triggers its metaphorical interpretation. As Lepore and Stone claim "[w]e must *recognize* that (1) is a metaphor, and *shape* our psychological response accordingly." (2010: 171).

For them, "the insights of metaphor and the distinctive import of words used metaphorically frequently go beyond the conventional rules of language." (2015: 162). We agree that dead or conventional metaphors must be excluded from the study of the interpretation of active, creative metaphors (Romero and Soria 1998, 2005a, 2005b). Conventional metaphors, following Lakoff and Johnson (1980)'s characterization, are metaphorical utterances that include either expressions of

[1] Not all theorists have defended the position that there are some peculiar characteristics of metaphor. For example, Sperber and Wilson have stated that "there is no mechanism specific to metaphors, and no interesting generalization that applies only to them" (2008: 84). According to them, metaphor is interpreted in the same way as other loose uses. The inferential process is guided by the Relevance Principle and results in the loosening or weakening of the lexical encoded concept by dropping part of its logical entry in the process of arriving at the intended interpretation (Sperber and Wilson 1986/95: 233–237). See our 2014 for critiques to this view on metaphor.

[2] This distinctive process has been argued by Richards (1936), Black (1954), Davidson (1978), Lakoff and Johnson (1980), Kittay (1987), Romero and Soria (1997/98) and many others. Nevertheless, not all theorists argue for this process. For example, relevance theorists claim that the metaphorical mechanism is a case of loosening (see the previous footnote) and Stern argues for metaphor interpretation as a case of saturation of a metaphorical operator (Stern 2000).

the used part of a conventional metaphorical concept or expressions of marginal metaphorical concepts that must be interpreted literally (or conventionally). A normal utterance of (2)

(2) The foundations of my theory are sure.

includes a used part of the conventional metaphorical concept THEORIES AS BUILDINGS: 'foundations'. The expression 'the foot of the mountain' in an utterance of (3)

(3) We reached the foot of the mountain.

is the only expression used from the marginal metaphorical concept MOUNTAIN AS PERSON.

The usual utterances of (2)–(3) are interpreted literally. They are called "metaphors" because they give expression to metaphorical concepts but they are "literal utterances" in the sense that they are identified as literal and must be interpreted literally.

For a literal proposition to be expressed, linguistic expressions must appear in normal linguistic and extralinguistic contexts. If we consider (4),

(4) [Sarah asks Marian where her pet is and she answers:] My cat is on the mat.

the context (linguistic and extralinguistic) of every word of the uttered sentence coincides with one of the potential contexts fixed for them in the linguistic competence of the speaker. In cases like this, a literal use of language is identified and the conventional interpretation is triggered. The same can be said about the so called "conventional metaphors", about usual utterances of (2)–(3). They must be interpreted literally although they involve conventional metaphorical concepts. Thus, conventional metaphors must be excluded from an account of metaphorical interpretation.

Novel metaphors instead demand a creative and distinctive process of interpretation. They are metaphorical utterances that include either expressions of the imaginative uses of conventional metaphorical concepts or expressions of new metaphorical concepts. In an utterance of (5)

(5) His theory has thousands of little rooms.

'thousands of little rooms' is an instance of the unused part of a usual conventional metaphorical concept, THEORIES AS BUILDINGS. In an utterance of (6)

(6) These facts are the bricks of his theory.

'bricks' is an extension of one of the used parts of THEORIES AS BUILDINGS: 'the outer shell'. Example (1) above is, according to Lepore and Stone (2015: 163), an unused part of a conventional metaphorical concept: LOVE AS A JOURNEY.[3]

[3] The conventional metaphorical concept LOVE (AS A JOURNEY) inherits the structure of the conventional metaphorical concept LIFE (AS A JOURNEY). "What is special about the *LOVE IS A JOURNEY* metaphor, is that there are two lovers, who are travellers, and that the love relationship is a vehicle. The rest of the mapping is a consequence of

In addition, in an utterance of (7)

(7) Classical theories are patriarchs who father many children most of whom fight incessantly.

'patriarchs who father many children most of whom fight incessantly' calls forth the new metaphorical concept, CLASSICAL THEORIES AS PATRIARCHS WHO FATHER MANY CHILDREN MOST OF WHOM FIGHT INCESSANTLY, which represents a new way of thinking. Any use of this concept is imaginative.

Finally, we also coincide in arguing that the special kind of perspective taking that characterizes metaphorical conceptualization is not propositional in nature. We maintain that analogical thinking delivers metaphorical concepts rather than propositions, concepts analogically restructured by means of the content of another: a conceptualization of one thing AS another. The imagery, the perspective we are taking on the subject-matter, as Lepore and Stone claim, does not proceed in many occasions constituent-by-constituent but across extended discourses. Understanding a metaphor involves improvising correspondences from properties and relationships of the source domain to corresponding ones in the target domain and the properties and relations of the source domain are often expressed by (simple or complex) expressions in clauses and may extend across a longer text showing cohesive ties. In any case, the metaphorical vehicle (the part of the metaphorical utterance that is metaphorically attributed to its topic) can, but does not have to, be a single word. This can easily be shown with (7) where the metaphorical vehicle is 'patriarchs who father many children most of whom fight incessantly'. Similarly, in (1), the metaphorical vehicle of metaphor is 'a snowmobile racing across the tundra and then suddenly it flips over, pinning you underneath. At night, the ice weasels come'.

But not all are agreements. We also disagree on several issues. Even if our disagreement does not entail a direct confrontation with their claim below,

> [w]e are defenders of grammar, of meaning, and of common sense. We are exponents of the richness of human experience and the creativity of language. We believe in drawing useful and principled distinctions. (Lepore and Stone 2015: v)

their omission of pragmatics (based on inferential intention-attribution) is our main source of disagreement. We are not only defenders of grammar, of meaning, and of common sense, we are also defenders of inferential pragmatics. Furthermore, the introduction of inferential pragmatics in the agenda is, in our opinion, crucial to draw an account of how metaphorical meanings are conveyed as intended and thus of how communication through metaphor is possible. If we are right, the distinction between the metaphorical and non-metaphorical use of lan-

inheriting the *LIFE IS A JOURNEY* metaphor." (Lakoff 1993: 223). As we can see in this quotation, for Lakoff and Johnson the term "metaphor" stands for the metaphorical concept or for the related expression rather than for the metaphorical utterance which is the one we favour.

guage does not depend on locating metaphor in the field of imagination and out of communication.

In the literature on metaphor, we often find a defence of the metaphorical mechanism as analogical mapping without denying a role for metaphorical meaning. Metaphorical meanings, if required for the correct explanation of the metaphorical use of language, have been considered as distinctive derived meanings that form a part of speaker meaning. Some theorists have considered that these derived or non-conventional meanings are involved in what is implicated, others have argued that they form part of what is said. According to the first position, the speaker makes as if to say one thing in order to mean another. According to the second, the speaker means what she metaphorically says. In both proposals, a pragmatic account of metaphorical meaning as a result or as a by-product of analogical thinking is possible.

Following the Gricean notion of implicature (Grice 1975/89) and Black's (1954) interaction theory on metaphor, Kittay (1987) elaborates her perspectival theory on metaphor with which she explains how the inferential analogical process reaches second-order meanings that intervene in metaphorical implicatures.[4] In contrast, we argue, following Indurkhya's (1986) mapping approach, that metaphorical meanings are improvised from the analogical reconceptualization of a domain (target domain) by means of a partial mapping of information from a different unrelated domain (source domain). This metaphorical reconceptualization of the target constitutes a new context of interpretation that delivers what is metaphorically said. A propositional content that includes what can be called, in Lepore and Stone's vein, "improvised" meanings[5] or, in our terminology, "provisional" meanings (ad hoc concepts) for the expressions used abnormally (non-conventionally).

Lepore and Stone (2010) argue that, if the notion of speaker meaning were developed coherently, there would be no room for any communicated propositional meaning in the metaphorical interpretation of utterances. In their opinion, the view of the metaphorical mechanism

[4] In general, we could say that the authors that have defended literalism in what is said have often argued for metaphor as implicature as well. In literalism, the sentence's linguistic meaning (understood as a compositional meaning that results from the combination of the conventional meanings of sentence terms) is closely related to what is said by an utterance of the sentence and the contextual information that is involved in the latter is demanded by the linguistic meaning itself. Since the pragmatic contribution needed to interpret metaphor is not closely related to the conventional meaning of sentence terms, it intervenes in an implicature and not in what is said. Authors such as Stanley (2005) argue for this proposal.

[5] Lepore and Stone (2015: 258) defend the idea that improvised meanings are possible as long as the speaker is able to specify and clarify this meaning enough for the purposes of the conversation so that it can contribute information to the conversational record. We think that metaphorical meaning is a type of improvised meaning, a case where conventions exhibit variation constrained by the identification conditions and the systematic interpretation mechanism of novel metaphorical utterances.

as analogical thinking is compatible only with the rejection of metaphorical meaning (Lepore and Stone 2015: 170). The previous proposals on the distinctive cognitive value of metaphor do not agree on the role of metaphor in speaker meaning. It can be formulated on the basis of the rejection of metaphorical meaning or on the basis of the discussion of whether metaphorical meaning is involved in what is implicated or in what is said by an utterance. All these positions are currently defended, as we show in Figure 1, but the crucial issue of this debate is whether there are metaphorical meanings or not.

Figure 1. Theories on the distinctive cognitive value of metaphor as analogy

No metaphorical meanings		Lepore and Stone (2010, 2014 and 2015)
Metaphorical meanings	What is implicated	Kittay (1987)
	What is said	Romero and Soria (1997/98, 2007 and 2013)

In our opinion, the notion of speaker meaning, rightly understood, can account for non-conventional uses of language at the explicit level since the speaker meaning may include metaphorical provisional meanings. To argue for that, we are going to focus on the rejection of two of Lepore and Stone's claims of their sceptic account on metaphorical meaning:

i. "(…) our insights in metaphorical thinking are prompted *just* by the literal meanings of utterances." (Stone and Lepore 2010: 175, our emphasis).

ii. The explanation of creative metaphorical meanings as a result of metaphorical thinking must be rejected because "this strategy fails to come to grips not only with the power of metaphor itself but also with the explanatory resources of traditional pragmatic theories" (Lepore and Stone 2015: 169).

To clarify why we disagree with these two points, we need to explain in more detail Lepore and Stone's view against metaphorical meaning.

3. *Lepore and Stone's views on metaphor*

Lepore and Stone say, quoting Emerson (1836: 34–35), that "Man is an analogist" (2015: 171n5). They, as many others, accept that the power of metaphor consists in using knowledge of one domain to give a perspective on something else.

Nevertheless, unlike those arguing that the application of this distinctive mechanism results in a peculiar propositional content for metaphorical utterances at the level of speaker meaning, Lepore and Stone claim that the recognition of the distinctive value of metaphor necessarily takes us to deny metaphorical meaning.

> [t]he information we get through a metaphor that comes from this process is not pragmatic, in the sense of not part of speaker meaning, not signaled by

> the speaker or recognized by the hearer. It becomes an extension of the external world, a place where our perceptions and demonstrations can inform our thinking and interaction we do with one another, but not part of our communication, that is, not part of the communicative enterprise. (Lepore and Stone 2014: 83)

Their rejection of metaphor as part of the communicative enterprise is based on arguing that metaphorical utterances are used to draw the hearer's attention to similarities. Speaker's goal is to make the hearer see similarities not to assert them. The hearer understands the proposition literally expressed and this understanding prompts him to look for certain similarities (Lepore and Stone 2010: 170 and 2015: 164).

Speaker meaning definitions developed by Grice (1957/89 and 1969/89) or by Lewis (1969 and 1979) permit Lepore and Stone to show why they reject the alleged metaphorical meaning. The nature of speaker meaning is, according to Grice (1957/89), determined by a group of intentions aimed to produce certain effects by means of the recognition of speaker's intention to produce them. The recognition of speaker's intention to produce certain effects is a necessary condition for the audience and a reason to reach them. In positing metaphorical meaning, this must be identified as part of what a speaker intentionally means and this requires, if we consider Grice's definition, an audience to recognize a specific content a speaker wants to get across. Grice's notion of speaker meaning is of no use to explain metaphorical meaning because the hearer's appreciation of the similarities is not achieved by means of the hearer's recognition of that intention (they take this as a reason parallel to Grice's reason to exclude Herod's showing John's head as an act with speaker meaning). The metaphorist has the intention that the audience appreciate certain similarities, but these are not reached by means of recognizing the speaker's intention to get across a propositional content in which a metaphorical meaning is involved.

Taking into account Lewis' characterization of speaker meaning does not change the situation. Lewis (1969) characterizes speaker meaning as the speaker's intention to coordinate with an audience to update the conversational record. Lewis considers not just speaker's intention but also the kinds of situation where agents face signaling problems. In his approach, the speaker means a propositional content by uttering a sentence if she intends to update the conversational record with that proposition by coordination. In positing metaphorical meaning, this must be a part of what a speaker intentionally means and this requires the signal of the metaphor to be used as the basis for the uptake of that content. Nevertheless, when a metaphorical utterance is produced by a speaker, her audience and she do not add a propositional content to the conversational record to satisfy their joint interests, including their interest in agreeing on the record. Once the utterance is recognized as a metaphor by means of the recognition of speaker's literal meaning, they do not have a mutual expectation of agreeing on adding a metaphorical content to update the record, but of appreciating a similarity.

The goal of a metaphorical utterance is "not for specific information to be exchanged, and interlocutors do not coordinate on the information itself or derive it directly by intention recognition." (Lepore and Stone 2010: 171). The speaker has a distinctive metaphorical thought but the hearer just gets the literal (absurd, irrational) linguistic meaning. If we reason from it, we would just rediscover the absurdity. The essential aspects of this activity cannot be characterized in terms of a communicative enterprise because although the literal proposition is made public, it cannot be considered as information to update the record. In their words:

> metaphors can shape our responses and guide our thinking, because of the particular kind of perspective taking they involve, without conveying information in the usual sense. This, more than anything else, is why we think metaphors must be explained in a distinctive way. (2015: 169)

Meaning or information is, for them, public content that underwriters interlocutors' joint inquiry into how things are. As there is "little evidence that metaphor ever contributes information in the sense of publicly accessible content that supports inquiry" (Lepore and Stone 2015: 170), they conclude metaphors do not contribute propositional content even if the speaker intended the hearer to focus on a series of similarities. These similarities are not part of any propositional content communicated by the speaker. Rather, they are part of the interpretive effects delivered by propositional content which is in no way special or metaphorical but literal. The literal meaning of the sentence uttered metaphorically prompts the similarity but does not underwrite interlocutors' joint inquiry into how things are. The hearer is invited to explore the implications of seeing one thing as another by using his imagination.

To support the plausibility of this view of metaphor they resort to a parallelism between metaphors and jokes. Speaker's goal to utter both jokes and metaphors is not to assert any propositional content it is rather to show some imagery. The comedian's utterance provides us with a humorous imagery but does not deliver any information that contributes to inquiry. Correspondingly, the metaphorist's utterance prompts the hearer's perception of a similarity but he is not informed of anything to update the conversational record.

These and other cases of "imagination" are taken as evidence that there is no room for the sort of communication based on reasoning and intention recognition. They reject the Gricean intentional framework to characterize interpretation as they think (i) that interpretive reasoning is more diverse and (ii) that linguistic meaning is broader in scope than Grice envisioned (Lepore and Stone 2015: 6). In their view, any traditional pragmatic account involves a reduction of meaning to communicative intention. As this is incompatible with (i) and (ii), they reject any pragmatic account of metaphor. They admit that intention recognition is necessary for all collaboration but

> interlocutors' contributions to conversation—just like their contributions to practical activity—carry things forward according to a dynamic that is antecedent to their intentions, and independent of them. (Lepore and Stone 2015: 6)

They consider it a mistake, therefore, that "the theory of CIs [conversational implicatures] eliminates the need to describe linguistic conventions and imaginative mechanisms in detail." (2015: 6). The inferential mechanisms recruited in language interpretation are quite diverse and eclectic and "it follows that overarching frameworks like the Cooperative Principle or the Principle of Relevance can't be the whole story." (2015: 83). A unitary account of all of the inferential mechanisms is not possible. They also claim that "the category of conversational implicature does no theoretical work. Pragmatics can be, at most, a theory of disambiguation; pragmatic reasoning never contributes content to utterances." (2015: 83). Consequently, they have to reject any pragmatic account of metaphor interpretation. The explanatory resources of traditional pragmatic theories cannot explain the alleged metaphorical meaning although it goes beyond the conventional rules of language.

If any pragmatic reasoning is ruled out, they must find "a way of thinking about the inquiry that interlocutors pursue in conversation as completely governed by linguistic rules." (2015: 6–7). Nevertheless, this is not bad news for them. They worry about the notion of *meaning* if we are too permissive with its use: "If we can locate metaphor elsewhere, it is good news for meaning." (2010: 179).

4. *A Defense of Pragmatically Derived Metaphorical Meanings*

Even if we agree on widening the scope of linguistic conventions and advocate for the distinctive character of novel metaphor interpretation, we think Lepore and Stone (2015) do not succeed in locating metaphor out of speaker meaning. In what follows, we give a pragmatic account for metaphor starting in 4.1 by giving the identification conditions of metaphorical utterances. These conditions block the literal interpretation of metaphorical utterances and show that metaphorical utterances go beyond the conventional rules. Thus they are an argument against Stone and Lepore's claim that our insights into metaphorical thinking are prompted just by the literal meanings of utterances.

We are aware that it is possible to accept our identification conditions and nevertheless hold that there is no metaphorical meaning. These conditions would merely prompt metaphorical thinking. Nevertheless, in our opinion, they can trigger the search for inferential information to solve the communication problem that arises from blocking the literal interpretation. Linguists often present examples that provide evidence that the inferential system does indeed generate content beyond the encoded meaning, and that its operation is grammatically constrained (Vicente 2010). Since semantic composition rules

drive the composition process, going beyond these rules triggers the search of content that eventually allows the composition, and it is the job of pragmatics to supply the specific conceptual addition or modulation. Modulation as required for metaphor will be considered in 4.2. In this section we expound how metaphorical thinking is a part of the interpretive context from which metaphorical provisional meanings are fixed. They are constructed by means of one type of the pragmatic sub-tasks typically involved in ad hoc concept construction: language-shifting.[6] The language–shift involved in interpreting a metaphorical utterance constitutes a new context of interpretation that at least delivers metaphorical meanings (ad hoc concepts) for the expressions used abnormally (non-conventionally) in the metaphorical utterance. Metaphorical meanings cannot exist without the metaphorical thinking from which they are conceived and thus the power of metaphor is not lost in their elaboration. Taking into account these metaphorical meanings, metaphorical utterances can convey propositional contents to agree or disagree, they can provide information that interlocutors add to the record.

4.1. *Metaphorical Identification*

As we have claimed in our first point of disagreement, we do not share the proposal that the analogical reasoning is triggered from the inadequacy of the literal meaning of the sentence as a whole when it is used metaphorically. For Lepore and Stone, the sentence used metaphorically has a literal meaning which is inadequate. For us, there is no need to get to that absurd interpretation because, among other things, many metaphorical utterances cannot be literally grasped.

According to Lepore, at least as it is explained in Cappelen and Lepore,

> any utterance succeeds in expressing an indefinite number of propositions. One of these, *the proposition semantically expressed, is easy to grasp*. (2005: 206)

But, in our opinion, there is no proposition easy to grasp at least in the first part of an utterance of (1), (1a)

(1a) Love is a snowmobile racing across the tundra.

For a proposition to be graspable it must be intelligible. (1a) does not have the minimal level of meaningfulness to be propositional. Sub-propositional expressions in (1a), for example 'a snowmobile racing across the tundra', give access to semantic information such as

[6] Different types of pragmatic processes are accepted in current pragmatic theories. For example, Recanati (2004) accepts, in addition to context-shifting, pragmatic processes of enrichment, loosening and transfer in what is said.

> A snow mobile is a sport vehicle like a car on skis that you drive through the snow.
> It's fun; it's exhilarating, and it gives a sense of adventure. A tundra is a frozen landscape with no trees, a place of relative safety. (Lepore and Stone 2014: 75)

Even the sub-propositional expression 'the ice weasels' in (1b), another sentence of (1),

(1b) At night the ice weasels come.

gives access to the semantic information also indicated by these authors: "Weasels are small predatory animals known for their fierceness and trickery." (2014: 75). But no propositional information is achieved from the literal interpretation of (1a). We then strongly disagree with what Lepore and Stone add:

> When you put this all together you imagine a prototypical course for a *love affair*, where it starts with a sense of adventure and excitement and then goes horribly wrong leaving you with a gnawing feelings of torture and pain. *What seems to be doing the work here is our ability to understand the sentence as described*; and then to draw an analogy between the experience of being in love and a certain kind of history that could happen. (2014: 75, our emphasis)

What is doing the work here is by no means our ability to understand the sentences included in (1) as described! What is literally grasped in (1a) is the meaning of the sub-sentential and linguistically meaningful complex expression 'a snowmobile racing across the tundra'. Even the sentences that follow (1a) in (1) are also literally grasped and could be compositionally added to the meaning of this complex NP. However, (1a) is semantically or literally unintelligible.

What we claim is that in the interpretation of sentences such as (1a) semantic composition is not available. Although (1a) is syntactically well-formed it lacks semantic coordination. There is a semantic mismatch since our semantic knowledge of the words 'snowmobile' and 'love' tells us that the composition of meaning is mediated by a semantic restriction in the word 'snowmobile' to the effect that its encoded concept (that denotes concrete objects to go on snow) cannot be normally used to talk about a feeling. The predicate 'is a snowmobile racing across the tundra' cannot make its semantic contribution to the clause since its meaning typically needs a concrete entity to fill in the semantic role of its subject and there is no element in (1a) with the feature [+concrete] to take such a role. As Asher says in relation to similar examples of semantic mismatch

> These predications should be precluded by normal type constraints, and we know that there are no reference preserving maps (in this world) from goats to talking agents or from trees to talking agents. (Asher 2011: 284)

As the predication in (1a) is precluded by normal type constraints, no resulting meaning is available to obtain an acceptable literal proposition. Besides, if the proposition semantically expressed by (1a) has to be easy to grasp, there is no proposition semantically expressed by (1a).

We must recognize that (1a) is a metaphor but the literal meaning of the sentence uttered cannot be the way of recognizing that (1a) is a metaphor since it cannot be obtained from the linguistic meaning of the sentence. What we know, as competent speakers, is that the term 'snowmobile' is used in an abnormal way in (1a), it could be taken as a category mistake or a non-conventional use of the linguistic expression. We call this "contextual anomaly" (Romero and Soria 1997/98) or "contextual abnormality" (Romero and Soria 2005a, 2005b and 2007) and we define it as the use of an expression in an unusual linguistic or extra-linguistic context. This means that there are two different modes of appearance of it:

(a) As an oddity between the terms uttered (as in the previous examples of novel metaphors: (1) and (5)–(7)).
(b) As an oddity between the occurrence of an expression in the actual unusual context and the implicit context associated to a normal use of this expression.

Abnormality of mode (b) is due to the incompatibility of the context of the utterance and the context conventionally associated to the uttered expressions. Let's change the context in (4)

(4) [Sarah asks Marian where her pet is and she answers:] My cat is on the mat.

to obtain (4').

(4') [Marian is reading at home, and her one-year-old son is playing on a mat with something he found on the floor. Sarah, a friend who knows the kid is a real mummy's boy enters the room and asks her where her son is. Marian answers:] My cat is on the mat.

Now there is a tension between 'cat' and the context. (4') is an utterance that concerns a boy rather than a cat. This tension blocks the literal interpretation. The composition between CAT and THE SPEAKER depends on any salient relation in the context. The only salient relation is the relation of motherhood that the speaker has with her one-year-old son. If the tension mentioned did not block the literal composition, we would get to an unintelligible result: THE CAT THE SPEAKER IS MOTHER OF. Now the proposition semantically expressed is not easy to grasp because (4') is not an utterance that concerns a cat. In our approach, our linguistic knowledge just guides us to semantic information but this information does not always have to lead by itself to a proposition semantically expressed by a sentence or its utterance.

Contextual abnormality is a necessary but not a sufficient condition for metaphor. It is also present in metonymical utterances such as (8)

(8) [In a restaurant, a waiter asks a waitress what to do next and she answers:] The ham sandwich is waiting for his check.

An additional identification criterion for metaphor is needed: a conceptual contrast. Conceptual contrast is the recognition that the speaker

is talking about a topic (target domain) using terms which normally describe another (source domain). This conceptual contrast occurs in (1a) but not in (8). In (1a) we detect as target domain the concept LOVE and as source domain the concept A SNOWMOBILE RACING ACROSS THE TUNDRA WHICH FLIPS OVER, PINS YOU UNDERNEATH AND LEAVES YOU EXPOSED TO ICE WEASELS AT NIGHT. In (4′), we detect as target domain the concept INFANT and as source domain the concept CAT.

Metaphorical identification achieved by both contextual abnormality and conceptual contrast blocks the literal interpretation to avoid a route that leads to no propositional content semantically expressed and triggers the metaphorical mechanism. Although a demand for pragmatic information is triggered by normal type constraints, the resolution of this situation cannot be treated as part of semantics. Nevertheless, this does not mean directly that no proposition can be conveyed by a metaphor as Lepore and Stone have argued (see section 3) and thus it is relevant to wonder if, by means of some conceptual adjustment in the interpretation of metaphorical utterances, some proposition can be expressed. Our positive answer depends on taking as input the evidence given by the speaker with her metaphorical utterance, which includes among other things the semantic information of sub-sentential expressions, to construct the analogical correspondences that are relevant to get as output the intended proposition. This proposition includes a metaphorical provisional meaning.

4.2. *A pragmatic account of metaphorical meanings*

If we accept metaphorical meanings, we have to take into account several features in our proposal of speaker meaning. First, we cannot forget that speaker meaning is an occasion-meaning that, according to Grice at least, fixes the occasion-meaning of the utterance and this may be non-conventional, as we argue metaphorical meaning is. Second, non-conventional meanings can be considered ad hoc concepts. These are ubiquitous in human cognition and can be constructed literally and non-literally. Third, pragmatic theories admit different types of pragmatic processes in deriving the proposition expressed and one of them, the language-shift, permits us to get the metaphorical meaning without losing the cognitive power of metaphor.

As Grice (1957/89, 1968/89 and 1969/89) shows, speaker meaning is an occasion-meaning. In addition, it is the basic semiotic concept, the concept from which the semiotic concept of the occasion-meaning of the utterance-type can be defined. This occasion-meaning of the utterance-type can coincide or not with one of its timeless meanings. If we consider Grice's example,

(9) If I shall then be helping the grass to grow, I shall have no time for reading.

we can point out that the timeless meaning of this "complete" utterance-type could be specified as 'If I shall then be assisting the kind of

thing of which lawns are composed to mature, I shall have no time for reading' or as 'If I shall then be assisting the marijuana to mature, I shall have no time for reading'. The word 'grass' means 'lawn-material' or 'marijuana'. But the speaker does not always mean by (9) one of these two timeless meanings. He could use (9) to mean 'If I am then dead, I shall not know what is going on in the world' and, in addition, 'one advantage of being dead will be that I shall be protected from the horrors of the world'. Nevertheless, the words 'I shall be helping the grass to grow' neither mean nor mean *here* 'I shall be dead'. In this case the utterer's occasion-meaning fixes an occasion-meaning of an utterance-type that does not coincide with one of the two timeless meaning of the words used (Grice 1969/89: 89–90). This means that occasion-meaning does not have to be a timeless or a conventional meaning. Arbitrary conventions are very useful in communication, especially in verbal communication, but reducing verbal communication to the linguistic meaning or timeless meaning is simply unacceptable both on empirical and theoretical grounds. Since Grice's use of the terms "occasion-meaning" or "speaker meaning", which by now have become conventional, we can no longer claim that conventional meaning is the only kind of meaning we can speak of.

The nature of metaphorical meaning can be accounted for if, in the theory of meaning, it is accepted that there are non-conventional uses of words which may acquire provisional non-conventional meanings. In the case of metaphorical utterances, words mean what they conventionally mean but speakers by using them convey propositions that involve new meanings for these words that are active only on the occasion of the utterance; they are merely provisional and explainable as the result of a productive and systematic mechanism: analogical reasoning. Although it is usual to understand, following Grice (1975/89), that these new metaphorical meanings form a part of conversational implicatures, not even Grice has always kept that claim. If we attend to his distinction between formality and dictiveness (Grice 1987/89: 361–362), we can recognize that there are cases of dictiveness without formality. Formality lets us know when the evaluated content belongs to the conventional part of the meaning of an expression while dictiveness discerns when a part of the content belongs to what is said. Grice exemplifies this with two examples, one of them being a metaphor.

> Suppose someone, in a suitable context, says 'Heigh Ho'. It is possible that he might thereby mean something like "Well that's the way the world goes". Or again, if someone were to say "He's just an evangelist", he might mean, perhaps, "He is a sanctimonious, hypocritical, racist, reactionary, money-grubber". If in each case his meaning were as suggested, it might well be claimed that what he meant was in fact what his words said; in which case his words would be dictive but their dictive content would be nonformal and not part of the conventional meaning of the words used. We should thus find dictiveness without formality. (Grice 1987/89: 361)

If someone utters 'He's an evangelist', what she means depends on what the words help us to say, although this does not coincide with its linguistic meaning, or part of its linguistic meaning. "Evangelist" does not conventionally mean 'sanctimonious, hypocritical, racist, reactionary, money-grubber' and in using this word, the speaker changes its meaning provisionally.

A provisional meaning or an ad hoc concept is, according to Barsalou (1983), one that is made up on the spot for a particular purpose and it contrasts with a ready-made concept. Ready-made concepts are associated with familiar words and well-established categories. They produce organized and easily recoverable knowledge which resides in long term memory. But not all the concepts that a speaker can represent mentally are ready-made and it is plausible to suppose that entities can be categorized differently to achieve some relevant aim. Our ready-made categorizations and conceptualizations can give way to ad hoc categorizations and conceptualizations (Barsalou 1983 and 1985). In fact, ad hoc concepts are ubiquitous in human cognition. They are non-lexicalized concepts and include knowledge associated to categories created for the occasion of the utterance, which does not reside in long term memory.

There are some varieties of ad hoc concepts and they can be represented linguistically either literally where the words of a complex expression interact by compositional semantics or non-literally by means of a pragmatic adjustment of a concept (expressed by a lexicalized expression or by a complex expression). The ad hoc concept GOOD THINGS TO STAND ON TO CHANGE A LIGHT BULB may be expressed literally by means of a complex phrase 'good things to stand on to change a light bulb'. In the same way, complex expressions such as 'a snowmobile racing across the tundra and then suddenly it flips over, pinning you underneath. At night the ice weasels come' express a literal ad hoc concept: A SNOWMOBILE RACING ACROSS THE TUNDRA WHICH FLIPS OVER, PINS YOU UNDERNEATH AND LEAVES YOU EXPOSED TO ICE WEASELS AT NIGHT. These concepts are "literal" in the sense that the words representing them keep their encoded meaning in the composition of the meaning of the phrase or extended text. But not all ad hoc concepts are literal, many of them can be constructed non-literally by a pragmatic adjustment of a literal concept (Romero and Soria 2010). The inputs of metaphorical adjustments are concepts in general (lexicalized or ad hoc). Sometimes, as in (4'), a lexicalized concept, CAT, is the point of departure for an ad hoc metaphorical concept or metaphorical provisional meaning but, other times, as in (1), a literal ad hoc complex concept is the point of departure. In each case, the metaphorical adjustment results in an ad hoc concept with a complete different denotation which is one of the constituents of the proposition expressed by the metaphorical utterance. But how is this meaning obtained?

We normally interpret utterances with respect to the context, *k,* in which they take place. But it is not always appropriate for us to inter-

pret them with respect to that context. On certain occasions we have to interpret them with respect to a context *k'* distinct from the context in which it is actually uttered. Context-shift, from *k* to *k',* is a pragmatic process and can be produced in several different ways, according to what aspect of context is shifted. If we represent a context, following Lewis (1980), as consisting of three parameters, a language, a situation, and a circumstance of evaluation, a context can be shifted by modifying one of these parameters. A context *k* is therefore analysed as a triple <*L*, *s*, *c*> where *L* is a language, *s* is a situation of utterance comprising a number of parameters corresponding to the situation of utterance (speaker, hearer, time, place, etc.), and *c* a circumstance of evaluation or a possible world.

There are examples of context-shifting which involve a situation-shift, a world-shift, or a language-shift, but, given the aims of this article, we will focus only on a case of language-shift such as (10)

(10) [It is mutually known to the speaker and his addressee that Paul is wrong about the use of 'paper session' that he understands with the meaning of 'poster session'. The speaker says:] Paul says he's due to present his work in the 'paper session'.

In (10), the context-shift can be described by a language-shift because the speaker of (10) does not use the expression 'paper session' in its normal sense but in the sense that expression has in Paul's idiolect, where it means the same as 'poster session' in its normal sense. Paul makes a deviant use of the phrase 'paper session'. The expression within the quotation marks, in this example, is not used with its standard meaning and so (10) has truth-conditions that differ completely from the truth-conditions of the utterance of the sentence when it does not include a quoted expression.

As we have argued, in metaphor there is also a language–shift (Romero and Soria 2007). This, in contrast to the one required for interpreting (10), is triggered by means of metaphorical identification and is guided by the metaphorical mechanism, by the development of an analogical thinking or, as we have understood it following Indurkhya (1986), by a coherent partial mapping of a set of features from source domain to target domain to obtain a metaphorically restructured target domain. Some properties of the source domain (only those relevant to get information for the characterization of the subject matter) are used as a source of information to describe the target. As Lepore and Stone claim "Metaphorical thinking often requires us to find many analogical correspondences simultaneously." (2015: 166).

Let's see an example of a mapping to interpret (4').The metaphorical mechanism links two separate cognitive domains in order to see one as the other: INFANT AS CAT. This link can be specified with a mapping, *M*, from the source domain, CAT, to the target domain, INFANT. A domain can be represented by both a set of terms which make up its vocabulary and a set of structural constraints which specify how these terms are

related to the information associated with the concept as we can see for CAT and for INFANT in Figure 2.

Figure 2

CAT	INFANT
$\boldsymbol{D_s} = \langle V_s, S_s \rangle$ $\boldsymbol{V_s}$ = {'cats', 'feline', 'walk', 'leg', 'pet', 'play', etc.} $\boldsymbol{S_s}$ = [1_s] Cats are small domesticated feline mammals, [2_s] Cats have soft fur, [3_s] The colour of the fur of each breed of cats varies greatly, [4_s] Cats play with anything available, [5_s] Cats are pets, [6_s] Pets need feeding and care, [7_s] Cats walk on four legs, [8_s] Cats scrutinise things carefully, [9_s] Cats are aloof with the unknown, [10_s] Cats are often used to catch mice, etc.	$\boldsymbol{D_t} = \langle V_t, S_t \rangle$ $\boldsymbol{V_t}$ = {'infant', 'human being', 'play', etc.} $\boldsymbol{S_t}$ = [1_t] Infants are young human beings, [2_t] Infants play with anything available, [3_t] Infants need feeding and care, [4_t] Infants are often unfriendly with the unknown, [5_t] At a certain stage, infants go on all fours, [6_t] Infants are at their early stage of their lives, etc.

The interpretation of (4′) entails finding a structural alignment between these domains. This alignment of consistent one-to-one correspondences allows the selection and partial mapping from a set of structural constraints from the source to the target domain. It is composed of a partial admissible function *F* from terms belonging to the source domain, arguments of the function, to terms that belong or will belong to the target domain. To continue with our example, *F* could be formed with the pairs: (cat → infant), (walk on four legs → go on all fours), (pet → infant), (aloof → unfriendly). The mapping is also composed of a subset of structural constraints of the source domain, *S*, which is coherently transformable by *F* to information associated with the target domain. In the example, *S* could be formed by structural constraints such as [4s], [6s], [7s], [8s], [9s] as we can see in the left column of Figure 3. These are transformable by *F* if each of its terms either belongs to the arguments of the admissible partial function *F* or belongs to the vocabulary of target domain directly. On transforming these structural constraints, we come across other structural constraints only in terms of target domain as we can see in the right column of Figure 3.

Figure 3

CAT S = [4_s] Cats play with anything available, [6_s] Pets need feeding and care, [7_s] Cats walk on four legs, [8_s] Cats scrutinise things carefully, [9_s] Cats are aloof with the unknown-.	INFANT^M or INFANT (AS CAT) [2_t^M] Infants play with anything available, (highlights 2_t) [3_t^M] Infants need feeding and care, (highlights 3_t) [4_t^M] Infants are unfriendly with the unknown(highlights 4_t) [5_t^M] Infants go on all fours, (highlights 5_t) [8_t^M] Infants scrutinise things carefully (new, coming from 8s)

If the union of the transformation of S with part of the information of the target domain is consistent, that is to say, if this union is true under at least one model, then the structural constraints of S have been coherently transformed by means of F in structural constraints of the target domain. Coherence is an inferential requirement for mappings: we can only transport the transformed information of the source domain that does not make our conception of target domain incoherent. The mapping M for (4') generates a metaphorically restructured conception of INFANT, INFANT^M or INFANT (AS CAT) characterized by the structural constraints in the right column of Figure 3. With [2_t^M], [3_t^M], [4_t^M], [5_t^M], nothing new is added to the target domain from the source domain but this target domain information is selected and reinforced by the relational similarities revealed by their alignment with the selected features in the source domain. As the information in [1_t], [6_t] (see the right column of Figure 2), is not selected, it is downplayed. In addition, when the description of the actual target domain is not aligned to the source domain in every feature, new information could be added to the target domain as it happens with [8_t^M]. Some similarities are created with metaphor.[7] In metaphorical interpretation, the target concept is described by means of some source concept-like relational properties which strengthen some properties that the concept already had and, in the more creative cases, by means of new properties that are added as a result of the analogical adjustment.

The mapping is shaped guided by the search of the properties that will enable the hearer to obtain the intended interpretive effects. The hearer does not explore the implications of seeing one thing as another just guided by his imagination. Rather, he is guided by his attempt to recognize what correspondences the speaker is intending him to make. This is not by magic, of course. The speaker must choose her words in such a way that her utterance makes her intention recognizable under the circumstances. The hearer conceptualizes the target as intended on

[7] Black (1954), Indurkhya (1986), Romero and Soria (1997/98) and Bowdle and Gentner (2005) have argued for this proposal.

the basis of the contextual evidence available, including the evidence provided by the speaker (both linguistic and non-linguistic), his previous world knowledge and personal experience, the specific situation, etc. Given the circumstances in (4'), Sarah can infer the analogical correspondences intended by Marian. The mapping is guided by inferential requirements (coherence, contextual relevance and intention recognition on the basis of the evidence provided) so that the target domain is provisionally restructured in the way the speaker envisioned and the parameter of language of the context of interpretation is shifted to *k'* as intended by the speaker. In *k'* new meanings are produced at least for the arguments of the function *F* where the metaphorical vehicle has to be included. INFANT (AS CAT) generates a new context of interpretation for (4') where the word 'cat' stands for an improvised meaning; this ad hoc meaning is sufficiently clear for the purposes of the conversation. The metaphorical provisional meaning of 'cat', the vehicle of metaphor, is the metaphorical provisional concept associated to 'infant' in that metaphorically restructured target domain, INFANT (AS CAT). The mapping approach allows us to know how certain words can change their meanings and what meanings they take on. This mechanism is used to produce non-conventional meanings in a systematic way, meanings or conceptions that are not lexicalized in the linguistic competence of the speakers of a linguistic community at that moment in the language. The cognitive power of metaphor often allows knowledge change (Gentner and Wolf 2000) and an increase of effability. We can entertain new concepts and express meanings not yet available in the system of the language and metaphorical utterances provide us with that expression.

By means of a language-shift, we can explain how the speaker can succeed in denoting when he uses the expression 'my cat' in (4') to denote the speaker's cat-like infant. The speaker in (4') can convey the information that her infant (conceptualized as a cat, for example, as an infant who is on all fours and scrutinizes things carefully) is on the mat by composing that content from, among other things, the sub-propositional metaphorical meaning. This meaning permits us to determine what is metaphorically said by (4'). She coordinates with the hearer about the information that her infant (conceptualized as a cat) is on the mat. Thus, the ad hoc concept INFANT (AS CAT) in *k'* can contribute information to the proposition that updates the conversational record. Sarah gets an answer to her question about her friend's son, she gets accessible information about how things are. The metaphorical conceptualization of one of the constituents of this proposition does not preclude the communication of a propositional content. What speaker and hearer cannot do is to coordinate on something her utterance does not concern, an imaginary cat and an imaginary mat, because when the hearer interprets an utterance such as (4'), both interlocutors know that the utterance concerns a real mat and a real boy.

This is very different from what happens in an utterance of a joke. Let us see a joke about cats.[8]

(11) A policeman in the big city stops a man in a car with a Siberian Lynx in the front seat. 'What are you doing with that Siberian Lynx?' He exclaimed, 'You should take it to the zoo.' The following week, the same policeman sees the same man with the cat again in the front seat, with both of them wearing sunglasses. The policeman pulls him over. 'I thought you were going to take that cat to the zoo!' The man replied, 'I did. We had such a good time we are going to the beach this weekend!'

If we compare the uses of the expression 'cat' in (4') and in the joke (11), we see that the joke does not concern a situation that includes any denotatum, it is a fictional situation with humorous intent. The hearer does not add to the record the information that a particular driver is going to the beach with his cat. The speaker is joking by exploiting the ambiguity of the linguistic expression 'taking someone to the zoo'. In jokes, we are presented with an idea or situation that is followed by an ambiguity or a twist, resolved in a clever way, simply to produce amusement while metaphorical utterances can and often do support inquiry. In (4'), we do have a particular boy denoted metaphorically by 'my cat' and a particular mat denoted literally by 'the mat'.

We would readily accept Lepore and Stone's (2010: 171) claim that in certain jokes, the sentences that the speaker uses mean that *p* but the point of the speaker is not to contribute information seriously to their interaction. However, to explain metaphor in this way is highly misleading. We do not think that the main point of the metaphorist is not to contribute information seriously to the conversation. Lepore and Stone can show a similarity between metaphor and jokes with (1) because this example is a humorous utterance by a comedian. In (1) the twist is presented and cleverly resolved by metaphorical propositions. Love is metaphorically described as a very positive and exciting experience (something that is widely accepted), however, it turns out to be a very negative and distressful one at the end of the sentence. In this case, the main point of the utterance is to produce amusement by the twist of the unexpected view of love in those negative terms. If taken as a joke, the speaker of (1) is not contributing the metaphorical proposition seriously. The joke is metaphorically told. If we did not get the mapping as intended, we would not get the metaphorical propositions and the humorous effect. We do not think, however, that metaphorical jokes should be the type of discourse selected if we want to find out whether the metaphorist can contribute information that supports inquiry. If that is the goal, it would be better to consider if the speaker can seriously use metaphorical utterances to answer questions, if she can agree or disagree about something that matters to the interlocu-

[8] Example taken from a google search: http://www.jokes4us.com/animaljokes/catjokes.html.

tors, since these are examples of utterances that typically support inquiry.[9] With metaphorical utterances interlocutors can seriously answer questions as in (4'). With a joke, by contrast, interlocutors do not identify questions that matter to them and they do not work to reach agreement about the answers.

Metaphorical utterances provide us with a content that can be negated while it would be absurd to negate a joke. Lepore and Stone's view of metaphor as just showing similarities is of no use to explain the difference between the interpretive effects of a metaphorical utterance and its negation. In metaphor, content is provided in an unusual way. But this is by no means a reason to reject that it is publicly accessible. Interpreting Donne's metaphor (12)

(12) No man is an island entire of itself; every man is a piece of the continent, a part of the main.

involves triggering new interpretations for 'island' and 'continent'. This does not miss the breadth and interconnections of Donne's imagery because the new meanings for 'island' and 'continent' depend on the metaphorical concept PEOPLE AS GEOGRAPHICAL PLACES and more particularly on HUMAN CONNECTIONS AS GEOGRAPHICAL CONNECTIONS, metaphorical concepts that act as part of the metaphorical context for the interpretation of (12). Without this new context of interpretation, characterized by a new conception of the target domain, it is not possible to conceive the metaphorical meanings for 'island' and 'continent'. The propositions *that no man is an island* and *that every man is a piece of the continent* are only conceivable from this new context. These propositions have truth-conditions. They indicate what the world must be like for the metaphorical propositions conveyed by (12) to be true. The former is true if and only if no object within the extension of the concept MAN belongs to the extension of the concept MAN (AS ISLAND). The latter is true if and only if every object within the extension of the concept MAN, also belongs to the extension of the concept MAN (AS CONTINENT). And speakers

[9] The role of metaphor in science, a clear case of speakers supporting inquiry, is well known to metaphor theorists. Gentner and Wolf (2000) explain how in the history of the scientific account of the atom, several metaphorical models have restructured the concept of atom: Wilson's plum-pudding model and Nagaoka's Saturnian-disk model guided Rutherford's research and finally led to incorporation of orbits as in the Rutherford-Bohr's solar-system model of the atom. Metaphorical content is publicly available, it is communicated and has an important role in knowledge change. This conceptualization of the atom allows utterances such as 'electrons travel in circular orbits around the nucleus' which include terms ('travel', 'orbit') of the source domain SOLAR SYSTEM with a transferred meaning and which makes it possible to "seriously" communicate the metaphorical proposition intended by the scientist. At Bohr's time, when the term 'orbit' was used to talk about the earth, it did not mean the same as when it was used to talk about the atom. The term 'orbit' contributes to the proposition with the metaphorical provisional meaning that was constructed for the utterance in the target domain ATOM (AS SOLAR SYSTEM). In this way, the speaker can convey propositions about the structure of the atom conceptualized as a solar system.

can agree or disagree with respect to these metaphorical propositions directly communicated by (12) as we can see in (13), an excerpt from the very beginning of the film *About a boy*, where we find the opinion of Will, the main character, expressed in the following way:

(13) If I may say so, [*that no man is an island* is] a complete load of bollocks.

Will disagrees with Donne's claim in (12). He denies that no man is an island, he denies that no object within the extension of the concept MAN belongs to the extension of the concept MAN (AS ISLAND). If the purpose of metaphorical utterance were, as Lepore and Stone claim, just to see the analogies created by PEOPLE AS GEOGRAPHICAL PLACES, how can the interpretive effects of the similarity be negated? We do not know how negative metaphors such as *no man is an island* can be explained from their account as it is quite absurd to negate that there is not a similarity that you have already entertained. With Lepore and Stone's view of metaphor as an invitation to see one thing as another and nothing more, it is not possible to explain the difference between what Donne and Will say respectively.[10]

With (12) and (13), the same metaphorical concept MAN (AS ISLAND) is prompted. The difference of their interpretive effects is merely that one asserts what the other denies. From our standpoint, we can say that we can interpret both (12) and (13) from the same metaphorical context *k*' and thus their speakers "metaphorically say" the opposite.

The opinion of Will is not exhausted by (13) and he adds (14).

(14) a. In my opinion, all men are islands.
b. And what's more, now's the time to be **one**.
This is an **island** age.
A hundred years ago, you had to depend on other people.
No one had TV or CDs or DVDs or videos...
...or home espresso makers.
As a matter of fact, they didn't have anything cool.
Whereas now, you see...
...you can make yourself **a little island paradise**.
With the right supplies and the right attitude...
...you can be **sun-drenched, tropical, a magnet...**
...for young Swedish tourists.
(…)
And I like to think that perhaps I am **that kind of island**.
I like to think I'm pretty cool.
I like to think I'm **Ibiza**.

He starts with (14a), a metaphorical utterance that expresses the proposition *that all men are islands*. This is true if and only if every object

[10] We are aware that, being a case of fiction, Will is not a real speaker, and the speaker meaning is really the meaning intended by the script author. However, we would like to use the example as if it were a real speaker. After all, this is what we usually do when we exemplify a proposal on speaker meaning.

within the extension of the concept MAN also belongs to the extension of MAN (AS ISLAND). He does not only deny (12) with (14a) but he also goes further saying that all men are islands. In addition, to interpret (14b), where the expressions in bold letters represent terms that come from the vocabulary of the source domain, involves a progressive alignment to elaborate a new mapping that results in the metaphorically restructured concept WILL (AS THAT KIND OF ISLAND). The source domain has become a more specific ad hoc concept AN ISLAND THAT IS SUN-DRENCHED, TROPICAL, A MAGNET…FOR YOUNG SWEDISH TOURISTS and the target is a specific man, WILL. This idea of progressive alignment of more specific aspects in the domains allows us to provide an explanation of extended novel metaphors that other pragmatic approaches cannot provide. In addition, it permits us to argue that the similarity the metaphorist invites us to construct depends on the recognition of his intention to get across the intended effects from a certain set of specific analogical correlations within the domains. He acts in a way that helps the audience to know what analogous properties of ISLAND will serve to describe a specific MAN.

The script writer of the film presents the fiction as an example of how a man can be a metaphorical island. The whole film is a metaphorical argument to show disagreement with Donne's metaphorical proposition that no man is an island. By showing Will's life, the script writer is providing an example of a MAN (AS ISLAND) with its peculiar characteristic as an individual, WILL (AS THAT KIND OF ISLAND). Interlocutors in the film identify a metaphorical question that matters to them and work to reach agreement about the answers.

When later on in the film, someone tells Will 'you can't shut life out. No man is an island', another friend exclaims: 'She's right, you know. Yeah, she is', and Will replies:

(15) No, she's not! She's wrong!
Some men are islands. I'm a bloody island!
I'm bloody Ibiza!

Will is insisting on his disagreement with a metaphorical proposition rather than making us perceive a similarity that we already had. Metaphorical propositions contribute to the conversational record. They contribute information in the sense of public content that underwriters interlocutors' joint inquiry into how things are.

When Lepore and Stone metaphorically defend that metaphor should not be on the scoreboard we understand what they mean and we disagree. In their book, the concept to build their central claim about what constitutes propositional content is metaphorically constructed. By doing this they have constructed a view of what is communicated as what can be registered in a (baseball) scoreboard and this metaphorical conceptualization of the conversational record can be part of the propositional contents they convey in their subsequent inquiry about communication. We understand their proposal, among other things,

from the analogical construction of the ad hoc concept, CONVERSATIONAL RECORD (AS A BASEBALL SCOREBOARD), in the way intended by them. The metaphorical concept is a part of the propositions conveyed. Metaphorical communication is successful only if we map the right features from source to target as intended by the speakers. In their works, 'scoreboard' is used with a metaphorical provisional meaning. Meaning that is possible in the new shifted context that its mapping provides and that can contribute to the propositions intended by the speakers (philosophers in search of truth rather than poets or comedians).

In this way, our position eliminates the difficulties that, according to Lepore and Stone, any non-sceptic account on metaphorical meaning share. In addition, our position provides both an effective proposal on metaphorical identification that appeals to sub-propositional conditions and a proposal about how to obtain metaphorical propositions that, without losing the power of metaphor, permits us to explain the speakers' agreements or disagreements when metaphorical utterances are involved. If we have a systematic metaphorical mechanism to combine encoded meaning and contextual information, our communication system is more powerful (increases the scope of generativity) and plausible (in terms of language acquisition and language change). Thus, we do not see in what sense, locating metaphor elsewhere can be "good news for meaning". As long as we stand clear about encoded meaning, linguistic composition rules and the linguistic constraints in relation to the demands of contextual information, we see no offense to semantics in talking about metaphorical propositional contents as part of speaker meaning (conceived as a combination of linguistically encoded, semantically constrained and contextually inferred information). As we acknowledge the important role of semantic constraints on the demand of contextual information marked in our metaphorical identification, our view coincides with Lepore and Stone' defense of a wider effect of linguistic rules on speaker meaning but this does not lead us to a sceptic approach to metaphorical meaning.

5. *Conclusion*

In this paper we have challenged Lepore and Stone's sceptic account on metaphor. We argue for metaphorical meaning as a variety of provisional meanings (ad hoc concepts) constructed by an analogical mapping and triggered by sub-propositional identification conditions.

By rejecting their claim that the metaphorical use of language is related, as any other, just to the literal meaning, our position provides an explanatory proposal on metaphorical identification which appeals to sub-propositional conditions: contextual abnormality and conceptual contrast. Utterances that are identified as metaphorical block the literal meaning of the sentence and trigger a pragmatic mechanism used to generate new conceptions of concepts (target concepts) resorting to conceptual mappings. The production and interpretation of metaphorical

utterances is not exhausted simply by entertaining the metaphorical concept (seeing one concept as another). These new conceptions change the interpretation context of the utterance in the language parameter. From this new interpretation context, provisional metaphorical meanings are determined for words, phrases or even extended discourses used as metaphorical vehicles in the utterance: they are the meaning of their counterparts in the metaphorically restructured target domain. These meanings, since they are determined in the new context of interpretation generated by the analogical thinking, are constituents of what is explicitly communicated to get the metaphorical propositional content communicated by the metaphor. What is asserted by the speaker is a propositional content in a new context. Thus the power of metaphor, its capacity to prompt the hearer to see one thing as another, is not lost.

In addition, our position provides a pragmatic account of how to obtain metaphorical propositions that permits us to explain the speakers' agreements or disagreements when metaphorical utterances are involved. In doing this, a theory of meaning needs to take into account something more than simply conventional meaning.

References

Asher, N. 2011. *Lexical Meaning in Context: A Web of Words*. Cambridge: Cambridge University Press.

Barsalou, L. W. 1983. "*Ad hoc* Categories." *Memories and Cognition* 11: 211–227.

Barsalou, L. W. 1985. "Ideals, Central Tendency, and Frequency of Instantiation as Determinants of Graded Structure in Categories." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11 (4): 629–654.

Black, M. 1954. "Metaphor." *Proceedings of the Aristotelian Society* 55: 273–294.

Bowdle, B. and D. Gentner 2005. "The Career of Metaphor." *Psychological Review*, 112: 193–216.

Cappelen, H. and E. Lepore 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.

Davidson, D. 1978. "What metaphors mean." *Critical Inquiry*, 5: 31–47.

Emerson, R.W. 1836. *Nature*. Boston: James Munroe.

Gentner, D. and P. Wolf 2000. "Metaphor and Knowledge Change." In E. Dietrich and A Markman (eds.). *Cognitive Dynamics: Conceptual Change in Humans and Machines*. Mahwah: Laurence Erlbaum Associates: 295–342.

Grice, H. P. 1957/89. "Meaning." In H. P. Grice 1989: 213–223.

Grice, H. P. 1968/89. "Utterer's Meaning, Sentence–Meaning, and Word–meaning." In H. P. Grice 1989: 117–37.

Grice, H. P. 1969/89. "Utterer's Meaning and Intention." In H. P. Grice 1989: 86–116.

Grice, H.P. 1975/89. "Logic and Conversation." In H. P. Grice 1989: 22–40.

Grice, H.P. 1987/89. "Retrospective Epilogue." In H. P. Grice 1989: 283–303.

Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.

Indurkhya, B. 1986. "Constrained Semantic Transference: A Formal Theory of Metaphors." *Synthese* 68 (3): 515–551.

Kittay, E. F. 1987. *Metaphor. Its Cognitive Force and Linguistic Structure*. Oxford: Clarendon Press.

Lakoff, G. 1993. "The Contemporary Theory of Metaphor." In A. Ortony (ed.), *Metaphor and Thought*, 2nd edition. Cambridge: Cambridge University Press: 202–255.

Lakoff, G. and M. Johnson 1980, *Metaphors We Live By*. Chicago: University of Chicago Press.

Lepore, E. and M. Stone 2010. "Against Metaphorical Meaning." *Topoi* 29 (2): 165–80.

Lepore, E. and M. Stone 2014. "Philosophical Investigations into Figurative Speech Metaphor and Irony." *ProtoSociology: An International Journal of Interdisciplinary Research* 31: 75–87.

Lepore, E. and M. Stone 2015. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford: Oxford University Press.

Lewis, D. 1969. *Convention: A Philosophical Study*. Oxford: Basil Blackwell.

Lewis, D. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8 (1): 339–359.

Lewis, D. 1980. "Index, Context, and Content." In S. Kanger and S. Öhman (eds.). *Philosophy and Grammar*. Dordrecht: Reidel: 79–100.

Recanati, F. 2004. *Literal Meaning* Cambridge: Cambridge University Press.

Richards, I. A. 1936. *The Philosophy of Rhetoric*. London: Oxford University Press.

Romero, E. and B. Soria 1997/98. "Stylistic Analysis and Novel Metaphor." *Pragmalingüística* 5–6: 373–389.

Romero, E. and B. Soria 1998. "Convention, Metaphor and Discourse." *Atlantis* 20 (1): 145–59.

Romero, E. and B. Soria 2005a. "Metaphorical Concepts and Language." In J. J. Acero and P. Leonardi (eds.). *Facets of Concepts*. Padova: Il Poligrafo: 185–208.

Romero, E. and B. Soria 2005b. "Cognitive Metaphor Theory Revisited." *Journal of Literary Semantics* 34 (1): 1–20.

Romero, E. and B. Soria, B. 2007. "A View of Novel Metaphor in the Light of Recanati's Proposals." In M. J. Frápolli (ed.). *Saying, Meaning and Referring. Essays on François Recanati's Philosophy of Language*. Basignstoke: Palgrave Macmillan: 145–59.

Romero, E. and B. Soria 2010. "Phrasal Pragmatics in Robyn Carston's Programme." In B. Soria and E. Romero (eds.). *Explicit Communication: Essays on Robyn Carston's Pragmatics*. Basignstoke: Palgrave Macmillan: 183–98.

Romero, E. and B. Soria 2013. "Optionality in Truth-Conditional Pragmatics." *Teorema: Revista Internacional de Filosofía* 32 (2): 157–74.

Romero, E. and B. Soria 2014. "Relevance Theory and Metaphor." *Linguagem em (Dis)curso* 14 (3): 489–509.

Sperber, D. and D. Wilson 1986/95. *Relevance: Communication and Cognition.* Oxford: Blackwell.

Sperber, D. and D. Wilson 2008. "A Deflationary Account of Metaphor." In R. W. Gibbs (ed.). *The Cambridge Handbook of Metaphor and Thought.* New York: Cambridge University Press: 84–108.

Stanley, J. 2005. "Semantics in Context." In G. Preyer and G. Peters (eds.). *Contextualism in Philosophy: Knowledge, Meaning, and Truth.* Oxford: Oxford University Press: 221–253.

Stern, J. J. 2000. *Metaphor in Context.* Cambridge, Mass: MIT Press.

Vicente, B. 2010. "The Role of Pragmatic Inferencing in Compositional Semantics." In B. Soria and E. Romero (eds.). *Explicit Communication: Essays on Robyn Carston's Pragmatics.* Basignstoke: Palgrave Macmillan: 58–74.

Weiltz, P. and Weiltz, Ch. (Directors) 2002. *About a Boy* [DVD]. Available in http://www.about-a-boy.com/

Croatian Journal of Philosophy
Vol. XVI, No. 47, 2016

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 Jutronić, 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 ϕing by ψing just in case it is common knowledge among members of G that they truly believe that they regularly ϕ by ψing, that they generally prefer to ϕ by ψing (as opposed to ϕing 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 ϕing 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 amount 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?
⇝ Serve me the french toast.

(5) Oil prices doubled and demand for consumer goods plunged.
⇝ The doubling preceded/caused the plunge.

(6) Well, it looked red. [said of a magician's handkerchief]
⇝ 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 ⇝ 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 thusand-such—or, more broadly, that they have performed an indirect communicative act of thus-and-such kind[8]—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 it if 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 extralinguisic 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 particularlized 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 thereore 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 of metaphors can be *mis*interpreted.[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) is 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 ceasefire, 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 ϕing. 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 ϕing in unconventional ways.
(ii) Somehow, out of the chaos, a pattern of ϕing by ψing emerges, at first for arbitrary or idiosyncratic reasons.
(iii) The pattern becomes increasingly standardized and self-reinforcing.
(iv) Finally, a convention of ϕing by ψing has come into being.

The problem for Lepore and Stone, is that this abstract story entails that for any convention of ϕing by ψing to develop in the first place, it has to be possible to perform at least some rudimentary acts of ϕing 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 Reimer, 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 aphonic 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.

Croatian Journal of Philosophy
Vol. XVI, No. 47, 2016

Cooperation with Multiple Audiences

MARILYNN JOHNSON
Department of Philosophy,
The Graduate Center, City University of New York, USA

Steven Pinker proposes a game-theoretic framework to help explain the use of veiled speech in contexts where the ultimate aims of the speaker and hearer may diverge—such as cases of bribing a police officer to get out of a ticket and paying a maître d' to get a table. This is presented as a response to what Pinker sees as the failure in H. P. Grice's influential theory of meaning to recognize that speakers and hearers are not always cooperating. In this paper I argue that Pinker mischaracterizes Grice's views on cooperation, and use this to refine a positive picture of what sort of cooperation is demanded by Grice's Cooperative Principle. This positive picture serves to insulate the Gricean framework from objectors—including Pinker—who overstate the obligations entailed by the adoption of the Cooperative Principle. I then argue that the cases Pinker presents are best treated by recognizing that in each instance the utterance is formulated with two intentions towards two different audiences and detail a resulting revision to Pinker's game-theoretic framework that reflects this proposal. I conclude by demonstrating how this proposed game-theoretic framework of cooperation with multiple audiences can be used to model the costs and benefits of other types of discourse, including political speech.

Keywords: Grice, cooperation, cooperative principle, multiple audiences, Pinker, game theory.

Introduction

Philosopher of language H. P. Grice is well known for his theory of speaker meaning, which is grounded in his Cooperative Principle and four maxims of conversation. Grice's work has influenced much current work in philosophy of language, linguistics, philosophy of law, evolutionary psychology, and many other areas. There is one area of Grice's program, however, that is frequently criticized: his Cooperative Principle. This principle is attacked by theorists across a wide range of

disciplines who point to the numerous ways in which certain communicative interactions are not cooperative[1] (Marmor 2011; Godfrey-Smith and Martínez 2013; Pinker 2007a; Pinker 2007b, 2011; Pinker, Nowak, and Lee 2008; Lee and Pinker 2010). However, as I will argue in this essay, Grice's Cooperative Principle is not as demanding as some have thought. Many interactions where the interlocutors have divergent ultimate aims are cooperative in the sense relevant to Grice's theory of meaning.

Perhaps the most robust and sustained objections to Grice on these grounds have been waged by psychologist Steven Pinker. In recent work, Steven Pinker (2007a, 2007b, 2011); Pinker, Nowak, and Lee 2008; Lee and Pinker 2010) has proposed a game-theoretic framework to help explain the use of implicature in contexts where the ultimate aims of the speaker and hearer may diverge. Pinker proposes that in such contexts implicature can be used as a way to avoid a number of social and financial costs by discussing examples—such as bribing a police officer to get out of a ticket and paying a maître d' to get a table—and provides a game-theoretic framework that is meant to model these costs and benefits. This is presented as a response to what Pinker sees as Grice's failure to recognize that speakers and hearers are not always cooperating.

In this paper I will argue that Pinker seriously mischaracterizes Grice's views on cooperation (see also Terkourafi 2011a; 2011b; Reboul Forthcoming a; Forthcoming b). In the course of doing so, I refine a positive notion of the sort of cooperation that is demanded by Grice's Cooperative Principle. Although I make Pinker my target, this positive notion could be used to dispel similar objections to Grice made by other theorists as well. I then argue that the cases Pinker presents are best treated by recognizing that in each instance the utterance was formulated with two different audiences in mind (See also Grice 1989: 37 and Neale 1999: 29-30 for discussion of multiple audiences): 1) the immediate audience, and 2) a future potential audience. The apparent obscurity arises because there is a clash between the intended communicated content with respect to these different audiences.

H. P. Grice

In works such as "Meaning", "Logic and Conversation", "Utterer's Meaning and Intention" (Grice 1989), Grice presents a theory of speak-

[1] Andrei Marmor writes "The standard model in the pragmatics literature focuses on ordinary conversations, in which the parties are presumed to engage in a cooperative exchange of information" (Marmor 2011: 83) but that "the enactment of a law is not a cooperative exchange of information" (Marmor 2011: 96). Peter Godfrey-Smith and Manolo Martínez write "Many theorists have seen communication as a fundamentally cooperative phenomenon. In an evolutionary context, however, cooperation cannot be taken for granted, because of problems of subversion and free-riding" (Godfrey-Smith and Martínez 2013: 1).

er meaning. Speaker meaning captures how a speaker can write something such as "Mr. X's command of English is excellent, and his attendance at tutorials has been regular" in a letter of recommendation and thereby mean "Mr. X is no good at philosophy" (Grice 1989: 33).

In this discussion Grice presents his Cooperative Principle. The Cooperative Principle states that conversational partners will make their conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange (29). This is a "quazi-contractual matter" (29). Further, Grice writes that "On the assumption that some such general principle as this is acceptable, one may perhaps distinguish four categories under one or another of which will, in general, yield results in accordance with the Cooperative Principle" (26). These are his four conversational maxims of quantity, quality, relation, and manner. The maxim of quantity states

> "1) Make your contribution as informative as is required (for the current purposes of the exchange) and
> 2) Do not make your contribution more informative than is required" (26).

The maxim of quality states

> "Supermaxim: try to make your contribution one that is true.
> 1) Do not say what you believe to be false, and
> 2) Do not say that for which you lack adequate evidence" (27).

The maxim of relation states

> "1) Be relevant" (27).

The maxim of manner states

> "Supermaxim: be perspicuous (clearly expressed or presented, lucid).
> 1) Avoid obscurity of expression,
> 2) Avoid ambiguity,
> 3) Be brief (avoid unnecessary prolixity), and
> 4) Be orderly" (27).

To understand what Grice was up to with his maxims of conversation we must pay attention to an important distinction between metaphysical 'determination' of meaning and epistemological 'determination' of meaning (Neale Forthcoming). That is, metaphysical determination of meaning is what makes it the case that some speaker's utterance has some meaning, and epistemological determination of meaning is the inferential processes hearers go through to work out that meaning (Neale Forthcoming). This is a fundamental point about the roles and aims of inquiries into meaning. Grice's theory of meaning is a theory of what metaphysically determines meaning; other theorists, such as Sperber and Wilson, are engaged in a project of developing a theory of the epistemological determination of meaning (Neale Forthcoming). This distinction allows us to clearly situate Grice's project and helps to explain certain features of his theory.

Grice's ideas are some of the most enduring in philosophy of language and have been accepted by many linguists and psychologists. His work has also been the subject of much debate and criticism. One source of criticism is of Grice's theory is the Cooperative Principle and the demands this principle places on interlocutors.

Steven Pinker

One such critic is psychologist Steven Pinker. In his 2007 book *The Stuff of Thought* and later papers in 2008, 2010, and 2011 Pinker proposes a number of instances of perceived non-cooperation as problem cases for Grice and his Cooperative Principle. These are sketched out in *The Stuff of Thought* and spelled out in mathematical detail (Pinker, Nowak, and Lee 2008) and experimental results (Lee and Pinker 2010) in later papers.

In setting up his argument against Grice in *The Stuff of Thought*, Pinker writes,

> We'll begin with a famous theory from the philosophy of language that tries to ground indirect speech in pure rationality—the demands of efficient communication between two cooperating agents. This Spock-like theory will then be enhanced by a dose of social psychology, which reminds us that people don't just exchange data like modems. (Pinker 2007a: 375)

In a paper published the following year Pinker and his coauthors write,

> Existing theories of indirect speech are based on the premise that human conversation partners work together toward a common goal—the efficient exchange of information, in the influential theory of H. P. Grice (5) … Yet a fundamental insight from evolutionary biology is that most social relationships involve combinations of cooperation and conflict. (Pinker, Nowak, and Lee 2008: 833)

He further presents his proposal in opposition to Grice writing,

> Our theory of the Strategic Speaker supplements the traditional approaches with the insight from evolutionary psychology that most social interaction involves mixtures of cooperation and conflict rather than pure cooperation. (Pinker 2011: 2866)

Pinker continues this characterization with,

> Grice came to conversation from the bloodless world of logic and said little about why people bother to implicate their meanings rather than just blurting them out. We discover the answer when we remember that people are not just in the business of downloading information into each other's heads… (Pinker 2007a: 379)

Pinker's criticism against Grice has grown into a sustained attack ranging across a wide body of work.

Pinker goes on to propose a number of instances of perceived non-cooperation as problem cases for Grice and his Cooperative Principle. The first is a case of bribing a maître d'. We are to imagine that someone wants a table in a busy restaurant. This person says something like "Is there any way you could speed up my wait?" while handing the maître d' $20 intending to implicate a bribe. Pinker has us consider the question, "What are the benefits of implicating the bribe rather than stating it literally?" The second case is of bribing a police officer. We are to imagine that someone gets pulled over for speeding. This person says something like "Can't we settle this here?" intending to implicate a bribe. Again, we are to consider, "What are the benefits of implicating the bribe rather than stating it literally?"

Pinker ties these cases to what he sees as the problem for Grice, writing that the police bribe case,

> … is inconsistent with the traditional idea that indirect speech is an implementation of pure cooperation: The driver here is using indirect speech not to help the honest officer attain that goal (viz. to enforce the law) but rather to confound that goal. (Pinker, Nowak, and Lee 2008: 834)

It is important to note that with these examples Pinker does not take himself to be presenting an analysis of bribing *per se* but uses these bribing cases as an instance of *perceived non-cooperation*—his true target. As such, my treatment of Pinker's cases here is not meant to be a treatment of these cases *qua* instances of bribing, but *qua* instances of perceived non-cooperation.

Pinker writes that "any scenario like that in which the best course of action depends on the choices of another actor is in the province of game theory" (2007a: 393) and presents game-theoretic frameworks such as the one below to show these costs and benefits:

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Ticket	Pay bribe; go free

The first matrix displays the options that are available to the speaker if implicature is not on the table.[2] In the first matrix if the speaker has

[2] Strangely, Pinker characterizes this first matrix as being what an agent acting in a "perfect" Gricean way, Maxim Man, would be faced with. This is a further serious mischaracterization of Grice that evinces Pinker's inadequate grasp of his theory—but which I will not go into in detail on here beyond this footnote. Pinker writes, "Consider a perfect Gricean speaker who says exactly what he means when he says anything at all. Maxim Man is pulled over for running a red light and is

an honest officer the best option is to pay the ticket; if the speaker has a dishonest officer the best option is to pay the bribe and go free. The second matrix shows the options available if implicature is on the table. Note that in the implicature row of the second matrix the speaker has the options of either paying the ticket, the best option if he has an honest officer, or paying the bribe and going free, the best option if he has a dishonest officer. On Pinker's framework, the speaker, in implicating the bribe, has as the possible outcomes the two best results given some officer.

Further, Pinker writes that for the second matrix to obtain, it must be the case that the speaker "knows that the officer can work through the implicature and recognize it as an intended bribe, and he also knows that the officer knows that he couldn't make a bribery charge stick in court because the ambiguous wording would prevent a prosecutor from proving his guilt beyond a reasonable doubt" (2007a: 394). This is a lot of knowledge for someone to have about a police officer he just met, especially considering that at the same time the speaker is not supposed to know whether the officer is honest or dishonest.

In such sections of *The Stuff of Thought* and later work building off these cases (2007a; 2007b; 2011; Pinker, Nowak, and Lee 2008; Lee and Pinker 2010) Pinker advances the following implicit argument.

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker's account instead.

In what follows I will argue that Pinker's premises 2 and 3 are both false, beginning with premise 3. This has the result that his argument is not sound, and we do not have support for the conclusion.

pondering whether to bribe the officer. Since he obeys the maxims of conversation more assiduously than he obeys the laws of traffic or the laws of bribery, the only way he can bribe the officer is by saying, 'If you let me go without a ticket, I'll pay you fifty dollars'" (Pinker 2007a: 393; See also Pinker, Nowak, and Lee 2008: 834 for this characterization again). The contrasts between this bizarre characterization of Grice and what his theory actually consists of will become apparent when I discuss the ways the maxims can be weighed against the others and in my positive proposal in later sections. To describe Grice's theory as one that does not allow for conversational implicatures demonstrates a misunderstanding of the most fundamental kind.

Third Premise: Pinker and Grice on the Cooperative Principle

Let me begin with my argument that Pinker's third premise is false. That is, I will now argue that it is not true that Grice only considers communication that occurs in contexts of pure cooperation, as given by his Cooperative Principle.

As noted in the previous section, Pinker contrasts his view with Grice's by saying that Grice's theory of conversation is "based on the premise that human conversation partners work together toward a common goal—the efficient exchange of information" (Pinker, Nowak, and Lee 2008: 833). Pinker characterizes Grice's view as relying on "pure cooperation" (Pinker 2011: 2866; Pinker, Nowak, and Lee 2008: 833) and writes that Grice overlooks the fact that most interactions involve both cooperation and conflict. He writes that "in trying to deduce the laws of conversation from a 'Cooperative Principle'" (note the use of scare quotes) Grice is "guilty" of assuming "the speaker and the hearer are working in perfect harmony" (Pinker 2007a: 392).

However, contrary to the way Pinker presents him, Grice does *not* demand full cooperation. What Grice actually says about cooperation is the following,

> Our talk exchanges do not normally consist of a succession of disconnected remarks, and would not be rational if they did. They are characteristically, to some degree at least, cooperative efforts; and each participant recognizes in them, to some extent, a common purpose or set of purposes, or at least a mutually accepted direction…at each stage, some possible conversational moves would be excluded as conversationally unsuitable. (26)

We see in this section that there are a number of places where Grice hedges or qualifies his position on cooperation in some way, using words and phrases such as "do not normally", "characteristically", "to some degree", "at least", "to some extent", "or at least", and "some…would be excluded". Grice is not making a sweeping claim that all conversation exchanges are always fully cooperative, but is making the tentative claim that they characteristically involve some degree of cooperation.

Grice goes on to be quite clear about what features a conversation must have to be cooperative[3] on his view, paraphrased here.

[3] It should be noted that Grice is not making a claim about the sort of cooperation that must have been required for language to have begun in homo sapiens as a species or with respect to particular human populations. Grice's arguments are made within a culture where we already have both the cognitive requirements for language use in general and a system where particular languages exist. There certainly are interesting questions to be asked about the sort of cooperation required for the cognitive capacity for language to evolve and be selected for, and for language use to be a sustainable system in a species (Dawkins and Krebs 1978; Sterelny 2003: Papineau 2005; Dessalles 2007; Tomasello 2010), but we can pull such questions apart from what I consider here.

Characteristic features that jointly distinguish cooperative transactions:[4]

1. Have some common *immediate* aim
2. Contributions of participants are mutually dependent
3. The transaction will continue until it reaches its natural terminus (the interlocutors will not just walk away in the middle of the conversation) (29)

Grice further elaborates on what he means by this requirement that interlocutors have some "common immediate aim". He writes, "The participants have some common immediate aim, like getting a car mended; their ultimate aims may, of course, be independent and even in conflict—each may want to get the car mended in order to drive off, leaving the other stranded" (Grice 1989: 29; see also 30). We see here that Grice explicitly allows for cases where the ultimate aims of conversational participants are in opposition.

With Grice's specific claims at hand let us reflect on one of the cases Pinker presented. Recall that Pinker wrote,

> [the police bribe case] … is inconsistent with the traditional idea that indirect speech is an implementation of pure cooperation: The driver here is using indirect speech not to help the honest officer attain that goal (viz. to enforce the law) but rather to confound that goal. (Pinker, Nowak, and Lee 2008: 834)

However, someone who does not want the officer to enforce the law as in Pinker's police bribe example can still be cooperative with respect to the conversation.

We could say the speaker is being "Communicatively Cooperative" although not "Ultimately Cooperative",[5] understanding these two notions in the following way:

Communicative Cooperation

Demanded by Grice's Cooperative Principle

1. Speakers abide by Grice's maxims of conversation of quality, quantity, relation, and manner in formulating their utterances.

2. The conversational participants have enough immediate aims in common that they will not abandon the conversation altogether.

[4] Dan Sperber and Deirdre Wilson argue that the notion of cooperation demanded by Relevance Theory is weaker than what is demanded by Grice. They write "… Grice assumes that communication involves a greater degree of cooperation than we do. For us, the only purpose that a genuine communicator and a willing audience necessarily have in common is to achieve uptake: that is, to have the communicator's informative intention recognized by the audience" (1989: 161). Although, it is not clear what this difference really amounts to, and how some set of interlocutors could share the desire that the audience recognize the communicator's informative intention if they did not have even a minimal common purpose.

[5] I will henceforth capitalize these terms to show I am using them in a technical way according to the definitions I provide here.

Ultimate Cooperation

Not Demanded by Grice's Cooperative Principle

1. Shared desire for some outcome with respect to the conversation, overall interaction, or some further long-term aim.

Having these two notions at hand can help to streamline discussion of cooperation and communication.

It becomes clear, for one thing, that participants often satisfy the requirements of Communicative Cooperation although they have divergent ultimate aims. Arguments, for one—which are certainly uncooperative in some sense—are very often Communicatively Cooperative according to the conditions just provided. Consider the following argument, between two participants, 1 and 2,

1: You never put the forks in the right place.
2: Well I think it's stupid to separate the salad forks from the dinner forks.
1: Just because you don't know the difference between a salad fork and a dinner fork doesn't mean I should have to live like a heathen!
2: A heathen? You're the one who leaves towels on the floor!
1: When's the last time I left a towel on the floor?

Notice that this conversation is closely in accordance with Grice's Maxims of Conversation. Specifically,

a) Each reply is directly in response to the previous comment (Maxim of Relation[6]).
b) They are consistent in tone (Maxim of Manner).
c) Each participant is expected to only say things that are true and are challenged if not (Maxim of Quality). And,
d) Each retort is about the same length (Maxim of Quantity).

This means all the maxims of conversation were abided by. At the same time, neither conversational participant abandoned the exchange. Thus, according to the conditions laid out, this is a clear example of Communicative Cooperation.

This example would still be a clear example of Communicative Cooperation even if each person were trying to achieve any of the following ultimate aims: a) make the other feel insecure, b) drive them to madness c) have the silverware drawer sorted so that they have easy access to the knife they plan to use to kill the other, and so on. None of these ultimate aims affects the fact that the dialogue is Communicatively Cooperative.

[6] As David Lumsden insightfully points out, there may be an important connection between what he calls the "extra-linguistic goals" and the "linguistic goals" insofar as the interlocutors' ideas about what the further aim of the conversation is could shape what counts as relevant to the conversation, as in Grice's petrol case. At the same time, there is not always this interplay between the two, as he also renognizes (Lumsden 2008: 1901).

The silverware drawer argument example just discussed is one that may at first seem to not be cooperative in the sense demanded by Grice, but in fact is, as becomes evident with a clearly defined distinction between Communicative Cooperation and Ultimate Cooperation.

In making his arguments against Grice Pinker conflates Ultimate Cooperation with Communicative Cooperation. The cases Pinker's presents as counterexamples do not, in fact, present a problem for Grice's Cooperative Principle. And neither would other arguments that conflate Ultimate Cooperation with Communicative Cooperation.

Thus, we are now in a position to return to Pinker's implicit argument.

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker instead.

I have argued here that the third premise—that Grice only considers communication that occurs in contexts of pure cooperation, as given by his Cooperative Principle—is false.

Failing to be Communicatively Cooperative

With the silverware argument example above I presented a case that might at first appear to be not cooperative in the relevant sense for Grice, but which, upon defining Communicative Cooperation technically, clearly is. This might raise the further question, "What, then, does it take for some utterance to fail to be Communicatively Cooperative?" This might raise the further, related question, "Is every utterance that appears to violate a maxim a case of the speaker failing to be Communicatively Cooperative?" The answer to this second question is 'no'.

A maxim of conversation may appear to be violated in cases of implicature. This is central to Grice's theory of conversation. The fact that a maxim appears to have been violated at the level of what the speaker said is what tips off hearers that the speaker may be implicating something at the level of what the speaker meant[7] (1989: 33; Neale 1992).

There are three types of implicature in Grice's theory:

1) Group A—those in which no maxim is violated, or at least in which it is not clear that any maxim is violated.
2) Group B—those in which a maxim is violated, but its violation is to be explained by the supposition of a clash with another maxim.
3) Group C—those in which a maxim is flouted, or exploited.

It is clear that Grice recognizes that if one fails to uphold a maxim it

[7] For more details on what the speaker said, or made as if to say, and what the speaker meant see Grice (1989: 33–34) and Neale (1992: 13–16).

may be justified (because of a clash with another maxim as in Group B or because a maxim is flouted as in Group C).

A speaker who makes a Group B or Group C implicature does so by abiding by the Cooperative Principle. Such cases are not counter-examples to Grice but features of his theory. If the speaker appears to fail to uphold a maxim and it is not a Group B or Group C implicature, this may mean 1) the speaker is *misleading* the hearer with an unostentatious violation—which includes behavior such as lying, or 2) the speaker is *opting out* of the conversation. This can be understood according to the following chart:

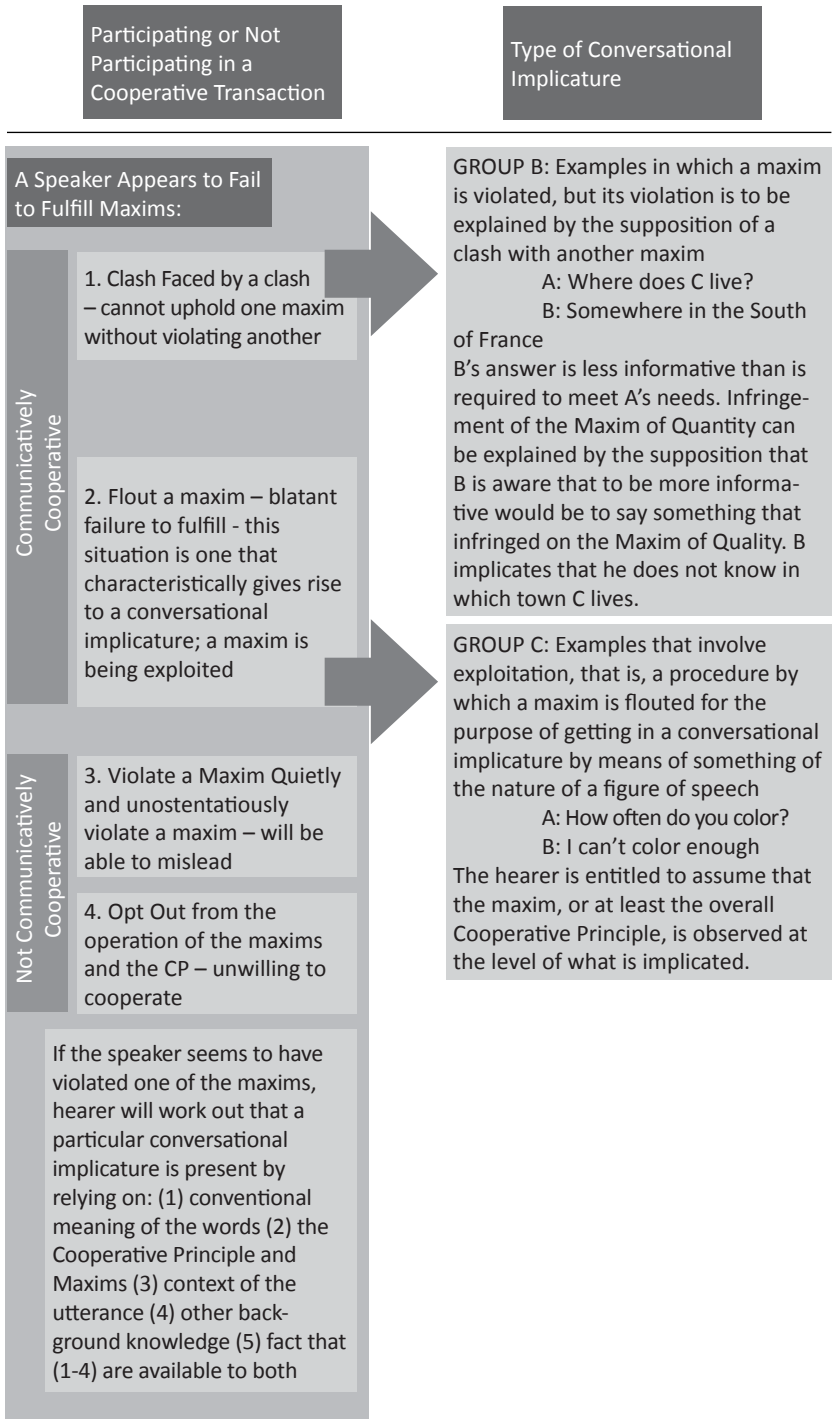
Participating or Not Participating in a Cooperative Transaction
Type of Conversational Implicature
A Speaker Appears to Fail to Fulfill Maxims:
Communicatively Cooperative
1. Clash Faced by a clash – cannot uphold one maxim without violating another
2. Flout a maxim – blatant failure to fulfill - this situation is one that characteristically gives rise to a conversational implicature; a maxim is being exploited
Not Communicatively Cooperative
3. Violate a Maxim Quietly and unostentatiously violate a maxim – will be able to mislead
4. Opt Out from the operation of the maxims and the CP – unwilling to cooperate
If the speaker seems to have violated one of the maxims, hearer will work out that a particular conversational implicature is present by relying on: (1) conventional meaning of the words (2) the Cooperative Principle and Maxims (3) context of the utterance (4) other background knowledge (5) fact that (1-4) are available to both
GROUP B: Examples in which a maxim is violated, but its violation is to be explained by the supposition of a clash with another maxim
A: Where does C live?
B: Somewhere in the South of France
B's answer is less informative than is required to meet A's needs. Infringement of the Maxim of Quantity can be explained by the supposition that B is aware that to be more informative would be to say something that infringed on the Maxim of Quality. B implicates that he does not know in which town C lives.
GROUP C: Examples that involve exploitation, that is, a procedure by which a maxim is flouted for the purpose of getting in a conversational implicature by means of something of the nature of a figure of speech
A: How often do you color?
B: I can't color enough
The hearer is entitled to assume that the maxim, or at least the overall Cooperative Principle, is observed at the level of what is implicated.

A quiet and unostentatious violation of a maxim is a case where the speaker is not being Communicatively Cooperative, but hopes the hearer will proceed as though the speaker is. The speaker formulates utterances with an audience that takes the speaker to be acting in a Communicatively Cooperative way. As Grice notes this can lead to manipulation of the hearer. Lying is one such form of manipulation achieved this way.

Opting out occurs when two potential participants in a talk exchange do not have a common immediate aim, and, in fact, may have immediate aims that are in conflict. For instance, imagine the following three scenarios. You 1) receive a letter in the mail that says "Action Required: Important Survey for Residents". You then notice the return address says it is from People for the Ethical Treatment of Animals. You'll not open it. You see PETA is trying to get you to open their mail by deceptive means and opt out, tossing the unopened envelope in the recycling. 2) A woman is walking down the street on her way to a date. A man sitting near the sidewalk says, "God damn". She will not reply or make eye contact. She opts out of the exchange. 3) You are walking by the Empire State Building on a visit to CUNY for a conference, trying to get to Starbucks on 33rd Street and back to 34th Street before the next talk starts. A man in a red vest carrying pamphlets about the Empire State asks you, "Going up?" You do not want to go to the top of the Empire State Building. You will not reply. Grice recognizes these uncooperative possibilities.

But, in any given case, before we can conclude that an apparently violated maxim means the speaker is being uncooperative, we must make sure the violation does not result from an implicature of the Group B or Group C kind.

With these cases in mind, we can now make a necessary final addition to our understanding of Communicative Cooperation.

Communicative Cooperation

Demanded by Grice's Cooperative Principle

1. Speakers abide by Grice's maxims of conversation in formulating their utterances.
2. The conversational participants have enough immediate aims in common that they will not abandon the conversation altogether.
3. If one appears to fail to uphold a maxim it is justified within Grice's theory of implicature—e.g. because of a clash with another maxim as detailed in the chart on the previous page.

Ultimate Cooperation

Not Demanded by Grice's Cooperative Principle

1. Shared desire for some outcome with respect to the conversation, overall interaction, or some further long-term aim.

Second Premise: Veiled Speech is Uncooperative

Equipped, now, with a clearer picture of Grice's theory, and what is demanded by his Cooperative Principle, we can return to Pinker's implicit argument. I will now argue that the second premise is false.

Pinker's Argument

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker's account instead.

In moving toward showing Pinker's second premise to be false we can next ask, "Are there any circumstances in which a speaker can be intentionally obscure while still being Communicatively Cooperative?"

In the section where Grice details a number of Group C implicatures he considers a case of apparent obscurity that arises as a result of the presence of a third party. He writes,

> Obviously if the Cooperative Principle is to operate I must intend my partner to understand what I am saying despite the obscurity I impart into my utterance. Suppose that A and B are having a conversation in the presence of a third party, for example, a child, then A might be deliberately obscure, though not too obscure, in the hope that B would understand and the third party not. (Grice 1989: 37)

For illustration of such a scenario, we can consider a scene from a fiction (King 2008), which Grice seems to have anticipated almost exactly. In this scene, four women—Miranda, Carrie, Charlotte, and Samantha—are meeting for brunch. Charlotte has brought along her young daughter, Lily. Carrie and Lily are coloring at the table.

> Miranda: How often do you guys have sex?
>
> Lily: Sex!
>
> Charlotte: Miranda, please! (Points to her daughter, Lily)
>
> Miranda: What? She's 3! She doesn't know what it means. I'm 41 and I still don't know what it means.
>
> Charlotte: I know, but she is repeating everything.
>
> Samantha: If I had known that girl talk was going to be on lockdown I wouldn't have flown 3,000 miles.
>
> Charlotte: No, we can talk. Let's just not use the word.
>
> Miranda: Fine. How often do you guys …
>
> Carrie: (looks up from coloring) … color?
>
> Charlotte: Thank you!
>
> Samantha: Well, I can't color enough. I could color all day, every day, if I had my way. I would use every crayon in my box.
>
> Carrie: We get it. You love to color. (Turns to Miranda.) Why are you asking? (King 2008)

We see in this example exactly the sort of case Grice mentions as a hypothetical. In this interaction there is intentional and careful obscurity used in the conversation in the form of the adoption of a new word in place of an ordinary word, and we see novel, metaphorical riffs off the new word by Samantha. The speakers in this conversation are being intentionally obscure in the hopes that the other women at the table will understand their meaning and the child will not.

Note that, at the same time, the substitute phrase Carrie comes up with is one that would not stand out to the child as being something that would be strange for people to discuss. If, on the other hand, Carrie had suggested that they discuss "how often they have noodles for arms" even the child would likely pick up on this as being strange and ask what they are talking about. Also, we see in Samantha's objection, a reason why the topic is not abandoned altogether for a time when the child is absent.

Of course, there are many techniques a speaker might employ if she has an utterance she wishes only one person to hear. She may meet with some audience alone. She may attempt to make her utterance known to only one audience member by manipulating acoustic means—that is, by whispering. However, whispering is recognized as behavior one engages when excluding some potential audience member who otherwise would hear, and, thus, if done the presence of company who feels entitled to be included, will not be used for reasons of politeness. A speaker could switch to a language one hearer knows and another does not. Although this can, again, be seen as rude and may be objected to by the excluded hearer. Thus, in the presence of audience members who for social reasons feel they should not be ostentatiously excluded from the conversation by means such as whispering, switching language, or asking to speak to someone alone, obscurity remains one of the least socially costly options for excluding some audience.

For Grice, a speaker can be intentionally obscure while still being Communicatively Cooperative in cases where there are two audiences the speaker has in mind. The speaker uses veiled or obscure wording so that one audience will understand the meaning of the utterance and the other will not. In the case of the child overhearing an adult conversation it is her presence that explains the use of veiled speech. The speaker is cooperating with the adult listeners insofar as the speaker works to ensure they can still understand what she means by the veiled speech. Such use of veiled speech requires formulating an utterance that will be understood by one audience but not the other.

This gap in interpretive understanding is to be explained by differences in background knowledge and interpretive tendencies of the hearers, which was gauged by the speaker. In the "coloring" example, if the speaker incorrectly gauges the background knowledge of her audience she may: 1) produce an utterance that the child understands and/or 2) produce an utterance that the adults do not understand.

Alternative Game-Theoretic Framework

So, we see that for Grice apparent obscurity can sometimes be explained in terms of a speaker wishing to be understood by one audience in terms of the literal content of their utterance and by the second audience in terms of the implicated content of their utterance. In the police bribe scenario Pinker has it be the case that the speaker implicates a bribe so that there is less chance he will get arrested. However, if the speaker properly implicated the bribe (taking into account the interpretive capacities of the hearer), there would be no lack of clarity from the perspective of the police officer. In order to implicate some content, *p*, a speaker must intend that his audience will recognize *p* (Grice 1989: 39). So *to whom* is the content of the implicature in this example really less clear?

I would like to propose that what is going on in the bribing case is that the speaker has formulated the utterance with two audiences in mind (as in the case Grice proposed). The implicature is clear to one audience and unclear to another audience.

The second audience the speaker has in mind does not have to be present. The speaker will think this person could later hear about their utterance and judge their behavior on the basis of it. If the second audience never in fact hears of the utterance this does not change what it meant (recall Neale on 'determination').

With this in mind, let us revisit the game-theoretic framework Pinker presents to explain the police bribe case.

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Ticket	Pay bribe; go free

Recall that one assumption of Pinker's framework is that in order for a speaker to implicate the bribe along the lines of the third row, it must be the case that the speaker "knows that the officer can work through the implicature and recognize it as an intended bribe, and he also knows that the officer knows that he couldn't make a bribery charge stick in court because the ambiguous wording would prevent a prosecutor from proving his guilt beyond a reasonable doubt"[8] (Pinker 2007a: 394). There is no reason to make this assumption.

This case can instead be understood first in terms of two different types of interpreters in the officers: 1) those officers who will understand the implicature, and 2) those officers who will not. This is not equivalent to the difference Pinker details between honest officers and

[8] Although Pinker later contradicts himself and writes that for certain bribes that are more heavily veiled there is "some risk that they might go over the head of a bribable officer" (Pinker 2007a: 395).

dishonest officers. For Pinker the difference between honest and dishonest officers is spelled out in terms of *the officer's response to the implicature, not whether or not he or she understand an implicature was made.* There is a gap between whether some person understands an implicature was made and how they act on the basis of this recognition. This is an important gap to recognize and the following framework will do so.

If we understand the officers as two different classes of interpreters that constitute two different audiences, we can see this case as a close parallel to the one Grice presents, where the officers who do not understand an implicature was made are akin to the child, in the sense that they understand the utterance in terms of its literal meaning, and those who do understand an implicature was made are akin to the adults, in the sense that they understand the utterance in terms of its implicated meaning.

This means that we must revise Pinker's game theory framework to reflect these two audiences.[9]

Audience 1. Savvy Police Interpreter—will understand the implicature

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Arrest for bribe	Pay bribe; go free

Audience 2. Naïve Police Interpreter—will not understand the implicature

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Ticket	Ticket

Veiled speech has the added benefit that some honest officers who might have otherwise arrested a driver for bribery will not understand that a bribe was made. At the same time, this framing highlights one of the risks of using veiled speech with an audience one is unfamiliar with: there is a possibility the intended audience will not work it out—in this case, losing out on officers who would take the bribe if they could work out one was offered.

A negotiation between these two possible audiences leads to (part of) the benefit of having some bribe be more or less explicit. A more explicit bribe increases a) the chances an officer who would take the bribe will work out one is being made and b) the chances an officer who will arrest for a bribe will work out one is being made. A more veiled bribe increases a) the chances an officer who would take the bribe will not

[9] Matrices modeled after Farrell and Gibbons (1989).

work out one is being made and b) the chances an officer who will arrest for a bribe will not work out one is being made.[10]

I am not arguing that this is what Pinker really meant but that what is really going on in the sorts of cases he discusses is best modeled with this framework that allows for more nuance, and makes clear the ways in which his account can be brought into conjunction with what Grice says about multiple audiences.

To most fully model the circumstances of the example we must also consider another audience. Should the driver be arrested, another audience becomes relevant: the jury. We should also consider the utterance with respect to the second set of audiences. In doing so, we separate two issues: 1) whether or not the police officer will understand the implicature, and 2) (bringing in a new potential audience) whether or not a jury will understand the implicature. If the speaker does formulate his utterance with two audiences in mind, as I suggest, this means that we must further revise Pinker's game theoretic framework to reflect these two audiences.

If it is right that the utterance was formulated with both of these pairs of audiences in mind—both 1) a savvy and naïve officer, and 2) a savvy and naïve jury[11]—the final game-theoretic framework would be a combination of two matrices.

[10] Let me provide a brief note on how this utterance could plausibly be received as something other than a bribe. When I presented this talk at the 2015 Language and Linguistics Conference in Dubrovnik, along with many of the other papers in this volume, there was an interesting discussion stemming from a professor who lives and works in Mexico struggling to understand how this utterance could be considered anything other than a bribe. In fact, in the U.S., where Pinker lives and works, the rules for how speeding tickets are paid vary by state. So, this example makes more sense in its setting in America where, because of the variance in rules by state, it is likely that there may be genuine ignorance or confusion on the part of some drivers about what the rules are in that state and where bribing is less common than in many other countries (Walton 2013).

[11] I realize these names "naïve" and "savvy" do not fully capture what is going on here but have decided to stick with them because they do capture some piece—although in perhaps too dramatic a fashion. "Literal" and "read-between-the-lines" could be another way of spelling it out. Also, this is not meant to be a fixed feature of the agent per se, but the agent faced with some particular utterance. Some police officer might be very good a recognizing a bribe in his native language, but more likely to interpret literally in his second language, for example.

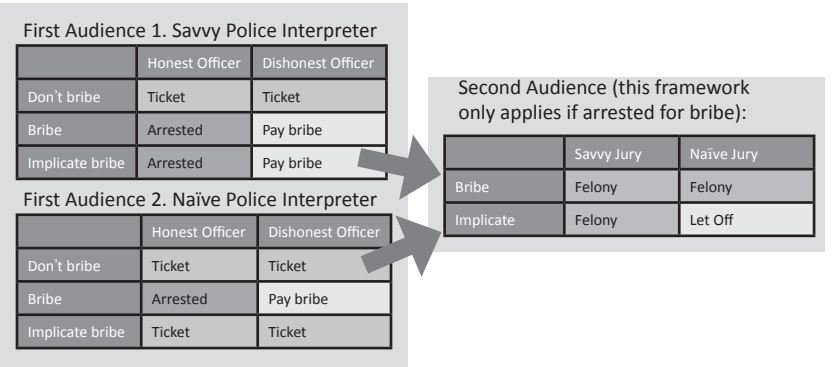

On the left, calculation of what the best move is in this step depends first on the values given in each of the cells and the ratio of honest and dishonest police officers. This is calculated within each matrix. Then the matrices are weighted with respect to the odds of a police officer being savvy or naïve. On the right, calculation of what the best move is in this step is done within the grid and depends on the ratio of savvy to naïve jurors.

The honesty or dishonesty of the police officer comes from the decision to put the utterance up for evaluation by the second audience. This feature is independent of this audience's ability to work out implicatures. A similar thing is going on in Pinker's maître d' case. In each case, who the relevant audiences are depends on who the speaker has in mind, which is a function of who the speaker believes may respond in a way that has costs or benefits for him/her.

Takeaway

We can now, once again, return to Pinker's argument.

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker instead.

I have argued here that premise 2 is false, and argued in previous sections that premise 3 is false. This means that Pinker's argument is not sound and that we do not reach his intended conclusion.

Despite the flaws in Pinker's argument, the first premise in his argument—which I have not argued against and which points to the inefficiency of indirect speech as an important puzzle—is true, and is an important point that has been overlooked in much literature on implicature. Implicature is costly: it is difficult for a speaker to produce, difficult for a hearer to interpret, and there is increased risk that the message will go wrong along the way (Kruger 2005). Because of these potential

costs to the use of indirect speech we ought to seek an understanding of why and when speakers choose to use implicature rather than state content literally. Pinker's idea of using a game theoretic framework to model the costs and benefits of certain utterance types—although flawed in the details of its execution and especially in his understanding of Grice—is a very helpful way to consider this question. The multiple audience framework I have presented here could be utilized to help us understand why implicature is used in a number of cases.[12]

The multiple audience framework provides one piece of the puzzle of when we use implicature. There are many other pieces of the puzzle and reasons for using implicature in addition to those discussed here.[13] It is a puzzle that anyone who makes use of Grice's framework of conversational implicature or other similar accounts of veiled speech should be interested in solving. Even if there is a quibble about my treatment of Pinker's specific cases here I hope to have succeeding in making the more general point against the general contours of Pinker's arguments against Grice.

Up to this point I have sought to stress the ways that Pinker mischaracterizes Grice on cooperation. I also presented a new game theoretic framework that deals with apparent uncooperative obscurity in terms of two audiences, and allows us to better explain Pinker's bribing case. My ultimate goal, however, is a more general one about the aims of the Gricean project, the tools at hand to help us meet those aims, and overlooked but worthy questions relevant to that project.

a. Deniability

I will now consider some objections and addendums to the proposal as presented thus far. Someone responding to the ideas put forth in this paper may object that this analysis is unnecessary because the benefits afforded by using veiled speech can be accounted for by saying that the veiled speech is used so that the implicated content is deniable. Pinker, himself, appeals to a notion of "plausible deniability" to explain the bribing cases.[14] However, if we are to say that some content is deniable, we must be clear about what is demanded by deniability. Four points of needed clarification emerge.

[12] I have in mind Grice's tea party and letter of recommendation cases.

[13] For example, one other piece of the puzzle might be in our judgments that those who express criticism with sarcasm rather than with literal speech are seen as being more in control (Dews, Kaplan, and Winner 1995; Pinker 2007a: 379).

[14] I seem to be at odds here, too, with Pinker, who writes that one who has uttered "They never seem to have enough salt shakers at this restaurant" can deny having asked for the salt, and that someone hearing such an utterance could simply ignore it without being rude. He also writes that for this reason, ostensibly following Brown and Levinson (1987), utterances such as "It's too dark to read" as a request for the lights to be turned on are more polite than direct requests, again dubious claim. It seems such a speaker, who assumes it is the duty of the hearer to quell all concerns, speaking to anyone other than a servant, would be quite rude and unpleasant.

First, in saying that some claim is deniable, what are we saying about the likelihood that this act of denying succeeds? Surely this objection rests on deniability being something stronger than just that some claim *could* be denied—for any claim whatever *could,* in principle, be denied. Consider the following utterances made by the same person:

1) I'm from Athens.
2) Oh—but I don't mean I'm from Athens.

Or, drawing from a recent political news event,

3) "Perhaps there are two Donald Trumps."
4) "I don't think there's two Donald Trumps. I think there's one Donald Trump." (Nguyen 2016)

Such a series of remarks would confuse and require clarification by the hearer. It seems not that the speaker has succeeded in denying the content of the first utterance but simply said two things that are contradictory. Deniability, then, seems to demand some reasonable chance of success of convincing an interlocutor that the speaker did not mean the denied content.

Second, we have a further point of clarification demanded of anyone who relies on a notion of deniability. That is, he or she must take a stance on whether succeeding in denying some utterance is an act of 1) getting the hearer to believe that the asserted proposition is not true, 2) getting the hearer to believe that the speaker did not mean the asserted content (although it may be true), or 3) Both 1) and 2)—that is, an act of getting a hearer to believe both that the asserted proposition is not true, and that the speaker did not mean the asserted content.

Third, at the same time, we would want to isolate deniability from clarification. Consider the following modification on two of the utterances above.

5) I'm from Athens.
6) Oh—but I don't mean I'm from Athens, Greece; I'm from Athens, Georgia.

In this case we do not have a true case of deniability because in this instance the speaker (assuming he had a genuine intention to speak about Athens, Georgia in the first utterance) is not denying the proposition he meant in the first utterance. What he was likely doing in this case is realizing after he made the utterance that he had been misunderstood, and wanted to rectify this by making a more explicit utterance. Such a case would result from a speaker not accurately assessing how he would be interpreted. Thus, we have a further distinction between deniability applying to 1) what the speaker meant or 2) what the hearer took the speaker to mean, or 3) what the speaker thought (nonfactive) the hearer took her to mean. Any appeal to deniability would require resolution of what the required target of the act of denying is.

Fourth, in saying that some content is deniable the further question also arises of to whom the content is to be deniable, especially if we are considering utterances formulated with multiple authors in mind.

b. Cancellability

Such discussions of deniability can be somewhat clarified by appealing to Grice's notion of cancellability. For Grice, if some content was implicated, rather than stated literally, the speaker should be able to cancel the implicature (39). Consider the following example:

7) Mr. X has wonderful handwriting.
8) Of course, I am not saying he is no good at philosophy.

In this example there is a reasonable chance that the speaker would succeed in cancelling the content of utterance 7—or at least a better chance than if he had stated the implicated content literally in the first utterance.

Grice's notion of cancellability, can, at times appear to be at odds with his spelling out of implicatures. For, in order for some speaker s to have formed a genuine meaning-intention with respect to some content p and some audience a by uttering some utterance r, s must expect that a will work out p on the basis of r. If this is right, why would s then think he could subsequently get a to believe $\sim p$?

One way around this apparent problem with cancellability is to say that what is going on here is not that s get a to believe $\sim p$, but that s get a second audience, a^*, to believe $\sim p$. We could say, then, that the cancellability of some content p is a feature of the ability to get a^* to believe $\sim p$ after hearing r. We see, then, that deniability—as understood in terms of Grice's notion of cancellability—is not a true alternative to the multiple audience framework but a concept that is already present in the Gricean framework and could be brought together with the multiple audience framework to be fully fleshed out.

c. The Demands of the Multiple Audience Framework

Some object to Grice because his account of speaker meaning is thought to be too psychologically demanding. Jennifer Hornsby, for example, writes the following of Grice's theory: "I think that this ought to seem ludicrous. Real people regularly get things across with their utterances; but real people do not regularly possess, still less act upon, intentions of this sort...notice that an enormous amount would be demanded of hearers, as well as speakers, if such complex intentions really were needed to say things" (Hornsby 2000: 95). The framework I present here is an especially complex application of the Gricean framework. What, then, of this objection to Grice, which seems doubled here?

There are many behaviors we engage in that seem dizzyingly complex when spelled out, such as athletes catching a baseball, musicians playing a saxophone solo, or firefighters deciding when to run out of a burning building (Kahnemann 2013). The complexity of these behav-

iors themselves is not enough to insist that we do not really engage in them, but further reasons to be in awe of what the human mind can do.

With regard to the multiple audience framework, there are undoubtedly many cases where, because of the potential costs and benefits, speakers do successfully navigate the waters of multiple audiences. The "coloring" example above is one such instance. A starting place of the objection—the belief that formulating some meaning intention with respect to multiple audiences is highly demanding—is right, and this is why it routinely goes wrong. However, the psychological demandingness appealed to in this proposal is not enough on its own to reject it.

Applying the Multiple Audience Framework to Political Speech

Failure to recognize that some utterance may be put up for interpretation by a second audience can lead to disastrous consequences. For my concluding example, I will draw on a real-world case that is an instance of such a disaster. This case makes it clear that my suggestion that a speaker could or should formulate an utterance while keeping in mind the possible interpretations of a second audience is not something cooked up for hypothetical examples, but something we routinely engage in.

Mitt Romney lost to Barack Obama in the 2012 presidential election. The blow that may have cost him the presidency resulted from failing to take into account the interpretation a second audience would make of one of his utterances. The gaffe occurred at a Boca Raton fundraiser, which would have been filled with very right-leaning donors to Romney from whom he would like to solicit as much money as possible.[15] With this audience it is most advantageous to use strong rhetoric. Romney told the crowd he was unconcerned with the 47 percent of the electorate who supported Obama (that is, "those people", "victims" who take no "personal responsibility") (Leibovich 2014).

However, unbeknownst to Romney, someone in the audience filmed his utterance and posted it online. Thus, this utterance that first took place "behind closed doors" became available for interpretation by a second audience, the general public, who were outraged by what Romney had said. A headline in *Mother Jones,* which broke the story, read, "Secret Video: Romney Tells Millionaire Donors What He REALLY Thinks of Obama. When he doesn't know a camera's rolling, the GOP candidate shows his disdain for half of America" (Corn 2012). After this footage came out, Romney's poll numbers slipped irreparably.

[15] Note here the two audiences are made up of two types of interpreters, not two individuals. The multiple audience framework can be applied either to individual audiences, or audiences grouped according to some similar interpretive tendencies.

On September 30, 2014, nearly two years after he lost the election, Romney did an interview with the *New York Times* in which he reflected on why he lost.

> 'I was talking to one of my political advisers,' Romney continued, 'and I said: 'If I had to do this again, I'd insist that you literally had a camera on me at all times' — essentially employing his own tracker, as opposition researchers call them. 'I want to be reminded that this is not off the cuff.' This, as he saw it, was what got him in trouble at that Boca Raton fund-raiser. (Leibovich 2014)

He continued, "My mistake was that I was speaking in a way that reflected back to the man," Romney said. "If I had been able to see the camera, I would have remembered that I was talking to the whole world, not just the man" (Leibovich 2014). As he later recognized, Romney would been better off if he had formulated his utterance by weighing the costs and benefits with respect to his multiple audiences. He felt he had erred in not doing so in the first place.

This tension is common in political discourse. Politicians are constantly negotiating the line between energizing supporters and making comments that will appeal to—or at least not greatly offend—the general public (Economist 2005; Fear 2007; Haney López 2014; Nguyen 2016). The proposed game theoretic framework of cooperation with multiple audiences could be used to model such negotiations. The framework presented here has wide applicability to a number of types of discourse.

Conclusion

In this paper I have argued against some of the claims made by Steven Pinker in his book *The Stuff of Thought* and in subsequent papers. In particular, I argued that Pinker is misguided in a number of objections he makes to the theories of philosopher H. P. Grice. In arguing for Grice I developed positive notions of the sort of cooperation Grice demands and explained how cases of obscure language can be cooperative within Grice's framework. Attention to the role multiple audiences play in leading to intentional, cooperative obscurity led me to return to some cases Pinker presented against Grice within Grice's treatment of multiple audiences. I concluded by addressing some possible objections to this proposal and by suggesting how the multiple audiences framework I presented could be applied to cases of political discourse.

Acknowledgements

This paper was presented at the CUNY Graduate Center Workshop on Relevance, Pragmatics, and Game Theory in New York in 2014, the Philosophy of Language and Linguistics Conference at the Interuniversity Center in Dubrovnik in 2015, and the Thursday Seminar Series at the Australian National University in Canberra in 2015. I am grateful to the audiences at all these talks for their helpful comments and ques-

tions. In addition, I would especially like to thank Justin Bruner, Ryan DeChant, Michael Devitt, Maita Escurdia, Michael Glanzberg, Michelle Johnson, Dunja Jutronić, Karen Lewis, Eliot Michaelson, Nenad Miščević, Stephen Neale, David Papineau, Ronald Planer, Anne Reboul and Deirdre Wilson.

References

Brown, P. and Levinson, S. C. 1987. *Politeness: Some Universals in Language Usage*. Cambridge: Cambridge University Press.

Dawkins, R. and Krebs, J. R. 1978. "Animal Signals: Information or Manipulation?" In Krebs, J. R. and Davies, J. B. (eds.) *Behavioral Ecology*. Oxford: Blackwell.

Dessalles, J. L. 2007. *Why We Talk: The Evolutionary Origins of Language*. Oxford: Oxford University Press.

Dews, S, Kaplan, J. and Winner, E. 1995. "Why Not Say It Directly? The Social Functions of Irony". *Discourse Processes* 19: 347–367.

Economist. 2005. "High Pitch, Low Politics". *The Economist*. March 23. Accessed online.

Fear, J. 2007. "Under the Radar: Dog-Whistle Politics in Australia". *The Australia Institute*. Discussion Paper No. 96. September.

Farrell, J. and Gibbons, R. 1989. "Cheap Talk with Two Audiences". *The American Economic Review* 79 (5): 1214–1223.

Godfrey-Smith, P. and Martínez, M. 2013. "Communication and Common Interest". *PLOS Computational Biology* 9 (11): 1–6.

Goltsman, M. and Pavlov, G. 2011. "How to Talk to Multiple Audiences". *Games and Economic Behavior* 72: 100–122.

Grice, H. P. 1989. *Studies in the Way of Words*. Harvard: Harvard University Press.

Haney López, I. 2014. *Dog Whistle Politics: How Coded Racial Appeals Have Reinvented Racism and Wrecked the Middle Class*. Oxford: Oxford University Press.

Hornsby, J. 2000. "Feminism in Philosophy of Language". In Fricker, M and Hornsby, J. (eds.) *The Cambridge Companion to Feminism in Philosophy*. Cambridge: Cambridge University Press: 87–106.

Kahneman, D. 2013. *Thinking Fast and Slow*. New York: Farrar, Straus and Giroux.

King, M. P. 2008. *Sex and the City: The Movie*. (Writer, Director). Film. New Line Cinema.

Kruger, J. et. al. 2005. "Egocentrism Over E-mail: Can We Communicate as Well as We Think?" *Journal of Personality and Social Psychology* 89 (6): 925–936.

Lee, J. and Pinker, S. 2010. "Rationales for Indirect Speech: The Theory of the Strategic Speaker". *Psychological Review* 117 (3): 785–807.

Leibovich, M. 2014. "Mitt Isn't Ready to Call It Quits". *New York Times*. September 30. Accessed online.

Lumsden, D. 2008. "Kinds of Conversational Cooperation". *Journal of Pragmatics* 40: 1896–1908.

Marmor, A. 2011. "Can the Law Imply More Than It Says? On Some Pragmatic Aspects of Strategic Speech". In Marmor, A. and Soames, S. (eds.)

Philosophical Foundations of Language in the Law. Oxford: Oxford University Press: 83–104.

Neale, S. 1992. "Paul Grice and the Philosophy of Language". *Linguistics and Philosophy* 15 (5): 509–559.

Neale, S. 1999. "Colouring and Composition". In Stainton, R. (ed.) *Philosophy and Linguistics*. Boulder: Westview Press: 35–82.

Neale, S. Forthcoming. *Linguistic Pragmatism*. Oxford: Oxford University Press.

Nguyen, T. 2016. "The Two Trumps: 5 Signs Donald Trump Is Having a Full-Blown Identity Crisis". *Vanity Fair*. April 22. Accessed online.

Papineau, D. 2005. "The Cultural Origins of Cognitive Adaptations". In O'Hear, A. (ed.) *Biology, Evolution, and Culture*. Royal Institute of Philosophy Supplement No. 56. Cambridge: Cambridge University Press: 291–324.

Pinker, S. 2007a. *The Stuff of Thought: Language as a Window into Human Nature*. London: Penguin.

Pinker, S. 2007b. *Games People Play: Indirect Speech as a Window into Social Relationships*. Lecture. Presented at the Institute of Philosophy, London, UK. October 17, 2007. Video accessed online at: http://library.fora.tv/2007/10/15/Steven_Pinker_Games_People_Play.

Pinker 2011 "Indirect speech, politeness, deniability, and relationship negotiation Comment on Marina Terkourafi's 'The Puzzle of Indirect Speech'". *Journal of Pragmatics* 43: 2866–2868.

Pinker, S., Nowak, M. and Lee, J. 2008. "The Logic of Indirect Speech". *PNAS* 105 (3): 833–838.

Reboul, A. Forthcoming a. "Is Implicit Communication a Way to Escape Epistemic Vigilance?" In Avrassimapoulos, S. (ed.) *Pragmatics at Its Interface*. Berlin: Mouton.

Reboul, A. Forthcoming b. *Communication and Cognition in Language Evolution*. Oxford: Oxford University Press.

Sterelny, K. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Oxford: Wiley-Blackwell.

Sperber, D. and Wilson, D. 1989. *Relevance: Communication and Cognition*. Oxford: Blackwell.

Terkourafi, M. 2011a. "The Puzzle of Indirect Speech". *Journal of Pragmatics* 43: 2861–2865.

Terkourafi, M. 2011b. "Why Direct Speech Is Not a Natural Default: Rejoinder to Steven Pinker's 'Indirect Speech, Politeness, Deniability, and Relationship Negotiation". *Journal of Pragmatics* 43: 2869–2871.

Tomasello, M. 2010. *Origins of Human Communication*. Cambridge: Bradford.

Walton, J. et al. 2013. "Map: Which Country Pays the Most Bribes?". *BBC News*. Accessed online. July 9.

Croatian Journal of Philosophy
Vol. XVI, No. 47, 2016

Coordinating with Language

JESSICA KEISER
Yale University, New Haven, USA

Linguistic meaning is determined by use. But given the fact that any given expression can be used in a variety of ways, this claim marks where metasemantic inquiry begins rather than where it ends. It sets an agenda for the metasemantic project: to distinguish in a principled and explanatory way those uses that determine linguistic meaning from those that do not. The prevailing view (along with its various refinements), which privileges assertion, suffers from being at once overly liberal and overly idealized. By parsing the most prominent aims we use language to achieve, noting their relations of dependence and the specific type of uses they involve, I arrive at a novel metasemantic account: facts of linguistic meaning are determined by locutionary action.

Keywords: Metasemantics, convention, illocutionary acts, speech acts, meaning.

> *If L is a possible language, and P is a population of agents, what relation must hold between L and P in order to make it the case that L is an actual language of P? The obvious answer is: the members of P must use the language L. But that answer is neither informative nor clear. They must use L in a certain way. If everyone in P used L by telling lies in L or by singing operas in L (without understanding the words), they would be using L but not in the right way; L would not be their language.* David Lewis (1969: 176–177)

> *I completely agree with Wittgenstein as far as he goes on that; the trouble is that he stops just short of the problem. I also think that the meaning [of an expression] is the use… I think this is one sufficiently broad characterization to cover the whole lot; the problem is, it's too broad… so it's a question of what use counts and what use is irrelevant and should be dismissed.... and there's where the problem begins, not where it ends.* W.V. Quine (2013)

Metasemantics is the study of what grounds facts of linguistic meaning. While it is uncontroversial that these facts are determined by use, given the myriad ways that any given expression can be used, this claim cannot be taken as ending systematic metasemantic inquiry; rather, it specifies its point of departure. One fundamental task of the metasemantic project is to provide a principled account of the sorts of uses that are meaning-determining, distinguishing them from uses that are not.

An initially plausible proposal, some variant of which has been either explicitly endorsed or tacitly assumed by the vast majority of philosophers and linguists for decades or more, is that *assertion* grounds facts of linguistic meaning. However, this view—along with its various refinements—suffers at once from being overly idealized and overly liberal. It fails both to account for standard and ubiquitous ways of using language that don't fit the idealized model of cooperative information exchange, and to sift out the kinds of use that are not relevant to determining linguistic meaning.

I offer a novel metasemantic account: facts of linguistic meaning are determined by *locutionary* action. My strategy is to parse the most prominent aims we use language to achieve, noting their relations of dependence and the specific type of uses they involve. This allows me to identify uses that meaning-determining and distinguish them in a principled and explanatory way from uses that are not. The resulting account provides a suitably robust—yet sufficiently fine-grained—account of linguistic meaning, while shedding light on the diverse roles that different kinds of uses will have in an overall theory of communication. Moreover, the account yields a new foundation for traditional topics in semantics and pragmatics which cannot be fully successfully theorized about independently of metasemantic considerations.

1. *Coordination Problems and Convention*

We use language as a tool for achieving a number of different aims. Many of these aims are situated in an order of dependence, in the sense that success in achieving one aim is a prerequisite for success in achieving another. We use language in different ways, corresponding to the different aims we are trying to achieve; this can lead to confusion in metsemantic theorizing, since not all of these uses are relevant to the question of what grounds facts about linguistic meaning. It is crucial, then, to get clear on the way these different ways of using language are related, and at what level each should appear in a full account of linguistic communication.

I will be making the simplifying assumption that language use is conventional, that a population speaks a certain language just in case there is a convention among them of performing a certain type of action.[1] This simplifies the task guiding metasemantic inquiry to that

[1] Most of my arguments (with the exception of one I make in 2.4) will not rely on this assumption—rather, they concern which type of action should appear in

of identifying the relevant action type, which—I will claim—is locutionary action. The notions of convention and coordination problems can be found in standard introductory economics textbook, and will be familiar to scholars working in any discipline that makes use of game theory. In the absence of uncontroversial yet rigorous definitions of these notions, I will be utilizing Lewis' account of both throughout the paper. Though there have been quibbles over the finer points of Lewis' definitions, they should nonetheless suffice here to provide a clear framework within which to talk about conventional uses of language.

A *coordination problem*, according to Lewis, is a situation in which two or more agents have a common interest in performing the same one of several alternative actions (or beliefs).[2] They must coordinate in order to achieve a mutually desired outcome, or equilibrium state. Coordination problems can theoretically be solved by pure luck; it is in principle possible that agents could inadvertently reach an equilibrium state by accidentally slipping on a banana peel or falling into a ditch. In most cases, however, coordination problems are solved by reasoning about another agent's mental states, predicting her actions, and then choosing one's own actions accordingly. In this way, agents are able to find a stable solution to a recurrent coordination problem by using the same strategy repeatedly on the expectation that others will do the same; in other words, a *convention*:

> A regularity R in behavior is a convention for a population P if it's common knowledge among the members of P that they follow R, that they expect one another to follow R, that R is a solution to the coordination problem that they face in S, and that there is another solution R′ that they could have conformed to instead of R. (Lewis 1969: 76)

In what follows I will tease apart several important coordination problems that we use language to solve, which happen to be nested inside of each other in the sense that solving one problem is a prerequisite for solving another. I think of these nested coordination problems like lay-

an account of linguistic meaning, regardless of whether the notion of convention is employed in such an account.

[2] In *Convention*, Lewis defines coordination problems in terms of action, going on to define conventions as regularities in action that are solutions to coordination problems. In "Language and Languages" he changes his account of convention so that the regularity is one of action or belief but does not make the relevant adjustment to his definition of coordination problem, which simply does not come up in that paper. This change in his account of convention was needed to accommodate the fact that in "Language and Languages" Lewis modifies his earlier account of language use as a convention of truthfulness (which is an action) to that of a convention of truthfulness *and trust* (which is belief). In this paper I will be using Lewis' amended account of language use and convention, which are given in terms of action *and belief* rather than merely action. Though he did not explicitly make the corresponding changes in his account of coordination problems, I do so here and consider this to be in keeping with his later views. For ease of exposition, I will often talk about coordination problems and conventions in terms of action only, but it should be implicitly understood that belief is also included (Lewis 1969: 24).

ers of an onion—peeling back one layer exposes another coordination problem underneath. I will suggest that up until this point, theorists of language have failed to peel the onion back far enough. The task is to peel back the layers of irrelevant use, stopping just at the point where the kind of use we have identified is robust enough to do the job of grounding linguistic communication. In what follows, I will consider roughly three layers of coordination problems—not because they are exhaustive, but because the first two correspond to the main metasemantic theories that have either been explicitly proposed or tacitly assumed in the literature up until now. The third, I will argue, is the one that corresponds to the *correct* metasemantic theory.

2. *Coordinating with Language*

Before teasing apart these different coordination problems and examining them more carefully, I'll first set up a bit of terminology. I will be using a broadly Austinian taxonomy of speech acts: The *locutionary* act is—roughly for now, but to be precisified later—to perform an utterance simply in order to transmit its conventionally encoded content (Austin 1975: 94–107).[3, 4] The locutionary act is successful when the audience entertains the content that the utterer has in mind. But merely transmitting a content rarely serves our goals—we often want to signal that we take some sort of attitude toward the content. Perhaps we want the audience to think that we believe a proposition, for instance, or to understand that we desire her to answer a question. So the locutionary act is generally in service of an *illocutionary* act; the performance of an utterance with some kind of force, such as assertion, command, query, etc. The illocutionary act is successful when the audience understands not only the content the speaker transmits, but the force that speaker attaches to that content. The illocutionary act, in turn, is usually performed in the service of the *perlocutionary* act; roughly, the act of doing something *by* uttering an expression, such as getting someone to close a door or to share one's beliefs, etc. The perlocutionary act is successful if the audience responds in the appropriate/desired way to the illocutionary act of the speaker, i.e., by producing the intended perlocutionary effect.

2.1. *Coordinating illocutionary act with perlocutionary effect*

Arguably, the aim of most ordinary language tends to be the production of perlocutionary effects—such as belief, action, and testimony. We navigate the world better when we are able to exchange information and

[3] For ease of exposition I will ignore context sensitivity here and make the simplifying assumption that the meanings of sentences are complete thoughts rather than functions.

[4] Some readers may find my interpretation of Austin to be controversial. While I believe I have provided a reasonable regimentation of Austin's taxonomy of speech acts, I don't have the space to defend that claim here. However, the fidelity of my interpretation of Austin should not be relevant to the arguments in this paper.

coordinate our behavior; it is hard to imagine how a community would have developed the use of language if not from pressure arising from this need.[5] This may explain early theorists' attraction to a metasemantic story in which perlocutionary effects play a primary role in linguistic convention. David Lewis' theory is a paradigm example of such an account: he worked out a rigorous metasemantic proposal, the general spirit of which is perhaps tacitly assumed by theorists who take the potential perlocutionary effect of an expression (e.g., change in context) to be indicative of its semantic value (e.g., context change potential). According to Lewis, parties in a conversation have a mutual interest in one member performing a certain perlocutionary response conditional on a certain state of affairs obtaining. They are able to achieve an equilibrium state if the utterer performs her illocutionary action only under certain circumstances (for instance, she will assert only true things, and command only actions that—if performed—will serve the interlocutors' relevant mutual interest) and the audience coordinates her perlocutionary response with the action of the utterer (by believing the content of the assertion, or obeying the command).[6] For instance, perhaps A and U ultimately have a mutual interest in both of them being able to unlock the house, though there is only one set of keys; then they will also share the more immediate interest for A to *believe* that the housekeys are hidden in the birdfeeder conditional on its being the case. Or for A to *hide* the housekeys in the birdfeeder conditional on its being the case that U will look for them there. Therefore, U will assert that the housekeys are hidden in the birdfeeder just in case they are, and she will command A to hide them there just in case that's where she is planning to look for them. In turn, A will respond to U's assertion that the housekeys are hidden in the birdfeeder by believing it, and will respond to her command to hide them there by doing it. Coordinating their action in this way serves their mutual interests: it allows them to navigate the world by exchanging information and acting together to achieve mutual goals.

While the motivation for this picture is clear, it does not provide a satisfactory basis for a theory of linguistic communication. As Austin

[5] I am uncommitted with respect to what was in fact the causal history of actual language use. Though, for instance, Noam Chomsky denies that language evolved as a response to a need for communication or information exchange, I take it to be uncontroversial that this is a primary function of language and strongly influences the development of individual languages over time. Moreover, in downplaying the conventionality of language, Chomsky is focused on syntax rather than semantics or pragmatics.

[6] Lewis uses a technical account of truthfulness and trust that is designed to describe the coordination between conversational participants in a way that generalizes across illocutionary act types. In each case one party is being truthful and one party is being trusting (in ways that will be cashed out differently depending on the type of illocutionary act). I will not go into the details of this account; what matters to the discussion here is that in each case the coordinating action on the part of the audience will be to perform a perlocutionary response (to believe the content of the assertion, fulfill the request, answer the question, etc.)

and Grice later started to point out, we can put language to many uses, not all of which determine conventional meaning or are even relevant to communication. What speaks decisively against the particular picture we are considering here is Austin's insight that communication has been successful when the hearer has identified the illocutionary action, regardless of whether she performs the corresponding perlocutionary response. If she knows what has been asserted, there has been successful communication regardless of whether she believes it. If she knows what has been commanded, there has been successful communication regardless of whether she does it. If she knows what has been asked, there has been successful communication regardless of whether she answers it, and so on. Furthermore, it seems that communication takes place even in situations in which we are actively and openly being *uncooperative* at the level which involves perlocutionary uptake.[7] So the lesson gleaned from Austin and Grice is that the layer of coordination involving perlocutionary action is irrelevant. The natural next step would be—and indeed *was*—to peel the onion back a layer and look instead to the problem of identifying the illocutionary act as a candidate for grounding language use. Something like this type of picture remains to be predominantly assumed in the literature today, though—as I will argue below—it is does not go far enough in stripping away irrelevant use.

2.2. *Coordination on identifying illocutionary act*

In order to coordinate her perlocutionary response with the speaker's illocutionary act, the audience must first identify which illocutionary act has been performed. Before forming a belief, for instance, she must understand what the speaker has asserted. Before answering she must understand the question and before obeying she must understand the command. This poses a separate coordination problem; the speaker and audience both want the audience to understand which illocutionary act the speaker is performing. They can solve the problem by relying on expectations and beliefs about the other's behavior: the speaker will generally encode her illocutionary act based on her expectations about how the audience will respond to her utterance, and the audience will generally form her beliefs about the speaker's illocutionary act based on her expectations of how the speaker will encode it. This is a coordination problem nested within a coordination problem, since it is only *after* the audience has successfully identified the speaker's illocutionary act that she can respond appropriately. But as we have seen, it seems as though successful communication occurs when the illocutionary act is identified, regardless of whether the audience responds in the appropriate or desired way. Indeed, agents need not even have an *interest* coordinating at the perlocutionary level in order to communicate suc-

[7] See Camp (forthcoming), McKinney (manuscript), and Pinker, Nowak, and Lee (2008).

cessfully. Such considerations make it appealing to look instead to the problem of identifying the illocutionary act to do the work of grounding linguistic meaning.

This general kind of picture which grounds meaning in illocutionary action is reflected in Grice and those following a similar framework, such as Stephen Schiffer and Brian Loar.[8] Differences in terminology can be misleading here: these philosophers took linguistic meaning to be grounded in acts of speaker meaning, but because they define speaker meaning in terms of a manifest intention to produce a propositional attitude (or other type of response) in the audience, it falls into the category of what I am calling *illocutionary action*, rather than what I am calling *speaker meaning*. Though Grice and his followers do not explicitly talk about the identification of the illocutionary action as a coordination problem, this way of framing things fits quite naturally with a Gricean account of communication. One of the key features of Grice's notion of speaker meaning, which sets it apart from natural meaning, is the involvement of coordination in choosing and interpreting the utterance. When a speaker *means p* she succeeds in communicating not simply by *causing* a belief, but by performing her utterance in a way that is meant to reveal her mental states, and to thereby provide the basis for the audience's reaction. In turn, the audience identifies the speaker's meaning not purely on the basis of external factors, but on her recognition of the speaker's intentions in producing that utterance, which serves as the basis for her response. The aspect of Grice's account of speaker meaning which makes it a fundamentally cooperative act is the requirement that the speaker intends for her communicative intention to be manifest to the audience—the recognition of which is meant to provide the audience a reason for identifying the speaker's illocutionary act one way rather than another. This kind of manifest, or reflexive intention functions as a mechanism for coordination. Though it is framed in a different way, the idea that communication fundamentally involves coordination by reasoning about each other's mental states and acting accordingly is as central to Grice's theory as it is to Lewis'. I take the general Gricean framework, then, to be a paradigmatic case of a metasemantic picture according to which language is grounded in conventions of illocutionary action.[9] Something in the vicinity also seems to be tacitly assumed by semanticists who posit illocutionary update potential as semantic values of expressions, as well as those who take illocutionary force to be relevant to locating the semantics/pragmatics boundary. As far as I can tell, this general type of picture seems to be widely taken for granted, but I will argue below that it is untenable.

[8] See Grice (1989), Loar (1981), Schiffer (1972).

[9] Grice had hesitations about appealing to convention in his account of linguistic meaning, and instead talked about members of a population having a certain procedure in their repertoires; however, because his account still grounds linguistic meaning in illocutionary action, the difficulties I outline here still apply (Grice 1989: 123–127).

This general approach runs into problems with something that is sometimes (somewhat disingenuously) called "non-standard speech"—specifically, cases in which language is used to illocute content that is non-identical to the standing meaning of the utterance, as well as cases in which language is used to do something else entirely from performing illocutions. Lewis was aware of these difficulties, which apply to his account as well (insofar as it grounds linguistic meaning *at least in part* in illocutionary action); he considers both cases in *Languages and Language*, though I will argue that his responses are inadequate. The two underlying problems are that linguistic meaning underdetermines illocutionary force, and that—like perlocutionary effect—illocutionary action does not seem to be necessary for communication or the use of language; these problems are intractable for a theory that seeks to ground linguistic meaning in conventions of illocutionary action.[10]

The first problem is that it seems as though there could be linguistic communities who speak non-literally more often than not; they primarily use sarcasm, metaphor, etc. According to the picture that grounds linguistic meaning in conventions of illocutionary action, however, such a community would not speak the language *L* that we'd naturally take them to be speaking; there would not be a convention among this population of using *s* to mean *L(s)*. Lewis responds by suggesting that the language we'd initially take to be the language of this population (call it *literal-L*) is only a simplified approximation to their actual language *L*, which may be obtained by specifying certain systematic departures from *literal-L*. (A simplified example: if *s* means *p* in *literal-L*, and is systematically used ironically by this population, then *s* means *not p* in *L* (Lewis 1975: 183)). The language of this community is determined by conventions of illocution, but they just happen to be speaking a different language than they initially appear to be.

But this response won't do: the most serious problem is that Lewis's solution assumes that the non-literal speech of this community will display enough systematicity to be able to derive *L* from *Literal-L*. However, this is at best extremely unlikely; it is certainly not necessary. Non-literal speech comes in many forms, and some of them—like metaphor—tend to be highly unsystematic; while we may be able to characterize the general features of different kinds of non-literal speech, there is no principled mapping from the meaning of a sentence *s* in Literal-English, for instance, to non-literal-English(*s*). The other issue is that the account rules out the possibility that for instance, Americans would share a language with the British in the case that the British started using irony to a sufficiently high degree—this seems like the wrong result.

[10] The objections can be modified to fit any theory that grounds linguistic meaning in illocutionary action, regardless of which/how many illocutionary acts are utilized in that theory. Furthermore, the objections will apply regardless of differences in the details of how these illocutionary acts are defined, as long as they meet the criteria outlined above in that they involve the speaker having a communicative intention over and above that of simply transmitting a content.

The second problem is that it seems that there could be linguistic communities who were not in the business of performing illocutionary actions at all; rather they may use a language *L* uniquely for the purpose of putting on plays, or telling jokes or stories.[11] If a population's language is determined by their conventions of illocution, this population does not count as speaking a language at all. Lewis' initial response here is analogous that given above; the language of this population *L* may be obtained by specifying their systematic departures from *literal-L* (which we would naturally identify as their language). In this case, however, Lewis recognizes that there may not be the requisite systematicity; he suggests an alternative solution, which is to restrict the account of the convention which governs language use to serious communication situations, which he defines as follows:

> We may say that a *serious communication situation* exists with respect to a sentence *s* of *L* whenever it is true, and common knowledge between a speaker and a hearer, that (a) the speaker does, and the hearer does not, know whether *s* is true in *L*; (b) the hearer wants to know; (c) the speaker wants the hearer to know; and (d) neither the speaker or the hearer has other (comparably strong) desires as to whether or not the speaker utters *s*. (Lewis 1975: 184)

The account would then be modified so that what determines which language a population speaks will be conventions of illocutionary action in *serious communication situations*. But again, this rules out plausible cases of shared language: we can imagine that a language was once used in a community to assert what were then taken to be religious truths, but which was later used to convey the *very same* stories, but this time as myths or parables. This situation does not reflect a change in language, but rather a change in religious belief. Perhaps more importantly, however, this restriction only serves to make salient the fact that serious communication situations are not the only contexts in which we use language. Information exchange is only one of many

[11] When joking around or telling stories, the speaker means something by her utterances—she intends her audience to attend to certain thoughts. Granted she is not *merely* intending to transmit these thoughts to the hearer; she will most likely have further intentions, perhaps including the intention to amuse her audience. However, this kind of further intention does not constitute an illocutionary act, since it is not a reflexive communicative intention. The speaker may intend that her audience be amused by the thought she has transmitted, but the basis for that amusement won't involve a recognition of the speaker's mental states (the fact that the hearer has these mental states will not typically be amusing); the source of amusement will be the transmitted thought itself. There need not be any coordination between the speaker and audience after the locutionary act has been identified. Grice touches on this point in *Utterer's Meaning and Intentions*, pointing out that while the audience's amusement may be partially *caused* by the recognition of the speaker's intentions (though I doubt that even this much is usually the case) it will not *be on the basis* of this recognition: "But though A's thought that U intended him to be amused might be a part-cause of his being amused, it could not be a part of his reason for being amused (one does not, indeed, have reasons for being amused)..." (Grice 1989: 92–3).

activities which we use language to facilitate. Our use of language is complicated and messy; we joke, we play, we say things we mean as well as things we don't mean—we often do these things unsystematically, and in the course of a single conversation. It is one thing to idealize away from complexities when constructing a model of language and language use, it is another thing to ignore or accept patently incorrect results for large swaths of perfectly ordinary use. Non-literal and "non-serious" speech are not peripheral cases; they are the norm, which any viable meta-semantic theory will need to account for.[12]

In the introduction I noted that metasemantic considerations will have significant implications concerning the viability of individual semantic theories. For instance, underlying any semantic theory which posits context change potential as semantic value seems to be the tacit assumption that meaning is in some way determined by illocutionary action or perlocutionary effect. This is because an update in the context, which can be modeled as the information state of the conversational participants, results from the recognition of the illocutionary action—and in some cases the subsequent production of the perlocutionary effect—in serious communication situations. For instance, assertion is standardly taken to effect a change in context when its intended perlocutionary effect (belief or acceptance of the asserted content by conversational participants) is achieved. In the case of questions and commands, the update in the context is not generally taken to require something quite strong as the production of the intended perlocutionary effect (an answer to the question or the performance the commanded action). However, it does involve recognition of the illocutionary action and a somewhat weaker response—for instance, the relevant conversationalists taking on the goal of answering the question or performing the commanded action. Given the underlying assumption that use determines meaning, semantic theories that posit context change potential as semantic values tacitly endorse the additional assumption that the *kind* of use that grounds linguistic meaning involves illocutionary action/perlocutionary effect. However, if the arguments above are correct and neither illocutionary action nor perlocutionary effect are necessary for language use, then it is hard to see what kind of metasemantics could support such a semantic theory.

2.3. *Coordination in identifying locutionary act*

So far we have considered two different coordination problems that language can be used to solve, one nested inside the other. The first was that of coordinating illocutionary action with perlocutionary response.

[12] In order to get around these problems one might try to incorporate mock illocutionary acts somehow (perhaps with a disjunctive account according to which meaning is grounded in conventions of illocutions or mock illocutions). I find little to recommend this approach. It is shamelessly *ad hoc*, and it generates the problem of giving a rigorous account of what mock illocutions amount to, exactly.

We saw that in order to reach the equilibrium state with this particular coordination problem, agents must first solve a different coordination problem: that of identifying the illocutionary act. But we can peel back yet another layer to find additional coordination problems nested within: that of identifying speaker meaning—and, where there is more than one level of meaning—identifying the speaker's *direct* meaning, which I will call the 'locutionary act'. I will argue that this is the coordination problem underlying language use. Though I am offering a novel metasemantic proposal, it is compatible (to varying degrees) with views about language and language use going back to philosophers like Locke, Austin and Strawson. Insofar as Locke considered language to signify mere contents of thought (devoid of illocutionary force), his view would naturally fit together with a metasemantic picture according to which language use is grounded in locutionary action, rather than illocutionary acts or perlocutionary effects. Austin thought that illocutionary acts were conventional, but that they were underdetermined by meaning; this seems to imply that the conventions of illocutionary action must be over and above those that determine meaning, at least in some cases. Strawson echoes this sentiment, agreeing with Austin that "we must refer. . . to linguistic conventions to determine what locutionary act has been performed in the making of an utterance, to determine what the meaning of the utterance is" but denying that illocutionary acts are always conventional (Strawson 1964: 442).

Given that it is possible for there to be language communities that employ language to different ends, it is important to identify what is common to all of these communities; I want to suggest that this is the act of speaker-meaning, which—to remind the reader—I define as follows:

> *Meaning*: By uttering *e*, U meant *m* iff for some audience A, U uttered *e* R-intending that A attend to *m* at least partially on the basis of her utterance.

The expression 'reflexive intention' is borrowed from Bach and Harnish (1979), and refers to intentions that are self-referential in the following way: a speaker has a reflexive communicative intention just in case part of her intention is that the audience recognize the full contents of that very intention. The speaker intends for her communicative intention to be manifest to the audience, the recognition of which is meant to provide the audience a reason for identifying the speaker's meaning one way rather than another. This definition of speaker meaning is adapted from Grice's account, but it differs crucially from his (as well as from others working within a broadly Gricean framework) in that it is less restrictive; a speaker means something merely by having the intention that her audience attend to a certain thought. The audience need not be intended to have any additional response, or to form any attitude toward the content over and above that of having it in mind.[13]

[13] Stephen Neale suggests a similar weakening in Neale (1992) p. 34–37. However, because he does not accept the implications that I take to follow from it—

If the speaker manages to get her audience to recognize her intentions and attend to the content that she has in mind, then this most basic of communicative acts—that of speaker meaning—has been successful. In contrast, more traditionally Gricean accounts of speaker meaning essentially involve some sort of additional communicative intention over and above what I am calling 'speaker meaning'—the intention that the audience do something with the content of the utterance after she has attended to it (such as to believe it, provide an answer to it, etc.) These additional communicative intentions constitute the illocutionary force of the utterance.

Of course, given that we have an interest in passing on information and coordinating action, speakers will often have communicative intentions that go above and beyond that of speaker-meaning, and instead constitute full illocutionary action. For instance, a speaker may utter a sentence not only with the intention that her audience *attend* to the proposition expressed by that sentence—but with the additional intention that the hearer recognize that the speaker believes the proposition, thereby coming to have reason to believe it herself. However, the performance and success of the act of speaker meaning are preconditions for the performance and success of an illocutionary act; so given that the agents have a shared interest in identifying the illocutionary act, they will have a shared interest in identifying the speaker's meaning in this more liberal sense. Here again we have a coordination problem nested within a coordination problem; in order to coordinate on solving the problem of identifying the illocutionary act, agents must coordinate on solving the problem of identifying what the speaker meant.

If we take the relevant coordination problem to be that of identifying speaker meaning, we get the desired results with respect to communities that do things with language other than perform illocutions. Paradigm cases in which a speaker means something by an utterance without performing an illocutionary act are conversational contexts in which the speaker is joking around or telling stories. In these contexts the speaker is not *merely* intending to transmit these thoughts to the hearer; she will most likely have further intentions, perhaps including the intention to amuse her audience. However, these further intentions do not constitute illocutionary action, since the source of amusement is typically not the recognition of the speaker's mental states, but rather the audience's impulsive reaction to her understanding of the thought being communicated. So while in this kind of situation there is no coordination needed to identify an illocutionary action, there is still coordination needed to identify the speaker's meaning. Since in non-serious communication situations speakers still mean what they say in the liberal sense that I've defined, by taking the language of a community to

for instance, that the literal content of indirect speech will be *meant* by the speaker this weakened sense—I assume he has in mind a more restricted account than the one I am proposing here.

be determined by conventions of speaker meaning, we will get the right results with respect to communities that use language solely in non-serious communication situations.

However, this does not provide the tools to deal with the other counterexample, which essentially involved indirection. In the example we considered, the indirection was with respect to the *illocutionary* acts being performed; that is, members of the hypothetical community used non-literal speech more often than not (in that, roughly, the content of their illocution tended to be non-identical to the semantic content of their utterance). But the example could easily be modified to create a problem for an account which takes the relevant language convention to be that of speaker-meaning. Indirection and non-seriousness are orthogonal issues; we could just as well consider a community of people who only joke around and tell stories, yet do all of this using indirect speech. In this kind of case, the speaker-meaning account will get the wrong results in regard to which language is the language of this population.

It is necessary to impose constraints on the account in order that the contents of implicature, ironical and metaphorical utterances and the like, do not get swept up by the theory—this can be done by restricting the account to *direct* meaning. Not all acts of speaker-meaning are on a par. With any given utterance, a speaker may perform multiple acts of speaker meaning; however she may use varying strategies to facilitate the hearer's recognition of the different things that she means. For a subset of the contents meant by the speaker, their communication will be linguistically mediated in a more direct sense than other contents that she meant in the course of that same utterance. That is, in the case that a speaker means multiple things with an utterance, there will be one content that she intends the speaker to entertain primarily on the basis of her utterance (together with shared background information), while there will be other contents that the hearer is intended to arrive at on the basis of other contents she recognizes to have been meant by the speaker. We can think of this subset as what is directly meant by the speaker. In the case that the speaker is part of a community with established linguistic conventions, what she directly means by her utterance will likely correspond with the content of what people have tended to think of as the locutionary act she performs with that utterance, in the sense that it will correspond to the conventional meaning of the expression she utters.[14, 15] In contrast, additional contents she

[14] This need not be the case. But the reason that it will often be the case is that usually the most efficient way for the utterer to get her audience to identify a certain content is to utter a sentence whose standing meaning is that very content.

[15] Again, for simplicity I am ignoring context sensitivity here, making the simplifying assumption that the meanings of sentences are complete thoughts rather than functions. In any case, the content of the locutionary act will generally bear some close relation to the conventional meaning of the utterance, though it may not be identity.

might have meant (for instance, those that she implicated with the utterance) neither correspond to the conventional meaning of the expression she used, nor were they directly meant. Rather, the hearer is intended to arrive at these additional meanings at least partially on the basis of other things meant by the speaker in that utterance. I will define locutionary action as direct meaning.

> *Indirect Meaning*: By uttering *e*, U indirectly meant *m* iff for some audience A, there is some content *k* (not identical to *m*) such that
>
> 1. By uttering *e*, U meant *k*
> 2. By uttering *e*, U meant *m*
> 3. U uttered *e* R-intending that A recognize (2) at least partially on the basis of (1).
>
> *Direct Meaning:* Whatever is meant: if it is not meant indirectly, is it meant directly.

My definition of locutionary action, then, differs from the traditional Austinian characterization in that the latter but not the former requires that there be some conventional meaning associated with speakers' utterance.

Again we find a coordination problem within a coordination problem: in the case of indirect speech, in order to identify the indirectly meant content *q*, the hearer must first identify that the speaker meant *p*, since it is on the basis of this fact that she is meant to identify the indirect content *q*. There can be many layers of meaning, but in each case the hearer must identify the content that is on the bottom-most level—the locuted content. Identification of the locutionary act, then, will be a part of any coordination problem of identifying speaker meaning. In the case that there is only one level of meaning, it will be the very same problem. In the case of identifying indirect speaker meaning, it will be a coordination problem within a coordination problem.

If we take the relevant coordination problem to be that of identifying the locutionary act of the speaker we get the right results not only with regard to the community that uses language only in non-serious communication situations, but also in regard to the community that uses indirect speech more often than not. But it is not simply a matter of making the right predicitions; the fact that this account manages to get the desired results in cases where the more traditional accounts failed is explanatory in that it is reflective of two broader facts about language use: (1) that locutionary action is common to all communication situations, and (2) that the locutionary act of the speaker is part of the mechanism by which indirect speech occurs—even when it not the "point" of the utterance, it is crucially involved in the communicative act. By grounding language use in locutionary action we peel back superfluous layers of use—following the lead of Austin and Grice, but continuing where they had stopped too soon—to find what is both com-

mon to all contexts of linguistic communication and robust enough to do the work of grounding linguistic meaning: locutionary action.

2.4. *Why not conventions of illocution as well as locution?*

As we have seen, identifying the speaker's illocutionary act involves coordination on the part of the speaker and her audience. Conversational participants use language to solve this coordination problem, in that the speaker chooses an utterance that she expects will provide the audience a way to identify her illocutionary act, and the audience will in turn identify the illocutionary act based on her expectations of the speaker. Surely, then, there can be conventions of illocutionary action; there is in principle no reason why a regularity could not develop within a community of using s to illocute L(s) and of taking others to do so. So why not take this convention—*in addition to* conventions of locutionary action—as relevant to metasemantic theorizing? We could get around the counterexample involving non-serious speech situations by introducing a disjunction into the definition: instead of the account I've offered—in which a population speaks a language L just in case they have a convention of using s to locute $L(s)$—we could say instead that a population speaks a language L just in case they have a convention of using s to locute $L(s)$ and/or using s to (directly) illocute L(s). The counterexample involving indirection could be handled by modifying the account to impose a directness condition on the illocutionary action involved. While on the face of it, it might look like this picture would get the right results, this strategy should be resisted for the following reasons:

First, what looks like a convention may in fact just be correlation. According to an old picture of sentential force, different kinds of sentence types all have the same kind of content; the difference in perceived meaning does not amount to a difference in content, but rather a difference in force. According to this picture it would be somewhat arbitrary which kind of force could be paired with a particular content on an occasion of utterance, given that there is only one *kind* of content (traditionally supposed to be propositional) which is compatible with each type of force. On any occasion of utterance, there would be alternatives; the speaker can choose the content and the corresponding force based on her interests and goals on a given occasion. However, this traditional Fregean view—according to which expressions with different sentential force share the same type of content—has been shown to be untenable. Because these expressions can embed, the strategy of taking them to have a uniform kind of content leads to difficulties in providing a systematic compositional semantics, and has been largely abandoned as a result (Groenendijk and Stokhof 1997: 1061–1075). It has been more promising to posit different kinds of contents corresponding to different kinds of sentences. If this strategy is correct, and indicatives, queries, commands, etc. have different types

of entities as their semantic contents, then the correlation we see between illocutionary action and sentence type might just be piggybacking on the convention of locution (which concerns merely content). For instance, we have conventions of locutions for indicative sentences that determine their contents to be propositions (or propositional functions); while indicatives may also correlate with the illocutionary act of assertion, it is not because there is a separate convention involved. Rather it is because of the nature of assertion; one can only assert contents that are propositions (and not, for instance, questions). *Mutatis mutandis* for queries with questions (semantic content of interrogatives), as well as requests with whatever kind of entity turns out to be the semantic content of imperatives. If my metasemantic proposal is correct, then in cases of direct illocutionary action we would expect to see the above correspondence. However, this correspondence would not indicate a convention because there would be no alternative; when the direct meaning is propositional, only assertion will be an option for the direct illocutionary act, when the direct meaning is a question, only a query will be an option for the direct illocutionary act, and so forth.

But let's suppose for the sake of argument that there could be conventions of direct illocutionary action. Even in this case I think that there are reasons—gleaned from the consideration of the counterexamples in *2.2*—to exclude these conventions from an account of what it is for a population to use a language. The first reason is that not all language communities need conventions of illocutionary action, though they all need conventions of locutionary action. Ideally, we'd want a simple theory that would generalize to all language speaking communities and all languages. The second is that there is an intuitive sense that two communities that share conventions of locution but not of illocution may share the same language. This seems to be evident from the hypothetical situation described above, in which Americans for the most part speak literally while the British primarily use non-literal speech. What they have in common is a convention of locution; they both use the same utterances to directly mean the same contents. They differ in their conventions of illocutionary force; whereas Americans directly assert (for instance) the locuted content, the British assert a different, indirectly communicated content. Yet there is the intuition in this case that even so, the two communities are speaking the same language. The same could be said if we were to compare a community that used language in serious speech situations in order to perform illocutionary actions, and one that did not—for instance, it seems as though *the very same language* that was used by one community to assert what they took to be religious truths (i.e., to perform illocutions), could later be used by another to convey the very same stories, but this time as myths or parables (i.e., to do something other than perform illocutions). These communities, again, would share the same conventions of locution but not the same conventions of illocution—and yet there is the intuition that they could share a language.

Acknowledgements

Many thanks to Zoltán Szabó, Jason Stanley, Liz Camp, Larry Horn, and the participants of the *Philosophy of Linguistics and Language Conference* (*Dubrovnik, 2015)* for helpful comments and feedback.

References

Austin, J. L. 1975. *How to Do Things with Words*. Oxford: Clarendon Press.

Camp, E. (forthcoming). "Insinuation, Indirection, and the Conversational Score." In D. Harris and M. Moss (eds.). *New Work in Speech Acts*. Oxford: Oxford University Press.

Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.

Groenendijk, J. and Stokhof, M. 1997. "Questions." In J. van Benthem and A. ter Meulen (eds.). *Handbook of Logic and Language*. Amsterdam: Elsevier: 1055–1124.

Lewis, D. 1969. *Convention: A Philosophical Study*. Cambridge: Harvard University Press.

Lewis, D. 1975. "Languages and language." In Keith Gunderson (ed.). *Minnesota Studies in the Philosophy of Science*, vol. VII. Minneapolis: University of Minnesota Press: 3–35.

Loar, B. 1981. *Mind and Meaning*. Cambridge: Cambridge University Press.

McKinney, R. Manuscript. Discursive Antagonism and the Value of Conversation.

Pinker, S., Nowak, M., and Lee, J. 2008. "The Logic of Indirect Speech." *PNAS* 105 (3): 833–838.

Schiffer, S. R. 1972. *Meaning*. Oxford: Clarendon Press.

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

Quine, W.V. 2013. *In Conversation with W.V. Quine – The Boolos Panel.* https://www.youtube.com/watch?v=2NmVVu7gF9c

Croatian Journal of Philosophy
Vol. XVI, No. 47, 2016

Superficially and Deeply Contingent A Priori Truths

MARCO RUFFINO
University of Campinas (UNICAMP) / CNPq, Brazil

In this paper, I review some standard approaches to the cases of contingent a priori truths that emerge from Kripke's (1980) discussion of proper names and Kaplan's (1989) theory of indexicals. In particular, I discuss Evans' (1979) distinction between superficially and deeply contingent truths. I shall raise doubts about Evans' strategy in general, and also about the roots and meaningfulness of the distinction.

Keywords: Contingent a priori truths, rigid designation, contingency, descriptive names, Gareth Evans.

Kripke's discussion (1980) of proper names and rigidity resulted in some well-known examples of truths that are putatively contingent and, nevertheless, can be known a priori by speakers in the appropriate circumstances (contingent a priori truths, or simply "CATs" from now on). Since then there has been an intense debate concerning the semantic, epistemic and the metaphysical status of these truths. Kaplan's later work on demonstratives (1989) presented some analogous examples of this kind of truths, but this time involving demonstratives and pure indexicals. Kaplan's examples also generated a discussion not only concerning the epistemic and metaphysical status of these truths, but also of the notion of proper contexts of utterances.

Many critics of Kripke tend to focus on the origin of his examples, i.e., his theory of proper names and direct reference. Some of them (e.g., Dummett (1973) and Hawthorne and Manley (2010)) take Kripke's example as a sort of *reductio* of the very thesis that names are directly referential. Others explain the phenomenon appealing to some natural features of language (e.g. Kripke, Kaplan). But the most prevailing line of approach is the one in which there is some weakening of the credentials of Kripke's cases either as genuinely contingent (Donnellan (1977), Evans (1979), Hawthorne (2002)) or as genuinely a priori (e.g., Soames (2003, 2005), Salmon (1986), Plantinga (1975)).

It seems to me that each of these approaches is, in one way or another, problematic for different reasons. In this paper I shall concentrate on one particular approach outlined by Evans (1979) and taken up by Hawthorne (2002). Since this approach is to a degree scale a reaction against Donnellan's ideas on the same issue, I start by reviewing the basics of his treatment. As I shall argue, Evans' approach has some serious drawbacks.

1. *CATs*

Kripke famously advocated a sharp distinction between metaphysical and epistemic modalities. He criticized both the thesis that all necessary truths are knowable a priori and the dual thesis that all contingent truths are knowable only a posteriori. Both theses, according to him, derive from the confusion between metaphysical and epistemic features of propositions (modalities). A proposition is necessary if it is true in all possible worlds, and contingent if true at some possible world and false at some other possible world. Hence, necessity is a metaphysical feature of some propositions. A proposition is a priori if it can be known independently of any empirical experience, and a posteriori otherwise. Hence, a priority is an epistemic feature of some propositions. It follows that the pairs a priori/necessary and a posteriori/contingent are not intensionally equivalent. But are they extensionally equivalent? Part of Kripke's discussion is an attempt to show that, as a result of the semantic phenomenon of rigidity together with the fact that descriptions might be used to fix the reference of names without thereby becoming their synonym, some contingent truths can be known a priori, and some necessary truths can only be known a posteriori. (We shall concentrate on the first kind of truths in this paper.) Kripke offers two celebrated examples:

i. Standard Meter Bar Case: The term 'meter' (which for all purposes can be treated as a proper name of a length unit) was historically introduced by a person or group of persons (the "baptizer") as naming the length of a certain standard platinum bar (call it S). Since the length of a metal bar varies with time, let us consider 't_0' as the exact instant in which the term was introduced. The baptizer was in a position to know that the following sentence is true without any relevant experience (i.e., without having to effectively measure the bar):

(M) The length of S at 't_0' is one meter.

(M) is a true identity sentence that can be known a priori. This is so because both sides of the identity sentence refer to the same length in the actual world (since the name 'one meter' was stipulated to refer to the same object as 'the length of S at 't_0"). However, 'one meter' is a rigid designator, which means that it designates the same length in all possible worlds, while 'the length of S at 't_0" is a non-rigid designator, which means that it might designate different lengths in differ-

ent worlds. Hence, although true in the actual world, (M) is false in a world in which the bar is not one meter long at 't_0'. It follows that (M) expresses a contingent truth.

ii. Neptune Case: In 1846 the French astronomer Leverrier, after carefully studying some small perturbations in the orbit of Uranus, was led to believe that there was a new and so far unobserved planet that should be the cause of such perturbations, and predicted the future position of it. He baptized this planet 'Neptune', that is to say, he gave the name 'Neptune' to the object corresponding to the description 'the celestial body that causes the perturbation in Uranus' orbit' (if there is actually one such celestial body). Only a couple of days after communicating his research to the Berlin Observatory, Neptune was effectively observed with a telescope for the first time. Hence, even before observing Neptune (and before having decisive evidence that the thing seen in the telescope was responsible for the perturbations), Leverrier was in an epistemic condition to know the truth of

> Neptune is the planet that causes the perturbation in Uranus' orbit

Actually, since there could be no planet that corresponds to the description (Leverrier himself postulated another planet called 'Vulcan' that, as it was later discovered, does not exist), this sentence could express no proposition at all if Neptune does not exist. Hence, what Leverrier really knew a priori was the conditional

> (N) If there is one and only one cause of the perturbation in Uranus' orbit, then Neptune causes the perturbation in Uranus' orbit.

Leverrier knew (N) a priori in the sense that no empirical experience (i.e., astronomical observation) was needed to know that it is true, since it came as a result of his original stipulation regarding the name 'Neptune'. One might thing that a great amount of empirical observation was necessary for Leverrier to come to postulate that there is such a planet (and hence that the knowledge of (N) cannot be a priori), but strictly speaking the experience is not necessary to know (N) but only to know that its antecedent (existential claim) is true. The conditional as a whole can be known simply as a result of the stipulation.

2. *Indexicals*

Kaplan (1989) famously distinguishes between two kinds of meanings that indexicals have, namely, the character and the content. The content is the extension (or intension, depending on the perspective) assumed by an indexical in each context of utterance, while the character is a rule (or a function) that associates an appropriate extension (or intension) to each context. The content of an indexical might change from context to context, but the character remains fixed, and it is usually identified with the meaning that a speaker understands independently

of the context of use. (Two occurrences of 'today' on two distinct days have the same character but two distinct contents.) Kaplan also distinguishes the cognitive significance of a sentence containing indexicals from the object of thought associated with it. The bearer of cognitive significance is the character, while the object of thought is the content (which is a proposition). Two persons thinking 'I am here today' in two distinct days are thinking of two distinct propositions (two objects of thought), but the cognitive significance of these objects of though, which is related to their character, is the same. Since a priori and a posteriori are epistemic properties, and since character is the bearer of cognitive significance, it is the character that is a priori or a posteriori. But contingent and necessary are metaphysical properties of propositions, and hence are related to the content. Some sentences containing pure indexicals have special characters in the sense that they produce true propositions in any context of use. Some good examples are:

I am here now.
I exist.[1]

Therefore, one can know in advance that these sentences will express something true whenever employed in any context. The character is a priori if it yields true contents (propositions) in any context of use; it is a posteriori if it yields sometimes true and sometimes false contents. Nevertheless, the proposition produced in each context is contingent. Something similar happens to some sentences containing demonstratives, e.g.,

Dthat [the German chancellor in 2016]
is the German chancellor in 2016

where 'Dthat' is an operator introduced by Kaplan that works as a paradigm of a demonstrative, only taking as argument a definite description instead of a real demonstration. The complex 'Dthat [the German chancellor in 2016]' is a rigid designator of the object selected by the description 'the German chancellor in 2016' (which is non-rigid). We know a priori (i.e., it will be true in any context of use) that both sides of the identity must refer to the same object. But 'Dthat [the German chancellor in 2016]' is rigid and, hence refers to the same object in all possible worlds, while 'the German chancellor in 2016' is non-rigid, and might refer to distinct objects in distinct possible worlds. Hence, although the sentence yields a true proposition in any context in which it is employed, that proposition will be false at some other possible world and, hence, is contingent.

[1] Actually, this depends on some assumptions on the kind of contexts that are admissible. Kaplan only allows what he calls "proper" contexts of utterance, i.e., contexts in which the agent of utterance is at the location of the utterance, at the time of the utterance and in the possible world of the utterance. Contexts that do not have this feature are called "improper", and need not be considered (in the same way that impossible worlds are irrelevant for modal semantics). The restriction to proper contexts is not uncontroversial. (See, e.g., Predelli 2005.)

3. *Donnellan's Criticism*

Perhaps the clearest and most influential criticism of Kripke's cases of CAT came from Donnellan (1977). This is a little ironic since Donnellan himself, along with Kripke, championed the direct reference theory of proper names. Donnellan's line of thought is that (M) and (N) could not represent genuine contingencies, i.e., be true in virtue of the way the world is. Genuine contingencies could not be known without experience and, hence, could not be known a priori. Donnellan first raises a deeper (although peripheral) concern as to whether the introduction of names by means of descriptions in fact yield, as Kripke claims, rigid designators. If the name 'Neptune' is introduced by means of the description 'the cause of the perturbation in the orbit of Uranus', this might be the result of two different processes: one of them is that the description is taken as a synonym of the name. The other is that the description merely fixes the referent of the name as being its denotation (if there is one). Now Kripke claims that:

i. Names can be introduced in the second way, i.e., with definite descriptions playing merely a reference-fixing role;

ii. It is part of our linguistic practices that most ordinary names are introduced in this way.

Donnellan claims that there is no way of deciding whether a name like 'Neptune' was introduced as a rigid designator (as opposed to the synonym of the description 'the cause of the perturbation in the orbits of Uranus'). His claim is based on a standard objection[2] against Kripke's so-called modal argument, which explores the fact that a sentence like

> (A) It could be the case that Neptune is not the cause of the perturbations in the orbits of Uranus.

is intuitively true. But if 'Neptune' were synonymous with 'the cause of the perturbations in the orbits of Uranus', the sentence that follows 'that' would be a contradiction, and hence (A) would be false. The intuitive reading of (A) as true counts as decisive evidence, according to Kripke, that 'Neptune' and 'the cause of the perturbations in the orbits of Uranus' have different modal behaviors and, hence, cannot be taken as synonymous. The standard objection is that this conclusion does not follow from the fact that (A) is intuitively true. For one can reconcile the intuitive truth of (A) with the descriptive nature of the name by supposing that proper names are descriptions with a special property, namely, that of always taking primary scope. This reading of (A) would be made explicit by

> (A*) The cause of the perturbations in the orbits of Uranus is such that it could be the case that it is not the cause of the perturbations in the orbits of Uranus.

[2] Originally made by Dummett (1973: 113–6).

Hence, Kripke's conclusion that names and descriptions cannot be equivalent would be a *non sequitur*. According to Donnellan, any appeal to differences in modal behavior of names and definite descriptions as evidence for the rigidity of the former would be open this standard objection (Donnellan call it an "evasion" (1977: 15)), i.e., that a counterfactual sentence like (A) is not decisive, since considerations of scope can always be raised as an alternative account of the intuitive reading.[3]

If the intuitive reading of (A) does not count as evidence for the rigidity of 'Neptune', does anything count? Donnellan thinks that nothing whatsoever at the grammatical level could indicates unequivocally that the name is rigid: only the intention of the speaker that introduces the name (in our case, Leverrier) could determine whether it is meant to be rigid or descriptive. He concludes that there is no reason to suppose that in ordinary language there really are rigid names introduced as part of normal linguistic practices. Only an explicit convention that a name should be taken as rigid would yield this effect. Hence, in the absence of an explicit convention, we could not tell whether rigid names are part of language or merely a theoretical possibility. But the mere theoretical possibility of rigid names is enough to pose a problem since it implies the theoretical possibility of sentences that would express CATs (like (N)). These considerations are independent from the thesis that names are rigid. What is under discussion is whether the rigidity of names is incompatible with the thesis that they can be taken as equivalent to definite descriptions. It is not incompatible, although names may be (and in fact are, for Donnellan) directly referential for other reasons.

Given that the introduction of rigid names by means of merely reference-fixing descriptions is at least a theoretical possibility, does this possibility imply the possibility of CATs in any interesting sense? The main point of Donnellan's paper is that it doesn't. And his argument for it starts with a distinction between (i) knowing that a sentence is true and (ii) knowing the truth that this sentence expresses. It is essential to keep in mind that the putative knowledge that one might have as effect of the stipulation, if it is not simply metalinguistic, must be *de re*. This is so because, by hypothesis, 'Neptune' is directly referential and, therefore, has no descriptive content at all. But is there such knowledge in these cases? Donnellan avoids giving a precise characterization of *de re* knowledge, but offers what he describes as two loose principles, two minimal conditions for it. Both principles are not exactly concerned with the relation between knower and the object of the knowledge, but

[3] If Dummett is right, a proper name like 'Neptune' is always equivalent to a definite description, and a distinctive feature of names *qua* descriptions is that they always take primary scope (differently from ordinary descriptions, for which there might be ambiguity of scope). This feature of the corresponding descriptions would simulate the rigidity of proper names. There is hardly any independent evidence for this claim and its only motivation seems to be to get around Kripke's modal argument.

are rather conditions on the knowledge report. The first principle is roughly the following: a true report 'S knows that n is F' (formulated in an idiolect in which 'n' is a name of an object or person and 'F' names a property) is of *de re* knowledge only if, in any other idiolect that contains a name 'm' for the same object or person and a translation 'G' for 'F', the report 'S knows that m is G' is also true. (In other words, the report is of *de re* knowledge with respect to some object or person if its truth does not depend on the particular name for that object or person.) The second principle is the analogue for indexicals: a true report 'S knows that n is F' (formulated in an idiolect in which 'n' is a name of an object or person and 'F' names a property) is of *de re* knowledge only if one can substitute 'that' (demonstrating the same object) or 'you' (in the presence of the person) for 'n' and the new report is also true. Taken together, both principles say that *de re* knowledge of an object requires that a true report of it be insensitive to the particular name or indexical used to designate it.

Now Donnellan claims that Kripke's Neptune case fails on the first requirement. He imagines a scenario in which there are inhabitants of Neptune having, in their idiolect, a different name for it, e.g., 'Enutpen', and knowing that their planet is responsible for the perturbation in Uranus' orbit. From there they observe Leverrier through a powerful telescope (before Leverrier observed Neptune) and see that he makes the stipulation that 'Neptune' refers to the cause of the perturbation in Uranus' orbits. The report 'Leverrier knows that Enutpen causes the perturbation in Uranus' orbit' seems to be intuitively false in their idiolect, although the report 'Leverrier knows that Neptune causes the perturbation in Uranus' orbit' is true in our idiolect. This is a sign, or so Donnellan thinks, that the truth or falsity of the later report depends on the particular name involved, and hence it cannot be a report of *de re* knowledge.

Where does the impression that Leverrier does have some sort of knowledge in virtue of the stipulation come from? What Leverrier has, according to Donnellan, is the merely metalinguistic knowledge that a certain sentence (N) is true, but not the knowledge of the truth that it expresses. The latter would require some sort of contact (acquaintance?) between Leverrier and Neptune, and hence could not be a priori.

Several other philosophers have taken a line of criticism that, although differing in the details, retain the spirit of Donnellan's approach (i.e., that no a priori knowledge is possible in the Neptune-like cases). E.g., Plantinga (1974), Carter (1976), Schiffer (1977), Salmon (1986), Soames (2003, 2005).

4. *Evans*

Evans (1979) has an almost entirely different approach. He develops both a general and a particular strategy to deal with what he describes as the "puzzle" represented by CATs. The general strategy is supposed

to deal with all forms of CATs. And the particular is meant to deal with the specific version of it discussed by Kripke and Donnellan. Let us look a t the particular solution first. It presupposes an incursion into a theory of what Evans calls *descriptive names* (i.e., names such as 'Neptune', whose reference is fixed by a definite description; Evans' favorite example is 'Julius', which he introduces as referring to 'the inventor of the zip'). Evans thinks that names like these are rare in ordinary language, but some cases do exist (this is an empirical claim), and this is enough to generate the puzzle. But, contrary to Kripke and Donnellan, he thinks that these names do keep their descriptive content after the reference has been fixed. The main obstacle to this claim is Kripke's modal argument, which is based on the fact that 'Julius is F' and 'the inventor of the zip is F' exhibit different modal behavior, since in every possible world 'Julius is not Julius' is false, while in some possible world 'Julius did not invent the zip' is true. (This is just another way of saying that 'Julius' is rigid, while 'the inventor of the zip' is non-rigid.) The modal argument presupposes that if two sentences differ in their modal behavior, then they must correspond to two different propositions and, if this is so, they have two different contents. In other words, it presupposes that content and proposition is one and the same thing, and hence if two sentences correspond to two distinct propositions, they cannot have the same cognitive content. But this is something that Evans wants to challenge. The same content can correspond to two distinct propositions, one of them a necessary truth, and another one a contingent truth. He is inspired by examples like the following pair of sentences:

(i) John is as tall as John
(ii) John is as tall as himself

(i) attributes to John a property that some, but not all, objects have, namely, to be as tall as John. But (ii) attributes to John a different property, namely, that of being as tall as himself. In symbols, the properties attributed to John in (i) and (ii) are, respectively,

(P) λx(x is as tall as John)
(P') λx(x is as tall as x)

Thinking of intensions as functions from possible worlds to extensions, they correspond to different intensions, since P will select, in each world, the class of objects that are as tall as John, while P' will select, in each world, the class that includes the totality of objects in that world. Anyway, Evans' strategy presupposes dissociating differences in the modal profile of expressions such as 'Julius is F' and 'the inventor of the zip is F' from differences in the cognitive content of both: they might behave differently in modal terms and, nevertheless, have the same cognitive content.

5. *Descriptive names*

As Evans' conceives them, descriptive names are a kind of monster, similar to those mythological creatures that result from a forbidden relationship between humans and gods. For they result from the equation of expressions belonging to two entirely different semantic categories: we have a proper name like 'Julius' (with some essential properties of names, like being referential and being rigid) having its content (or Fregean sense) given by a definite description (which is neither referential nor rigid). Definite descriptions, for Evans, are not referential expressions: they are binary quantifiers. 'The inventor of the zip is British' is, according to his suggestion, similar to 'All inventors of the zip are British', 'Some inventors of the zip are British', 'Most inventors of the zip are British', etc. These can be represented by second-order operators that associate pairs of conceptual expressions to truth-values, i.e., as 'All(I, B)', 'Some(I, B)' (where I and B abbreviate 'inventor of the zip' and 'British', respectively). By parity, 'The inventor of the zip is British' has the form 'The(I, B)'. It follows that 'Julius' is actually equivalent to 'The (I, X)', which stands for a second-order operator from pairs of concepts into truth-values. So, the Fregean sense of a descriptive name in Evans' conception is not given by another referential expression (as we would expect), but is the sense of a second order operator applied to a conceptual expression. The sense of the name is, roughly speaking, a way of presenting its reference as being the unique object that satisfies the description.

The above characterization of 'Julius' as a descriptive name involves two claims: first that Julius is a referential expression and, second, that 'the inventor of the zip' (that gives the former its Fregean sense) is not a referential expression. To ground that, two questions have to be answered:

(i) How do we know that 'Julius' is a referential expression? Evans offers two reasons. The first, and most important, is that we can easily make 'Julius' fit into a minimal and general theory of reference. (The same does not hold for 'the inventor of the zip', as we shall see.) A minimal theory of reference (in the sense that it contains all that is necessary and sufficient to characterize what is essential in reference) for names (such as 'Max Freund') can be given by the homophonic clause:

'Max Freund' refers to Max Freund.

Why is this clause enough? Because it gives everything that one should expect from any adequate theory of reference, i.e., it can be combined with the notion of truth and satisfaction in order to make the following principle true:

If 'a' refers to o, and 'Fa' is atomic, then 'Fa' is true iff o satisfies 'F'

In our case,

> If 'Max Freund' refers to Max Freund, and 'P(Max Freund)' is atomic, then 'P(Max Freund)' is true iff Max Freund satisfies 'P'.

In other words, no matter what the foundational reasons are for saying that Max Freund is the reference of 'Max Freund' (e.g., that it corresponds to the intention of someone using the name, or that it is causally tied to the name, or whichever theory one might have that connects Max Freund to 'Max Freund'), and for saying that Max Freund satisfies 'P', all that matters for the concept of reference, according to Evans, is that is combines properly with the concept of truth and with the concept of satisfaction. The second reason is simply an appeal to the intuition that 'Julius' is used rigidly, i.e., that we normally consider sentences like the following as false:

> If you had invented the zip, you would have been Julius.
> If Julius had not invented the zip, he would not have been Julius.

(ii) How do we know that descriptions are not referential expressions? Here Evans offers a meta-philosophical argument (which is different from Russell's reasons for not considering descriptions as referential; for Russell an expression is not referential if a sentence that contains it has truth conditions even under the assumption that the expression lacks a reference, which means that a sentence containing a referential expression lacks truth conditions if the expression has no reference). Assimilating descriptions into the category of referential expressions would require making some adjustments in the theory of reference in order to leave room for some phenomena that are typical of descriptions, such as their non-rigidity and ambiguity of scope when embedded in sentences containing negation (such as 'The current king of France is not bald') or modal operators (such as 'The firs man in space could have been an American'), or epistemic operators (such as 'George IV wants to know whether Scott is the author of Waverley'). But this would introduce a great discontinuity and artificiality in the theory of reference, since paradigmatic referential expressions (like pronouns and ordinary names) never really take up the possibilities left open by this new (modified) theory.

6. *Free Logics and Descriptive Names*

Consider Kripke's sentence (N) again:

> (N) If there is one and only one cause of the perturbation in Uranus' orbit, then Neptune causes the perturbation in Uranus' orbit.

In symbols

> (N) $\exists!x \ \ Fx \rightarrow (F(n))$

'n' is supposed to be a name that is rigid and directly referential, and that's why the truth expressed by the sentence is contingent: there could be a possible world in which the antecedent is true (i.e., there is

one unique cause of the perturbation of Uranus' orbit) and the consequent false (i.e., Neptune is not the cause of the perturbation in that world). If 'n' were a description, (N) would be a necessary truth because the consequent would be a tautology. However, as Evans notices, (N) can only correspond to a proposition if 'n' refers. In the absence of a reference, there is no proposition to be known. And this is so despite the fact that, because the name is descriptive, in the absence of a reference a sentence containing it might still have a Fregean sense (and, hence, can be understood). Some sentences W(n) containing the name 'n' can be true even if n does not exist (e.g., 'It is not the case that Julius is F', when there is no inventor of the zip). But, according to classical logic, we must be able to apply existential generalization in the 'n'-position, thereby getting $\exists x\ W(x)$, which might be false. Hence, if one wants to claim that (N) yields knowledge, one must not allow as a general rule the inference from W(n) to $\exists x\ W(x)$ unless there is a guarantee that the name 'n' refers (i.e., that $\exists x(x=n)$), which means that one must adopt a logic that is free of existential assumptions, i.e., one must adopt some form of free logic.

Evans actually suggests something stronger, in the form of two theses:

(i) The acceptance of descriptive names requires the acceptance of free logics (p. 166)

(ii) The acceptance of free logics requires the existence of descriptive names (p. 173)

Both taken together imply that the assumption of free logic and the assumption of descriptive names go together, or so Evans seems to think. In other words, either (N) is formulated using a free logic (i.e., without the assumption that 'n' refers), or there is no clear candidate for the corresponding CAT. And being formulated using free logic presuppose the existence of descriptive names, i.e., names that are rigid but have a descriptive meaning even in the absence of a reference; this makes it possible that a sentence containing the name might be understood (i.e., have its truth-conditions formulated) even if there is no referent. But, as Evans points out, Donnellan is not willing to accept the existence of descriptive names in this sense. And this is so for two reasons. First, for Donnellan, even if a name has its referent fixed by a description by an explicit convention, the description does not remain attached to the name. (In other words, if there is no referent of the name, there is nothing to be understood in a sentence containing the name.) Second, and more importantly, the understanding of a name requires, for Donnellan, some sort of contact with its referent, and therefore there cannot be such understanding if the name does not refer. Hence, Evans' first point against Donnellan is, as he himself calls it, *ad hominen*: if (N) is to be a candidate for contingent a priori knowledge, then it must not allow for existential quantification in the position occupied by the name. But if this is so, one must accept free logic. But free logic (at least in

Evans' version of it) requires descriptive names, i.e., names such that sentences containing them do not depend on the existence of referents to have meaning. Donnellan does not accept such names, and so, according to Evans, he should not recognize a puzzle in the first place.

7. *Contingency: Deep and Superficial*

However deep and ingenious the reflections on reference and descriptive names are (and whether or not they yield a solution to the puzzle generated by CATs), these reflections turn out to be not strictly necessary for dealing with the puzzle, since there is a second, simpler and more general strategy, this time not coming out from a theory of reference, but from a theory of contingency. Evans introduces a distinction between superficially and deeply contingent truths. A sentence P is superficially contingent iff it is false at some possible world, i.e., '$\Diamond\neg$P' is true. And it is deeply contingent if a verifying fact is not guaranteed by semantics alone, i.e., by understanding the sentence one is not thereby assured that there is a verifying fact that is a "contingent feature of reality" (p. 185). If we can talk of contingency as a sort of requirement for truths, superficial contingency and deep contingency require different things. Superficial contingency requires a certain minimal modal profile, while deep contingency requires something from the truthmaker, i.e., that it is not generated by semantics alone.

Evans claims that the cases presented by Kripke are, at best, only superficially contingent, since the semantic stipulation gives a guarantee of a verifying fact for (N). (This is also presumably so for Kaplan's cases.) And all cases of superficially CATs have the same source. Since propositions are true or false at possible words, we might also think of them as properties that hold (or do not hold) of possible worlds, or as requirements for being true in each possible world. Now a property or requirement of this kind might be based on a contingent feature of the actual world. In this case, on the one hand it trivially applies to the actual world (because it was extracted from it) and, on the other hand, it is not a property that all possible worlds have (because it is based on a contingency) and hence does not apply trivially to any world. Let C be any contingent fact of the actual world (e.g., that the sky is blue), and consider the property (P) of possible worlds expressed by

(P) λw(w includes C iff @ includes C)

(where @ is the actual world). '@ includes C' expresses a necessary truth, since in any possible world it will be true that @ includes C. The result is that this is a property that not all possible worlds have, but @ certainly has. Now consider a different property

(P*) λw(w includes C iff w includes C)

P and P* correspond to two properties, since the first is false of some possible worlds, but the second is true of all of them. However, Evans claims that they yield cognitively equivalent propositions when applied

to @, since both result in saying that @ includes C iff @ includes C (which is not really a new thing to know). Here, again, what is operative is the distinction between proposition and cognitive content.

This is supposed to be the general strategy for dealing with all cases of CATs, and reducing them to things that do not represent substantial, but only trivial, knowledge about the actual world. Better said: it reduces them to an attribution of a non-trivial property to the actual world that is epistemically equivalent to the attribution of a trivial property. How does this general strategy apply to Neptune-like cases (i.e., cases involving descriptive names)? For we saw that the first strategy, trading on specific features of descriptive names, was actually not strictly necessary, since there was something broader and more fundamental to solve the problem of CATs in its greater generality. Now remember Leverrier's sentence:

> (N) If there is one and only one cause of the perturbation in Uranus' orbit, then Neptune causes the perturbation in Uranus' orbit.

In formal notation:

$$(N)\ \exists!x\ Fx \rightarrow ([n]F(n))$$

Here '[n]' is the scope indicator of the name 'n', and it has the effect of not allowing existential quantification in the 'n" position over (N), but only over 'F(n)', which is, according to Evans, a necessary condition for there being the puzzle of CATs in the first place.[4] Let's rewrite (N) by unpacking 'n' and using '@' for the actual world. A definite description must have a position open for the possible world in which it is being considered, and 'n' has, by convention, its reference fixed by the description taken in the actual world. So (N) becomes:

$$\exists!x\ Fx \rightarrow (F\ (\text{the } x\ F(x, @)))$$

Now if we consider that 'Fx' is true or false under an assignment of x only in relation to a possible world w, we might rewrite it as a binary relation 'F(x,w)' between objects and possible worlds, and hence we have

$$\exists!x\ F(x,w) \rightarrow (F\ (\text{the } x\ F(x, @), w))$$

We might see this as a property of possible worlds:

$$(\lambda w)(\ \exists!x\ F(x,w) \rightarrow (F\ (\text{the } x\ F(x, @), w)))$$

This property requires of a possible world w that, if there is one and only one object that is F in w, then the one and only object that is F in @ is F in w. This property is not satisfied by all possible worlds: it is

[4] If we had

$$[n](\exists!x\ Fx \rightarrow (F(n)))$$

instead, we should allow for existential generalization, i.e.,

$$\exists y(\exists!x\ Fx \rightarrow (F(y)),$$

but this cannot possibly be known a priori, since it involves the existence of an object with a contingent property.

false in those worlds w in which there is one and only one F, but the only object that is F in @ is not F in w. If this property is applied to the actual world it yields

$$\exists!x\ F(x,@) \rightarrow (F\ (\text{the } x\ F(x, @), @)))$$

which says of @ that if there is one and only F in @ then the one and only F in @ is a F in @. This has the same content as the application of the property

$$(\lambda w)(\exists!x\ F(x,w) \rightarrow (F\ (\text{the } x\ F(x, w), w))))$$

to @, which is a property that absolutely every possible world has (i.e., if there is one and only one F in w, then the F in w is a F in w).

We can see the intended effect leaving untouched the definite description. But if we want to follow Evans and treat it as a binary quantifier, we can rewrite

$$W\ ((\text{the } x\ F(x, w), w)$$

as

$$(Ix)(\ F(x, w);\ W(x, w))$$

(where '(Ix)' is the binary operator corresponding to 'The'), and the property above becomes

$$(\lambda w)(\exists!x\ F(x,w) \rightarrow (Ix)(\ F(x, @);\ F(x, w))$$

which, applied to the actual world, yields

$$\exists!x\ F(x,@) \rightarrow (Ix)(\ F(x, @);\ F(x, @))$$

This is, according to Evans, epistemically equivalent to the attribution to @ of the trivial property

$$(\lambda w)(\exists!x\ F(x,w) \rightarrow (Ix)(\ F(x, w);\ F(x, w))$$

8. *Contingency and Existence*

An important aspect of Evans' general strategy is that it does not have to appeal to a claim like the following: if a statement is contingent, it must be existentially committing and, hence, cannot be a priori at all (since matters of existence of ordinary objects are not knowable a priori). For the source of contingency is not necessarily related to the existence of objects with such and such features. There is a trivial and a non-trivial reason for this. The trivial is that there might be CATs that do not involve singular terms and, *a fortiori*, do not require the existence of anything to be true. E.g., the following sentence

All whales are mammals iff all whales are mammals in @.

There is no singular term here, only a universal quantifier and two conceptual expressions. Hence, its truth does not require the existence of any object. The biconditional is knowable a priori, and is contingent (because the left side of it is true in the actual world but false in some possible worlds, while right side of it is true in every possible world). The non-trivial reason is compressed in an apparently enigmatic passage:

> [I]t does not follow from the fact that a sentence is contingent because it is formulated with the use of a referring expression, that its contingency is due to the contingent possession of a certain property by the object to which the expression refers. (p. 175)

Evans is thinking about a sentence like our

(N) $\exists!x\ Fx \rightarrow ([n]F(n))$.

Now, the fact that it includes a referring expression (which is, therefore, rigid) implies that

(N*) $\Diamond(\exists!x\ \Phi x\ \&\ \neg[n]\Phi(n))$

(If 'n' were not a referring expression, (N) would have been a necessary truth since the consequent of the conditional would be a necessary truth, and (N*) would be false.) Which means that it is contingent (at least superficially contingent, in Evans' sense). But because 'n' is taken with narrow scope we are not allowed to infer

(N**) $\exists y\ (\exists!x\ Fx \rightarrow (F(y)))$

which is contingent in virtue of the fact that its truth demands the existence of an object with a specific property (i.e., the property corresponding to

$\lambda y\ (\exists!x\ Fx \rightarrow F(y)))$.

(N**) could only be derived if the name in (N) had wide scope. So, the only existential claim allowed by free logic to follow from (N*) is

(N***) $\exists!x\ Fx \rightarrow (\exists x\ F(x))$

which is knowable a priori.

In a nutshell: the puzzle requires that (N) is such that:

(i) It is formulated with a name (referential)
(ii) The name is descriptive
(iii) The name takes narrow scope

For: (i) – if it is not a name but a description that appears in (N), then (N) expresses a necessary (and not a contingent) truth; (ii)-if the name is not descriptive, the existence or not of a proposition to be known would be dependent on a contingency (but, as it seems, a proposition cannot depend, for its existence, on a contingency); (iii) – if the name takes wide scope, it should be open to existential quantification and, hence, the proposition could not be a priori, but only a posteriori. Although (i)–(iii) are necessary and sufficient conditions for (N) being contingent, the latter is not contingent in virtue of the existence of the object corresponding to 'n'. The contingency is (only) a superficial one, due to the modal profile of (N), but not due to the existence of an object with a particular property in the world since (N) does not require such existence to be true.

9. *Some Drawbacks*

Evan's general strategy has some drawbacks. A minor one is that all cases envisaged by him involve in one way or another an actuality operator (actually P is true iff P is true in the world that is treated as actual) or reference to the actual world, which enables one to create a property that is trivially satisfied in the actual world (because it uses an empirical fact that obtains in the actual world) and not satisfied in some other world. In the same way that 'the actual DD' is a rigidified description, 'actually P' is necessarily true (if P is true in the actual world) or necessarily false (if P is false in the actual world). But reference must be made to the actual world in the shaping of such properties; otherwise we do not get the same effect. The strategy was meant by Evans to be completely general and not involving aspects belonging to a theory of reference. But the actuality operator is an indexical (indeed, a pure indexical, according to Kaplan's terminology), with all the properties that indexicals have, including rigidity and direct reference. So, the strategy is not, after all, independent from aspects of the theory of reference, and not quite as general as Evans might have thought.

A more serious drawback is that Evans' strategy relies heavily on the distinction between cognitive content and propositions. But this distinction is at best highly controversial. For it implies, in talking about the cognitive content, eschewing aspects related to modal behavior. This is not clearly reasonable, unless one sees this as anticipating (without fully articulating) Kaplan's later distinction between character and content. As we saw, in Kaplan's (1989) framework, we can explain away the strangeness of CATs by attributing distinct roles to each of these dimensions. The character is a priori (in case all propositional contents generated by occurrences of that sentence in any context are true propositions) or a posteriori (in case some propositional contents generated by that sentence in some contexts are true and some are false). And the content is the proposition generated by the character in that context, and it is the content that is contingent or necessary, according to its modal profile. If this is so, Evans can be understood as holding that all cases of CATs are due to semantic properties of indexicals[5], and also that all CATs derived from indexicality are only superficially contingent. Hence, that there are only superficially CATs.[6]

But what if there are cases of indexical-free CATs? It seems that they would have to be treated as deeply contingent in Evans' sense. Williamson (1986) discusses an example that would have this effect. At first sight this seems to be an easy task, since sentences like (N) contain proper names but appear to contain no indexicals (provided the

[5] Hawthorne (2002: 247) and Soames (2003: 417) seem to agree with Evans on this.

[6] Contrary to Evans, Kaplan does not treat characters that always yield true contingent propositions in virtue of indexicals as less interesting (or superficial) than those that do not always yield true propositions. Kaplan draws no distinction between superficial and deep contingency.

verb is given a tenseless reading). But this easy alternative would not be indexical free because 'Julius' has its reference fixed with the help of a definite description taken in the actual world, and hence somehow the actuality operator (which is an indexical) is behind the name. (Since he places no restriction, Williamson seems to think that this is so for all proper names whose reference is fixed by descriptions, even those that are not descriptive in Evans' sense; so presumably he is leaving out of the picture names whose reference is fixed by indexical-free description such as 'the first prime number greater than the cardinality of the continuum'.)

Here is Williamson's example:

> There is at least one believer

There is no proper name and no indexical here. The 'is' here has to be taken in its tenseless reading, for otherwise a temporal indexical would be present.[7] Presumably one can know a priori (B), for believing it (without further empirical evidence) makes it true. And it is contingent, given that there could have been no believers. One might think that (B) can only be known by inference from

> (I) I am a believer.

But (I) has the indexical 'I', and this would mean that knowledge of (B), like knowledge of 'Julius invented the zip', is also derivative from knowledge concerning indexicality. However, there is another way of deriving (B) that does not appeal to (I), namely, by means of the following principle:

> (M) If there is a valid deduction of P from the premise that someone believes that P, then one is assured in believing P.

(M) is hyperreliable in that any belief acquired by employing it is guaranteed to be true. So, if one comes to believe a proposition P by just employing (M), that belief is certainly true. On the other hand, if there is a valid deduction of P from the premise that one believes P, then no empirical evidence at all is needed to know P, just believing it is enough to transform that belief into knowledge by using M. There certainly is a valid deduction of 'there is at least one believer' from the premise that 'someone believes that there is at least one believer'. Representing 'There is at least one believer' formally as $y $ x(B(x, y)) (where 'B(x,y)' abbreviates the binary relation x believes y), we have the premise

> ∃ x(B(x, ∃v∃tB(t, v)))

from which, by existential generalization, we get

> ∃y∃x(B(x, y))[8]

[7] One might raise doubts as to whether the effect intended by Williamson (i.e. something being made true by someone just believing it) can be obtained by taking the verb as tenseless. I will not discuss this issue here.

[8] We can skip here some minor complications such as the fact that the first order variables have both believers and propositions in the domain, and that the existential

and, hence, by (M), one is entitled to believe a priori that there is at least one believer. In order to deal with the objection that this might be a necessary truth (in case one takes God's existence to be necessary), Williamson proposes the following modified version of (B):

(B') There is at least one fallible believer.

Consider the premise that someone believes that there is one fallible believer. There are two possibilities for this someone who believes this: he or she might be either fallible or infallible. If he or she is infallible, then the content believed must be true, and it follows that there is at least one fallible believer. If he or she is fallible, then it follows by existential instantiation that there is at least one fallible believer.

10. *Going Deeper on Superficial Contingencies*

Williamson's example seems to show that Evans' thesis that all CATs are superficially contingent is questionable. But we can go further and question whether the distinction between weakly and strongly CATs makes sense in the first place. (And, even if it does, we can ask whether it can be motivated independently from the fact that it provides a way of discrediting CATs or whether it is simply *ad hoc*.) Evans and some of his followers (e.g., Hawthorne (2002)) seem to think that the distinction is clear enough, and it suffices to show that there can be no coherent account of cases of genuine knowledge of deeply CATs.

It is not completely clear why knowledge of weak contingency should be less interesting than that of deep contingency. (Especially if we adopt a two-dimensional semantics, like Kaplan does, and attribute cognitive significance to one of the dimensions.) From the examples given by Evans it seems that all cases of superficial contingency follow a certain pattern. We can know some superficially contingent truth about something if we consider a property by reference to this same thing. Hence, 'a has the same size as a', 'a has the same color as a', 'a is as old as a' all state contingent properties of a (in the sense that not every object possess the properties that a has). This is no different when the actuality operator is used, since it applied in a possible world means 'as it is in this world'. Hence, somehow reflexivity seems to be the source of superficial contingencies in Evans' sense, while no reflexivity is involved in deeply contingent truths that, presumably, cannot be a priori for him. (We have seen from Williamson's example that this is probably not true.) But there are different ways of stating a property of an object by reference to itself. Some of them yield trivial properties of the object, while others yield properties that are not trivial. Evans seems to base his general approach on properties of a of the form

a has the same P as a.

generalization would, strictly speaking, require the previous transformation of the existential sentence in the corresponding proposition by means of a that-operator.

But now consider:

> Hesperus is as large as Hesperus.
> Hesperus is as large as Phosphorus.

Are they both sources of the same property? Or two distinct properties, one of them holding trivially of Hesperus? Having Kripke's example in mind, consider

> Paderewski has the same DNA sequence as Paderewski.

Does it state a trivial property of Paderewski or something that one might be surprised in discovering (e.g., if one thinks that they are different persons, as in Kripke's example)? This brings us into the vicinity of Frege's problem, i.e., the problem of explaining informativeness when reference to the same object is made repeatedly. One is not guaranteed that by defining a property of an object by reference to this same object we thereby get something that is trivially known of this object, even if reference is done using the same name-type. This might depend on there being coordination in language or thought between different tokens of names.[9]

References

Carter, W. R. 1976. "On A Priori Contingent Truths." *Analysis* 36 (2): 105–6.

Donnellan, K. 1977. "The Contingent A Priori and Rigid Designators." *Midwest Studies in Philosophy* II: 12–27.

Dummett, M. 1973. *Frege. Philosophy of Language*. Cambridge: Harvard University Press.

Evans, G. 1979. "Reference and Contingency." *The Monist* 62 (2): 161–89.

Fine, K. 2007. *Semantic Relationism*. Oxford: Blackwell Publishing.

Hawthorne, J. 2002. "Deeply Contingent A Priori Truths." *Philosophy and Phenomenological Research* 65 (2): 247–69.

Hawthorne, J. and Manley, D. 2010. *The reference Book*. Oxford: Oxford University Press.

Kaplan, D. 1989. "Demonstratives." In Almog, J., Perry, J., and Wettstein, H. (eds.), *Themes from Kaplan*. Oxford: Oxford University Press: 481–564.

Kripke, S. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.

Plantinga, A. 1974. *The Nature of Necessity*. Oxford: Clarendon Press.

Predelli, S. 2005. *Contexts. Meaning, Truth, and the Use of Language*. New York: Oxford University Press.

Salmon, N. 1986. *Frege's Puzzle*. Cambridge: MIT Press.

Schiffer, S. 1977. "Naming and Knowing." *Midwest Studies in Philosophy* II: 28–41.

Soames, S. 2003. *Philosophical Analysis in the Twentieth Century. Vol. 2: The Age of Meaning*. Princeton: Princeton University Press.

[9] This thesis is developed in Fine (2007).

Soames, S. 2005. *Reference and Description*. Princeton: Princeton University Press.

Williamson, T. 1986. "The Contingent A Priori: Has It Anything to Do With Indexicals?" *Analysis* 46 (3): 113–7.

Croatian Journal of Philosophy
Vol. XVI, No. 47, 2016

Book Reviews

Mark G. E. Kelly, *The Political Philosophy of Michel Foucault*, New York and London: Routledge, 2009, 186 pp.

Jon Simons wrote that "commentary on and critique of Foucault's notion of power has become an intellectual industry in itself" (129), in his book entitled *Foucault and the Political*. But only a few of those books which engage with Foucault have an original perspective and coherent trajectory. The main reason underlying this may be Foucault's being "on a perpetual slalom course between traditional philosophy and the abandonment of any pretension to seriousness" (93) as Maurice Blanchot said. Foucault permanently changes his position from theme to theme and never sticks to any political or philosophical constant. His fluid thinking style always challenges and compels his commentators.

Mark G. E. Kelly's book is one of those few works that distinguish themselves: he elaborates Foucault's oeuvre as a *coherent whole*. Actually Kelly tries to show that the philosopher's thought is consistent in itself. In this respect, *The Political Philosophy of Michel Foucault* which locates itself against Eric Paras' *Foucault 2.0* is very attractive for both beginners and scholars. The book cuts across Foucault's political positions, arguments and reasoning as a whole in seven chapters; thus, we encounter a new image of Foucault by the end of the book.

The first chapter of the book, which is entitled "Epistemology," begins with the first studies of the philosopher such as *The Archaeology of Knowledge* and *The Order of Things: An Archaeology of the Human Sciences*. Kelly *thinks that these books* are foundational. The main motive of this chapter is locating the philosopher's thoughts within a "materialism". But this materialism has a very different structure and features than other ways of thinking called materialism. Because of that, Kelly calls it "a materialism of the incorporeal". This term seems paradoxical at first glance, but according to Kelly, "the incorporeal, here typified the event, is not necessarily immaterial since it can be something that occurs in the material world" (13). "A materalism of the incorporeal", as a specific foundation of Foucault's oeuvre, is also a kind of ontology which never looks like other ontological approaches in general. Kelly continues his task by providing a reading of theorists and philosophers such as Althusser, Nietzsche and Derrida. Thus, he discusses Foucault's epistemological and ontological attitudes with these figures' intellectual inventory.

Another important aspect of this chapter is depicting the influence of "May 1968" on Foucault. When the "events" of '68 happened, he was living in Tunis. However, according to Kelly, the events changed him in two ways. Firstly, "Foucault threw himself into political activism for the first time in his life" (17). The second change was actualised at a theoretical level and he became seriously interested in Marxism. During this time, he never described himself as a Marxist, but the Marxist influences never shade away especially in his later works. He once mentioned that "Marx is our Machiavelli: the discourse does not stem from him, but it is through him that it is conducted" (243) in *Security, Territory, Population*.

The subsequent chapters (second and third) deal with power. The second chapter investigates Foucault's early power concepts which ended in late 1970's. The third chapter focuses on the transformation of his thoughts and concepts about power. The second chapter not only brings the main fulcrums of "the analytics of power" to light, but also tackles different aspects of power such as power-knowledge linkage, "power as war", and "technologies of power" etc. Thus, Kelly explicates what "a need to cut off the king's head" means in political theory. Moreover, he marks Foucault's principal lines of conceptualising power and he compares the essential characteristics of power in *The Will to Knowledge* and *Discipline and Punish*. In the third chapter, Kelly explores how Foucault's late works interact with his early thoughts on power. He questions the concept of "game" in order to do that. Kelly's book distinguishes itself at this point because he produces a new interpretation against commentators who argue that there are two different power models in Foucault's thought. According to Kelly, one cannot posit a change in the model, because these terms should be understood at a *metaphorical level*. The essential difference between "war" and "game" models is their level: "war occurs at a grand, societal level, whereas the game occurs at an interpersonal one" (59). From "war" to "game", he tries to assess that what power means and how it works in Foucault's terminology.

In the context of this debate, Kelly examines the concept of "government" in more detail. Especially, he focuses on Foucault's well-known article entitled "The Subject and the Power". According to Kelly, this article is a touchstone of Foucault's re-conceptualisation of power and subject. The main theme of this article is relationality which brings out the real character of Foucault's conception of power. There, Foucault defines power's relationality for the first time. After that, according to Kelly, power is defined as the capacity for making someone to do something by Foucault. This approach gives an opportunity for linking the concepts of "power" and "conduct".

Kelly examines the concept of "resistance" by focusing on "the reversibility of power relations", which he discusses in more detail in chapter five. Regarding power relations, one should use this phrase cautiously, because, according to Foucault, all social relations are not power relations, but all power relations are co-extensive. Kelly discusses power as a capacity in terms of reversibility. In this regard, his analysis perceives power is a two-way street: not only domination, but also resistance.

In chapter four, which is entitled "Subjectivity", Kelly argues that there is no rupture between Foucault's earlier works and later works on subjectivity. Especially, he strikes out at Judith Butler regarding her misreading.

Drawing on Gilles Deleuze and Friedrich Nietzsche's thoughts, he offers a new conception of Foucault's subjectivity theory which depends on materiality of the body. In this context, he elaborates the "death of man" and "return to the subject" that are attributed to Foucault. According to Kelly, the philosopher's position was never "anti-subjectivist"; but his endless quest for the "subject" should always be understood relating to the concept of "experience". Therefore, one should keep in mind that "Foucault is not interested in looking for the origin of the subject" (85).

The most important discussion in this chapter is about Butler's misreading. Kelly claims that she confused terms such as "subjection" and "subjectivation". Tracing the trajectory of how these terms are misused, Kelly criticises Butler from a very solid theoretical position. The rest of this chapter's debate revolves around "psychoanalysis" and the concepts of "interpellation" and "self".

Resistance is an inseparable part of Foucault's power theory. Therefore, Kelly discusses to the co-extensiveness of power and resistance in chapter five. He sets out to develop the concept of "resistance", which proceeds from the contention that "where there is power, there is resistance". According to Kelly, this well-known formulation points out the "potentiality" of resistance. After that point, he focuses on the relation between micro- and macro-resistances. On the one hand, he tries to make a proper analysis about these resistance types; on the other hand, he underlines that "Foucault does not himself make his distinction between micro- and macro-resistance explicitly, as in the case of power" (110). In this respect, Kelly argues that the "will" is a key notion in terms of resistance because –if I have any right to deflate the famous motto– "*We* do not even *know* what a *will* can do." Kelly's emphasis on "will" is not as decisive as he thinks, for Foucault avoided this term and never elaborated on its meaning.

The last two chapters engage with "Critique" and "Ethics" respectively. These chapters interrelate to each other in regard to practical implications of Foucault's political ontology. According to Foucault, critique is the main political duty of a philosopher. Likewise for Kelly, "The central political role of the intellectual is to advise as to the *possibilities* of political action, though an analysis of the strategies of power." In short, the critique is a kind of key for understanding historical relations between politics, subject and truth. One should mention the debate on "Bodies and Thoughts" in this chapter. There, Kelly follows Hinrich Fink-Eitel's critique about Foucault's expression in *The Will to Knowledge*, "Against the device of sexuality, the fulcrum for the counter-attack ought not to be sex-desire, but bodies and pleasures". Kelly's position is slightly different than Fink-Eitel's. He argues that "Foucault is doing something very simple, referring to bodies and pleasures themselves, the actual bodies we are/inhabit, and pleasures we experience/constitute" (147). This assertion actually stresses Foucault's materialism.

The concluding chapter of the book deals with the meanings of ethical practices in Foucault's thought. The chapter begins with a debate about "critical ethos" and then "ethics of the self". Kelly argues that Foucault uses "ethics" at least in two different senses in his later works: practices of self-relation and permanent resistance. According to Foucault, "freedom is the

ontological condition of ethics". From this point of view, ethics is determined by politics for Foucault.

What is striking about Kelly's book is how he forges Foucault's works on power into a coherent theory by using philosopher's own conceptions and reasoning. In particular, his emphasis on materialism is considerably important for discussions on Foucault's thought. As a final remark, one may conclude that the hidden matrix of the book is the conceptualisation of the connections between Marxism and Foucault. Yet, Kelly never clarifies which Marxism he particularly refers to. Nevertheless, the most visible contribution of the book to Foucault studies is its original approach developed in the second and third chapters. Kelly succeeds in developing a new perspective from Foucault's oeuvre by making use of a series of hints embedded in his works.

UTKU ÖZMAKAS

Noël Carroll and John Gibson (eds.), *The Routledge Companion to Philosophy of Literature,* London: Taylor & Francis, 2016, 520 pp.

Under the editorial wisdom of Noël Carroll and John Gibson, *The Routledge Companion to Philosophy of Literature* brings forward 40 newly commissioned essays dedicated to philosophical exploration of the wonderfully rich and excitingly intriguing phenomenon of literature. To my knowledge, this is one of the most encompassing books dedicated to analytic philosophy of literature, and the breadth of coverage testifies to the extent to which the discipline has grown and to the variety of problems it is concerned with. The outstanding selection of contributors (difficult as it was to make it, as the editors lament, given the amount of first rate philosophers who work on literature and literature-related issues), in itself indicates that this book is a must have/must read for everyone interested in and infatuated by literature.

Ranging from the forefathers of analytic philosophy of literature, to philosophers who have expanded the field by throwing light on not so often discussed specimens of literature such as popular fiction, poetry and screenplay, to people who helped deepen the field's interest in certain themes, such as emotions, imagination, empathy and character, and people who have strengthen the field's connection with other philosophical areas or have introduced literature to new areas of research such as neuroscience, the *Companion* brings together the most respected philosophers of art today, whose tireless work on philosophical challenges raised by our artistic practices is at the very foundation of contemporary philosophical approaches to art. Their contributions provide an excellent mapping of the 'philosophy of literature terrain' and give insightful summaries of the main positions, arguments and thesis. Consequently, this *Companion* is excellent

as a compendium of historical development of certain problems, positions and arguments, as a pointer on 'who said what' in philosophical debates on literature, and as a way mark on where to go next (most of the essays include a section 'recommendations for further reading'). Due to the *Companion*'s immensely informative aspect, it is valuable to those who are taking up philosophical challenges of literature for the first time, to get an idea of what these problems are and to get oriented on the main discussants in debates, problems and claims. Those more engaged with literature will profit from having 'all the eggs in one basket', as there is hardly any topic pertinent to contemporary philosophy of literature that is not discussed here (with the exception of 'philosophy in literature'). Given the depth and rigour of analyses, essays will be more than illuminating for philosophers of art generally, as there are many themes discussed that are relevant for other artistic practices.

The *Companion* is unique in several respects. First, it explains historical roots of philosophy of literature, showing how philosophers' reflections on literature helped define philosophy of art and establish literature's artistic status. Second, though each contributor is concerned with a distinctive issue, in many ways they engage with and challenge one another and, considered jointly, address some of the same problems, but do so from different perspectives. Thus, the *Companion* on the whole tackles literature from every philosophical angle, showing ways in which literature raises interesting questions and offers valuable insights for philosophy across the board. Third, the *Companion* gives space to some of the topics that have gone unmentioned in similar anthologies and textbooks. Among these, I take the relation between literature, neuroscience and theory of mind, especially from the perspective of evolution, to be particularly important. On the one hand, insights from these researches can prove immensely valuable for the longstanding discussions on the emotional engagement with literature, and, in particular, for the alleged cognitive benefits of literature (in terms of its potential indirect influence on and benefits for our cognitive economy, conceptual framework, inferential abilities, emotional and moral sensibility and the like). Although, as the contributors working on these points make clear, we are far from having conclusive evidence on whether or not literature can in fact aid us in becoming smarter, wiser and better at thinking, contemplating and understanding ourselves and others, the results of the research done so far, presented here, are already sufficiently rich to justify future interest in these topics and deeper probing of literature from these perspectives. Undoubtedly, whatever results we eventually come up with, will be of immense importance for arts generally, as well as for various humanistic sciences concerned with our cultural legacy. On the other hand, philosophical interest in literature was for a very long time influenced by the impact of Plato's negative and Aristotle's positive accounts of literature. Armed with researches done along the lines of neuroscience, psychology and evolution, we can re-evaluate those ancient arguments regarding the role of literature (and arts) within our societies, our education and our culture, as well as the impact of literature on our understanding of who we are as individuals and as society. Without neglecting or diminishing literature's artistic value, the essays gathered in the *Companion* show how philosophical explorations of

literature matter for many of our other intellectual disciplines, such as psychology and anthropology. It is safe to say that the *Companion* will redefine not only our philosophical approaches to literature and the way we read, take pleasure in and appreciate literary works, but also the way we think of and understand our cognitive and aesthetic engagement with reality, mediated as it undoubtedly is, through literature.

"It is our hope" say Carroll and Gibson, "that we have presented enough of the field to inspire readers to discover the rest and to chart out new territories in their own research" (xxiii). I have no doubt that every reader of this amazing *Companion* will feel that this hope has been fulfilled.

In what follows, I present a very brief summary of each of the essays.

The *Companion* is divided in six parts. It opens with an overview of *Historical Foundations* of the philosophical probings of literature. *Stephen Halliwell* delivers an array of challenges that the greatest philosophers of the antiquity left to philosophy of literature. Plato's is a well known challenge issued at poets who, in order to be readmitted into the perfect state, have to account for poetry's value to the wellbeing of individuals and for the community. Aristotle left us with the task of explaining the gap between the intelligibility of the casual structure in the plot and the lack of such intelligibility with reference to our lives. Negative attitudes towards (philosophy of) literature are left behind by Epicureanism. With Seneca, the challenge becomes that of addressing the philosophical lessons found in literature. This challenge is further bolstered by Plutarch, whose application of philosophical agenda to interpretation of literature is indicative of the potential difficulty of finding the balance between imposing philosophy on a work, rather than extracting it from it. *Paul Guyer* analyzes developments in philosophy of literature in the 18th century and claims that, though the field itself did not officially exist, all the major aestheticians based their general theories of art on literature. Two main issues emerged at the time, that of experiencing pleasure in tragedy and the one concerning comparison between literature and painting. To provide background to how these issues were thought about, Guyer first discusses what the perceived goals of literature were. He meticulously explains the role that Locke's theory of meaning and his theory of the association of ideas had in explaining the emotional impact of literature, before turning to the paradox of tragedy as problematized by Du Bos and David Hume. Du Bos' views are also relevant for the 'poetry versus painting debate', which was concerned with accounting for different ways in which poetry, paining and music are sources of beauty. *Allen Speight* takes us back into the 19th century, which saw a clash between two dominant views on literature, that is, on poetry. The first one, idealist, inspired by Hegel, influenced Schelling, A.C. Bradley, T.S. Eliot, American New Critics and it is best understood along the lines that M. H. Abrams calls objectivist approach. The second one is traced back to J.S. Mill, whose views on poetry were influential for the romantic strain (Wordsworth, Tennyson, Carlyle, Wagner, Nietzsche) nowadays associated with lyric expressivism. Speight explains how each of these two approaches developed and analyzes the impact of each on the theories of literature and various literary genres in the 20th century. Although the 19th century views were temporarily shaded by structuralist, post-structuralist and ideologically grounded

approaches to theory, and (within analytic aesthetics), by Frege's views on language, Speight claims that they are still visible in contemporary thinking about art, most notably in John Gibson's elaboration of literature's humanistic character and in the claim that literature presents other worlds of aesthetic creation. Three main theoretical approaches to literature in the 20th century, claims *Kristin Gjesdal*, have to be understood as a reflection on modernism. According to the Marxism and the Critical theory, what matters for literature is its relationship with the society and its potential political power to issue real social changes. The second approach, Phenomenology, centres on ontological questions – what is a work of art, what is literary meaning, how to understand the truth in the literary work – and ignores political dimension of literature. Finally, Existentialism sees literature as first and foremost a space of freedom. Gjesdal's essay is informative in summarizing the main theoretical clashes of thinkers such as Lukas and Adorno, Heidegger and Gadamer and Sartre and Merleau-Ponty. *Kristin Boyce* illuminatingly explains development of the analytic philosophy of literature, marked as it was by a certain paradox of philosophy of literature: as objects of philosophical reflection, literature and arts have been marginalized and neglected, yet they had special methodological importance for philosophers from the early days of analytic tradition. Boyce sees this as a consequence of the attempts of analytic philosophy to establish itself as close to science, rather than to arts and the humanities. She discusses ways in which modernist literature and philosophy stand to one another, and is primarily focused on analyzing the relationship of both towards language. Henry James and Frege both share deep worries regarding the ordinary forms of expression and the way it fosters degeneration. Throughout the essay, Boyce discusses the relationship between New Criticism and Logical Positivism, and invites us to acknowledge the deep connection between modernist literature and philosophical enterprises.

The second part of the *Companion*, *What is Literature*, begins with *M.W. Rowe's* essay on literature. Rowe first explains the special kind of aesthetic attention needed to attend to a work of art and explains how it leads to aesthetic pleasure. He then moves on to explicate literature's capacity to convey knowledge. Unlike scientific knowledge, which is objective, literature offers humanistic knowledge, that is, subjective knowledge about how something seems to someone from a certain point of view. Literature gives knowledge through pleasure, but only provided that the reader is fully engrossed and absorbed in a work. *Robert Chodat* discusses novel and the way it intensifies the ancient quarrel between literature and philosophy. On his view, the novel is different from other literary forms (particularly epic) in that it does not present moral heroes but instead raises religious questions, pushing forward metaphysical agnosticism and challenging our moral duties. The novel, unlike works in philosophy, does not treat individual's actions and decisions as isolated events but rather places them into the stream of life showing how particular decisions are made in particular circumstances. *Anna Christina Soy Ribeiro* defends poetry's connection to truth, claiming that the language of knowledge used to be poetic – from poets and philosophers to scientist and mathematicians, all sorts of insights were delivered in poetic forms. This changed when the invention of press and similar tech-

nological achievements made the process of learning easier and diminished the need to memorize texts. If anything, Ribeiro claims, given how poetry makes various potentialities of language salient, such as phonetic, syntactic and semantic features, it enlarges our potentialities for thought, feeling and expression. *Susan L. Feagin* compares and contrasts two accounts of plays, the one which treats them as dramatic works that have to be read in order to be appreciated, and the one that sees them as scripts, i.e. works that need to be performed in order to be appreciated. Explaining the development of the modern theatre, she argues that, given that performances nowadays deviate from the scripts (as explained by James Hamilton's model according to which scripts are but one ingredient in the production), they can no longer be taken as reliable access into works themselves. That however does not mean they cannot be enjoyed as performances. In fact, the reading of scripts and the viewing of performances are mutually supportive and we can appreciate both modes of gaining access to the work, since both are capable of offering new insights. *Aaron Meskin* provides an intriguing overview of the philosophical challenges put forward by popular fiction, showing that the distinction between literary fiction and popular fiction is a substantial one, non-reducible to art/non art or good fiction/bad fiction dichotomies, and, due to the complexities involved in defining, interpreting and evaluating it, deserving of and inviting more philosophical interest. *Ted Nannicelli* defends screenplays's literary status by showing that the two most commonly held arguments against such status are wrong. According to the Ingredient hypothesis, screenplays are not independent, autonomous objects but constituent parts, belonging to a wider production context and ingredients in the motion pictures with which they are associated. According to the Incompleteness hypothesis, screenplays cannot stand as a work of art prior to the process of production and the creation of a motion picture associated with the given screenplay. Drawing on the discussion on the ontological status of theatrical scripts, offering an array of unfinished literary works, and showing how, for any given definition of literature, there are at least some screenplays that can satisfy it and therefore classify as literature, Nannicelli insightfully discards both hypotheses, urging philosophers to start taking screenplays seriously as literature. *Steven Davies'* contribution to the *Companion* is a masterfully succinct and immensely informative summary of the varieties of views that count as evolutionary approaches to literature. Two main ones are evocriticism, i.e. the application of the theories of evolutionary psychology, socio-biology and the like to the interpretation of literature, and Literary Darwinism, which sees literary behaviour grounded in our evolved human nature either as adaptation or as a byproduct of some other adaptation. Davies pointedly reveals problems with these views, carefully delineating issues that are sometimes neglected by proponents of evolutionary approaches, but that need to be settled and further clarified if evolutionary approaches stand a chance as plausible explanation of our literary practices. *Stein Haugom Olsen* discusses concepts of canon and tradition and shows how they relate to each other, as well as to (the notion and practice of) literature itself. The origin of the concept of canon is to be explained with reference to sacred scripts, which reveals two of its main aspects: authenticity and authority. Within literary criticism and literary theory, the

notion of canon can be employed in three different senses, as writings of a secular author accepted as authentic, as a sanctioned or accepted group or body of related works and as equivalent to the concept of literature itself. Olsen is extremely meticulous in explaining differences in these three senses and the role that authenticity and authority play in each, particularly within the wider context of the value and definition of literature as such. Finally, he gives an account of literary tradition, explaining how it differs from the canon in virtue of being tied to a practice (understood as a way of doing things, a way of writing, painting and a way of reasoning that has built into it a set of standards and a notion of skill), of having a continuity, in being anonymous, in the sense that there is no authority responsible for its creation and development, and in being culturally embedded.

Matthew Kieran's insightful essay on creativity introduces the third part of the *Companion*, *Aesthetics and Appreciation*. Pointing to the numerous examples of genre fiction, Kieran shows how authors can be creative without thereby being original, thus refuting the traditional views, according to which the conditions of creativity are novelty and value. He discusses the views of Margaret Boden, Berys Gaut, Noël Carroll and David Novitz and confronts these to the Colleridge-inspired view according to which creativity is related to irrationality. Particularly intriguing part of the essay concerns Kieran's take on the empirical data regarding the connections between creativity and mental disorders. Kieran concludes by outlining and defending the view of creativity as a virtue. *Paisley Livingston* surveys three dominant accounts of authorship. According to the causal conception, authorship amounts to performing certain kinds of actions (writing, composing) and deciding that the work has been completed. On this conception, authorship is reducible to the actions that proximately cause a work to be created. Attributionist conception, defended by Foucault and Barthes and criticized by analytic philosophers, is premised on the claim that writer's or composer' action do not suffice for authorship, as a system of authorial attributions is also required. Finally, on the fictionalist conception of authorship, the actual author does not matter for the interpretation and appreciation of a text, as it might be the case that readers find the text more interesting when they focus on the implied author, i.e. an author based on the features of the text that does not have to correspond to the actual author. *Peter Lamarque* brings together two views on poetic expression, the Romantic view, according to which poem expresses author's personal emotions and experiences, and the Modernist view, which grounds the autonomy conception of poetry by divorcing author's intentions, emotions, experiences and attitudes from the writing itself. Lamarque analyses an array of views on poetry and expression and concludes by offering an account which professes evidence of convergence of the two views. In his immensely informative essay, *Wolfgang Huemer* discusses theoretical approaches to style (style as choice, style as signature, style as expression of the author's personality) explaining how each of these figures in the way we think of literature. The choice is to think of style as identical to some features of a text, or to see it as pertaining to the actions performed by the writers and account for it in psychological terms. However, all of these theories discuss style in descriptive sense, which, Huemer claims, fails to do justice to the style as an

aesthetic category. As Huemer sees it, style 'distinguishes', it grounds aesthetic value and gives rise to the evaluative judgments we make in reference to literary works. He concludes by claiming that a reflection on style can be fruitful for the way we think of and analyze questions regarding the cognitive value of literature, the role of emotions in fiction and various others issues that sometimes do not take into consideration aesthetic dimension of literature. *Eileen John* discusses theme. Themes articulate what literary works are significantly or importantly about, and this relates to the operation of a theme within the work, where it emerges from the work's subject as that which unites and structures the work, and to the operation of a theme within overall human concerns, which relates to the fact that literature is often seen as a source of humanly important issues. With reference to the operation of a theme within a literary work, questions emerge regarding the recognition of a theme and the ways it can be expressed. Themes also serve to ground the cross-textual reference, given that different works can share the same theme. Because themes are not fictional, but relate to general human concerns, questions arise regarding the cognitive value of literature. *Garry L. Hagberg* brings together Aristotle's theory of character (from *Nicomachean ethics*), Raimond Gaita's doctrine of 'truth as a need of the soul' and Bernard Williams' views on the process of making sense of oneself, of events and actions, in order to explicate the notion of a character. His main idea is that the process of understanding one's own, and other people's character is the same as the process of understanding fictional characters, and both arise from our desire to know the truth, i.e. to know what is true, real and genuine, and what is fake, counterfeit and false. Against the background of *Othello*, Hagberg shows how the fact that both, our ethical deliberations (in the widest sense of the word) and literature use language to articulate things, is to be taken as illuminating on the fact that literature can help us come to terms with our true selves. Life and literature are connected, and the insights we gain from literature are applicable to our lives, with the patterns of influence going the other way as well. We can make sense of fictional characters because we recognize in them real people. *John Gibson's* contribution is on empathy. As he defines it, empathy is a form of imaginative, essentially other-directed perspective taking which makes possible an especially intimate and powerful form of identification with another human being, where one comes to identify oneself with the object of empathy. Gibson first explains the relation between empathy and sympathy and proceeds to explain his own account, which is deeper and more encompassing then the accounts currently at disposal (the simulation theory or 'in your shoes' theory). Central to his account are imagination and narrative, which work together in order to enable one to achieve the target's first person perspective, that is, to feel 'as and because' another person does. In the second part of the essay, Gibson explains how theoretical insights regarding empathy matter to philosophical inquiries about literature, appreciation and the reading experiences. Two areas of overlap are prominent: first, empathy is relevant for our capacity to emotionally connect to the characters, particularly those which are morally dubious. Second, the notion of empathy matters for the way we think about literature's cognitive value. It is widely held that literature offers knowledge about the possible

experiences, and via empathy, we can come to learn something about what it feels like to undergo these experiences. In this context, Gibson adds another layer: empathy enables us to connect not only to the characters who undergo these experiences, but with what he dubs the perspective of a work itself. He ends by discussing potential problems with his theory. *Damien Freeman'*s essay deals with the paradox of fiction. Freeman's approach is original in that, after briefly explaining the paradox and charting the main theoretical reactions to it (those offered by Walton, Lamarque, Carroll and Smith, and what he calls, pseudo-Coleridge), and the problems these theories face, Freeman concentrates on the notions of paradox, emotion, response and fiction, and even offers a succinct but informative survey of the philosophical theories of emotions. In addition, Freeman shows how the point of the paradox changes, depending on the approach taken. Aristotle's was a practical approach, in that in *Poetics*, he gives prescriptions on how to write tragedy so as to enable the audience to experience those emotions considered appropriate for tragedy. With Hume, the problem becomes psychological, given that he was primarily concerned with explaining why human psychological set up is such that we take pleasure in tragedy. Colin Radford's approach is logical, since, what is really at stake on his view is a contradiction. The problem of fictional reactions to fiction is pressed even further by *E. M. Dadlez***,** whose focus is on the paradox of tragic pleasure. Not only tragedy, but genres such as horror, melodrama and suspense give rise to pity, fear and distress, prompting the question of why we enjoy these genres, that is, why do we willingly subject ourselves to painful emotions. Tracing the problem back to its origins in Aristotle, Dadlez compares and contrasts two main approaches to the paradox. First, to eliminate charges of incoherent emotional reactions, some philosophers claimed that pleasurable and painful emotions are directed towards different objects. The second approach, originating in Hume, is centred around the claim about the causal dependence of the positive emotion on the negative one. Dadlez concludes that neither conceptual nor causal account can alone provide the full explanation for the unique ability of these genres to arouse negative and positive emotions. Therefore, rather than focusing on shortcomings of these approaches, we should recognize ways in which they demonstrate the potential of literature to enrich our lives. *William P. Seeley* turns to neuroscience to explain our engagement with literature, particularly our capacity to understand narrative stories, formulate expectations about what will happen, construct fictional world beyond the outline provided by narration, come up with moral evaluation of the characters and their events and experience emotional appraisals of what we read. The introduction of neuroscientific research into domain of art is relatively recent, but Seeley offers abundant evidence to support his claim that philosophy of literature can only profit from opening itself up to insights from this field. As it turns out, our capacity to understand stories is underlined by the same recognitional mechanism that enables us to understand events and people in the real world. Seeley explains how the knowledge and familiarity with the genre and conventions guides our engagement with works of art.

The first essay in the fourth part of the *Companion*, *Meaning and Interpretation*, is *Noël Carroll'*s, who provides an account of what is distinctive of

narratives as opposed to other types of representations, fictional and non-fictional, such as state and event descriptions, series, annals, and chronicles. Carroll defends a causal account of narratives, according to which something is a narrative if it is a representation, in a temporally perspicuous, forward oriented ordering, of at least two events and/or states of affairs concerning the career of at least one unified subject where the earlier events in the representation are portrayed as causal contributions to the later events. To shed lights on a much debated question of causality as a criterion for narrative, Carroll introduces the notion of a narrative connection, and claims that in order for something to count as narrative, it has to exhibit this kind of connection, which comes in degrees and depends on the genre. Narrative is also the focus of *Daniel D. Hutto*, who discusses narrative understanding. Contrasting it with theoretical and logico-scientific understanding, Hutto claims that narrative understanding is a sui generis type of understanding, achieved by making sense of happenings and their significance by situating them within a wider range of possibilities. Narrative understanding is not reducible to a causal account, it is primarily and distinctively concerned with particulars, contextualizing them and casting them in a certain light, revealing personal perspectives taken towards them. Therefore, narrative understanding is closely connected to our everyday practice of making sense of a people's actions in terms of reasons. Hutto's essay is immensely informative on the contemporary debates on narrative understanding, theories of mind and folk psychology. An array of issues connected to interpretation is displayed in another of *Noël Carroll*'s essays. Carroll starts off by providing a role for interpretation in literary criticism, evaluation and appreciation, and proceeds to discuss more contentious issues: is there one single right interpretation or are there multiple, perhaps even contradictory, acceptable ones, the status of interpretive claims, constructivism with reference to the objects of interpretation etc. The question of the relevance of authorial intentions is at the centre of Carroll's essay. Defending actual intentionalism against the value-maximizing approach to interpretation and hypothetical interpretation, Carroll ends by masterfully summarizing arguments for and against each of these theories. *Stephanie Ross* discusses the role of criticism in literary practice. Taking the practice of restaurant and movie recommendation as instances of good criticism, Ross starts by asking what resources are available to a reader wondering what works of literature to consume. She first offers summary accounts of the four most discussed views on criticism. First, taxonomies developed by Monroe Beardlsey, and in more recent times, Noël Carroll, who both share the idea that via various critical activities (on Beardsley's account, these include explication, elucidation and interpretation, on Carroll's, description, classification, contextualization, elucidation, interpretation and analysis), critics provide a reasoned, summary evaluation of a work. On Isenberg and Sibley's accounts, critics' task is to help appreciators detect the aesthetic qualities. Wollheim's 'criticism as retrieval' sees the goal of criticism as a reconstruction of creative process that lead to the creation of a work. Finally Hume's account is focused not on the practice and process of criticism but on the persona of an ideal critic. Ross ends by offering her neo-Humean account which unites insights from all the theories examined,

but is only useful provided readers only seek recommendations from critics whose taste they share. Accordingly, on Ross' account, Hume's trained, sensible, experienced, prejudice-free critic (who also posses emotional responsiveness and imaginative fluency) is best suited to recognize and alert readers to value-making properties, to conduct interpretive activities and reconstruct artworld influences. *Ernie Lepore* and *Matthew Stone* bring the resources of the philosophy of language to clarify the status of poetry in a broader account of speakers' knowledge of language and their linguistic practices. Crucial in their account is the notion of poetic imagination, which they define as a specific kind of interpretive engagement that poetry demands. Lepore and Stone compare poetry with quotations, claiming that both privilege and problematize linguistic form in relationship to meaning. Once an utterance is understood as poetic (which, on their view, is a matter of employing distinctive kind of interpretative practices), one's attention is turned to how the poem is articulated, i.e. to the formal organization of the poem itself. Understanding the difference between interpretive practices is important for understanding the distinctive experience of poetry within philosophy of language. The main argument in the essay concerns the special kind of insights that poems deliver, insights which are not to be understood as the contents of any level of linguistic meaning, including pragmatic level of meaning, and are prompted by our focusing on the form and content of a poem. *Elisabeth Camp's* essay, though primarily focused on metaphors, is much wider in scope and concerns more general questions having to do with deciphering and understanding the meaning of a literary work, the relation of author and readers mediated via text and the potentials of literature to deliver cognitive benefits to readers, primarily when a literary work is seen as showing a distinctive perspective towards the world. Camp claims that metaphors in literature do not differ from metaphors in other contexts, and briefly describes how various schools of criticism treated the problem of interpreting metaphors. According to her account, metaphors are one among three poetic rhetorical devices used to present perspectivally laden contents (other two are exemplifications, which include telling details and stories, and thick terms, such as stereotypes and slurs) which are often open-ended, evocative, experiental and/or imagistic. Metaphors differ from other tropes in that there is typically no explicitly available content which the author endorses, since the metaphor's proffered content is available only indirectly, and in their twofoldedness, which includes an experiential awareness of the representing frame and the represented subject. The notion of twofoldedness is given a more detailed explication, as Camp sees it as having both, cognitive and aesthetic consequences.

Part five, *Metaphysics and Epistemology* brings first *Amie L. Thomasson*'s essay on the ontology of literary works. Separating ontological issues –those concerning existence, survival, identity conditions and modal properties of literary works – from issues regarding the definition of literature and evaluative questions concerning good vs. bad literature, Thomasson informatively summarizes leading views on the ontological status of literature, giving pros and cons of each. She discusses theories put forward by Roman Ingarden, Richard Wollheim, Nelson Goodman, Nicholas Wolterstorff and Guy Rohrbaugh, to show difficulties involved in explaining

whether a literary work is a material object to be identified or equated with some specific copy of a work, or some kind of mental entity related either to author's or reader's mental states. The so-called action-centered theories, primarily Gregory Currie and David Davies', associate a work with types of actions or individual token actions, but Thomasson offers reasons to doubt these theories. With her discussion of Peter Lamarque's theory, questions concerning the ontology of literature are expanded to include cultural contexts within which works are created. In the second part of the essay, Thomasson discusses more general, methodological issues relating to ontology of art generally and even some meta-ontological discussions concerning the status, aim and appropriateness of philosophical discussions of ontology. In his third contribution, *Noël Carroll* discusses fiction and the multiple ways in which fiction is distinguished from nonfiction. Among others, these include identifying fiction via formal devices employed in fictional narratives (such as free indirect discourse, where authors provide us with information that is, in any other contexts, unavailable to us, such as reports of the character's state of mind and thoughts), via the content (fiction is about possible world, nonfiction about the actual) or via the relation to truth (fiction does not aspire to be true while nonfiction does). However, none of these devices is exclusively employed in fiction, which shows that a distinction between fiction and nonfiction cannot be grounded in any of them. Another set of suggestions on how to distinguish the two is inspired by pragmatist accounts. These include Searle's pretended speech act theory and Beardsley's fiction as imitation of illocutionary speech acts, which were criticized by, among others, Ken Walton, who offers his 'fiction as make-believe' theory. Pragmatist account also includes theories inspired by Grice's account of meaning and communication. The last two theories Carroll scrutinizes are Derek Matravers and Stacie Friend's. *David Davies* discusses the role that fictional truth plays in our engagement with fiction. One way in which the notion of fictional truth is central to our engagement with fiction concerns the problem of figuring out what is true in fiction. While understanding what is going on in a fictional narrative is continuous to our ability to understand nonfictional narratives, being clear on what is true in the fictional narrative is not straightforward. Not only are narrators sometimes unreliable, deceptive or ignorant, they can also use various techniques, such as metaphors or irony, that can harden our ability to track fictional truth. Another set of worries relates to factors that determine what is properly taken as unstated background, given that fictional narratives are not delineated by what is true in the actual world. Even more, fictional narratives are often incomplete, in that we are not told, and have no way of discovering, all the relevant information. In discussing these issues, Davies summarizes theories which deal with this problem, namely David Lewis' possible world analysis and Gregory Currie's idea that it is the beliefs of 'fictional authors' that determine what is true in the story. The second sense in which fictional truth is problematic concerns the much discussed issue of whether we can learn truth and gain nontrivial knowledge about the world through fiction. Those who claim that we can, literary cognitivits, claim that various kinds of knowledge are available through fiction (knowledge of particular facts, general principles, categorical understanding, affective knowledge) but Da-

vies exposes a wide range of issues that cast doubt on literary cognitivism and ends by analyzing in more details one particular way of trying to salvage cognitivist intuition, namely the clam that fiction is analogous to thought experiments. Literary cognitivism is discussed in more details in *James Harold*'s essay. Harold first explains the relevance of the main idea behind literary cognitivism (the claim that literature is a source of knowledge) for the liberal humanistic education and proceeds to discuss various views that all fall under literary cognitivism. On its strong formulation, literature is a source of propositional knowledge, and those who defend this claim have to first explain how readers deduce relevant propositions from a given work, and then show what makes literature epistemically reliable. On the weaker reading, literature is a source of nonpropositional knowledge, such as experiential and perspectival knowledge. With this interpretation, the problem is to explain how these kinds of cognitive benefits add up to truth and knowledge. Harold expands discussion of literary cognitivism to questions regarding the nature of knowledge itself, the meta-ethical status of ethical claims (since the dominant tendency of cognitivists is to claim that literature is a source of ethical knowledge and that engaging with literature is beneficial for our moral development), the concept of literary value (in order to see whether literature's capacity to insert knowledge adds up to or is neutral toward the value that literature has) and the psychology of reading (to analyze whether reading can in fact potentially be morally corruptive). *Jonathan Gilmore* discusses the role of imagination and kindred phenomena (pretense, make-believe and simulation) in our engagement with and experience of literature. He starts off by discussing similarities between our imaginings and our beliefs, particularly when it comes to figuring out what is true in fiction. Similarities include the fact that our imaginings, like our beliefs, aim at consistency and are dependent on our background knowledge. A distinction between imaginings and beliefs is most evident in the fact that we generally do not tend to act on the basis of what we imagine. In the second part of the essay, Gilmore explains how our imaginings are connected to our emotional reactions to fiction and presents theories on our empathic reactions to characters and the phenomenon of imaginative resistance. This phenomenon is given a profoundly rich treatment in *Tamar Szabó Gendler* and *Shel-Yi Liao*'s essay, who offer not only a probing analysis of what imaginative resistance consists in, but provide us with a detailed overview of the development of philosophical approaches to this problem. As they make evident, there is a great variety of disagreements surrounding the phenomenon, regarding its scope, the mechanisms for evoking it, over its psychological components and over its nature. Gendler and Liao summarize three main approaches to imaginative resistance, describing their development as the first wave of philosophers' tackling imaginative resistance. On Cantian theories, imaginative resistance should be understood as a breakdown of authorial authority, as one simply can't imagine as one has been invited to. On Wontian theories, one will not imagine the relevant content, and on Eliminativists theories, there is no such thing as imaginative resistance per se. The second wave, beginning in 2010, is characterzed by the turn towards questioning substantive and methodological assumptions of the first wave's theories. In the second part of the essay, authors

describe the relevance of imaginative resistance for moral psychology and meta-ethics, for the questions of fiction's capacity to morally educate and for the ethical criticism of art, for our cognitive architecture and for modal epistemology. *Gregory Currie* closes this section by explaining the connection between literature and the theory of mind, i.e. our capacity to think about, manage, monitor, manipulate and take into consideration other people's mental states (beliefs, desires, emotions). Currie refers to this capacity as mentalising and claims that it is a central feature of our social and ethical life, as well as a central feature of literature. His main argument is that literary works are for the most part concerned with mentalising activities people engage in on an everyday basis (he sees *Gilgamesh* as the only exception to this). Currie's discussion is set against a wider background of evolutionary approaches to literature, as he is concerned with accounting for the enjoyment readers take in the literary representations of mentalising. He discusses several potential explanations (we enjoy mentalising fictions because it portrays characters similar to us, there are adaptive advantages conferred by a taste for fictions which represent mentalising, fiction improves mentalising and capacities which depend on it, enjoyment we take in mentalising fiction is a by product of some other adaptive development, preference for mentalising fiction was selected for but it is no longer an advantage) and concludes that more research is needed to support the traditional humanistic belief into fiction's capacities to influence, improve and refine our cognitive capacities, particularly those pertaining to the theory of mind.

The last part of the *Companion*, *Ethics and Political Theory*, offers first *A.W. Eaton*'s essay on the ways in which literature and ethics are related. Eaton focuses on two main topics that figure in this context within the Anglophone philosophy of art: literature's morally relevant influences and the relationship between moral and aesthetic value. Regarding literature's capacity to morally influence readers, the core question is what makes literary works the proper object of moral judgment. Eaton offers an array of theories designed to explain literature's impact on readers (the catharsis model, the imitation model, the conditioning model and the literary moral cognitivism) and concludes by showing how this discussion is further complicated, relative to whether one takes the empirical-casual (i.e. descriptive) or interpretative-teleological (i.e. normative) take on it. Regarding the much debated question of the connection between a work's moral and aesthetic value, Eaton compares and contrasts dominant views (moralism, ethicism, moderate moralism, immoralism, cognitive immoralism, robust immoralism) underlying the main arguments adduced in support of each. *Espen Hammer* discusses the Marxist literary criticism, explaining its origin in the ideas of Karl Marx, and tracing its influence on the subsequent literary theories, particularly those developed by Althusser, Lukacs, Hauser, Adorno, Bloch, Sartre, Barthes, Foucault, and Jameson. Hammer's key claim is that Marxist literary criticism offers powerful and relevant tools for reflecting on the relation between literature and society. As he explains, a dominant tension in Marxist approaches to literature is a dual function attributed to a literary work: it is understood as an ideological representation of the ruling class's interest, and a reflection upon the society and therefore a source of

significant moral critique. Given Marxist belief that art originates within a certain socio-historical background, they reject the idea that literature expresses universal human concerns and values, and instead argue that the proper understanding and interpretation of literary works is only possible if one is familiar with the culture in which the work was produced. *Paul C. Taylor* develops a framework within which to consider the overlap of philosophical interest in literature and in race. Though seemingly literature and race have little in common and are rarely brought together under the same philosophical umbrella, Taylor shows just how immensely valuable insights are gained when these two fields are scrutinized jointly, insights having to do with literature, formation of canon, evaluative judgments, culture, peoples, race and race-related sets of problems, both theoretical and practical (the part of the essay where Taylor summarizes philosophical race theory is particularly insightful for these matters, as it helps to situate philosophical approaches to race and racialism into a wider social and political context). Mostly however, Taylor's analysis probes into the very foundation of race, and how we come to think of, and judge, different races and their artistic achievements, making us reconsider the connection between literature and civilization, civilization and barbarism, humanity and non-humanity when the notion of being human is judged from the perspective of one dominant culture. The essay is wider in scope than issues regarding colonial and post-colonial literature reveal, as Taylor discerningly brings together an array of questions that are relevant for the literature-race nexus, questions concerning the racialized meanings of literacy, the availability of particular linguistic resources across racial boundaries, questions about whether and how to open previously closed traditions to each other, how to change the literary cannon so as to insert capital works from other cultures, questions having to do with interpretations of works from other cultures, ethical dimension of literature that is or contains racists elements, etc. The final essay in the *Companion* is *Mary Bittner Wiseman's* discussion of literature and gender. Focused primarily on the female perspective, Wiseman is concerned with two notions: the experience of reading and the distinctive ways in which, via engaging with literature, women can come to recognize different possibilities for who they can be. She starts off by claiming that literature has the power to change readers' beliefs, including those beliefs pertaining to the role that gender plays in one's sense of the self, and dedicates the first part of the essay to the analysis of the words' meaning and the meaning-making practices embedded in Saussure's tradition and developed in the Barthes' idea that full engagement with the work is possible only if one reads it as a producer of, rather than as a consumer of meaning. With this in mind, Wiseman proceeds to her account of the female reader, and, drawing upon her own reading experiences, describes ways of engagement with literature women can embrace in order to create their own meanings. Because there is a sense in which reading implies taking the perspective of a man, paying close attention to what words could mean and becoming active producer of these meanings enables women to imaginatively discover all that women have been, felt and done.

IRIS VIDMAR
University of Rijeka, Croatia