

The problem of lexical innovation

Josh Armstrong¹

Published online: 3 March 2016
© Springer Science+Business Media Dordrecht 2016

Abstract In a series of papers, Donald Davidson (Synthese 59(1):3–17, 1984, The philosophical grounds of rationality, 1986, Midwest Stud Philos 16:1–12, 1991) developed a powerful argument against the claim that linguistic conventions provide any explanatory purchase on an account of linguistic meaning and communication. This argument, as I shall develop it, turns on cases of what I call *lexical innovation*: cases in which a speaker uses a sentence containing a novel expression-meaning pair, but nevertheless successfully communicates her intended meaning to her audience. I will argue that cases of lexical innovation motivate a *dynamic* conception of linguistic conventions according to which background linguistic conventions may be rapidly expanded to incorporate new word meanings or shifted to revise the meanings of words already in circulation. I argue that this dynamic account of conventions both resolves the problem raised by cases of lexical innovation and that it does so in a way that is preferable to those who—like Davidson—deny important explanatory roles for linguistic conventions.

Keywords Conventional meaning · Linguistic communication · Lexical innovation · Semantics-pragmatics interface

1 Introduction

Theoretical discussions of language have long seen important explanatory connections between natural language semantics and social convention. It has, for example, been widely maintained that the semantic properties of the words in a

✉ Josh Armstrong
jarmstrong@humnet.ucla.edu

¹ Department of Philosophy, UCLA, Box 951451, Dodd 321, Los Angeles, CA 90095, USA

language depend upon the linguistic conventions that prevail among a population of language users. Likewise, it has been widely maintained that background linguistic conventions tightly constrain the ways in which speakers can use language in interpersonal communication. I will call the conjunction of these claims about semantic dependence and linguistic communication *semantic conventionalism*.

In recent years, many have come to reject semantic conventionalism as fundamentally mistaken. My focus in this paper is on one kind of consideration—famously developed by Davidson (1984, 1986, 1991)—that has led to this trend. The argument turns on cases of what I call *lexical innovation*: cases in which a speaker uses a sentence containing a pairing of a basic expression and a meaning that is novel to the speaker or hearer of that utterance.¹ As we shall see, lexical innovations are ubiquitous in our everyday speech and regularly result in episodes of successful linguistic communication. These facts seem to show, and indeed have been widely taken to show, that there is no deep explanatory connection between linguistic meaning and social convention.²

In this paper, I will argue that a commitment to a robust explanatory role for semantic conventions is fully consistent with the existence of lexical innovations. What is required, and what I will provide, is a better account of the nature of linguistic conventions. In particular, I argue that cases of lexical innovation motivate a rejection of the widespread assumption that the semantic conventions are stable bodies of facts whose properties are largely determined independently of the speaker and hearers engaged in a conversation. In place of this static picture, we must invoke a *dynamic* account of linguistic conventions.³

The dynamic account I develop makes two related claims concerning the nature of linguistic conventions. First, linguistic conventions are *adaptive*: prior conventions may be rapidly expanded to incorporate new expression meanings or revised to shift the meanings of expressions already in circulation. Second, linguistic conventions are *local*: different semantic conventions may prevail among groups of agents within a wider population, and individual speakers may utilize different semantic conventions when engaging in conversations among different groups. I show how each of these features of linguistic conventions are natural byproducts of the influential theory of convention developed by Lewis (1969). The resulting account is, I argue, flexible enough to explain what is communicated by lexical innovations, while also capturing the pervasive constraints on the manner in which this process unfolds.

As will become clear, my account offers a new perspective on the relationship between semantics and pragmatics. The account I develop maintains that the

¹ I borrow this use of the term innovation from Gleitman and Gleitman (1971).

² See Nunberg (1979), Moravcsik (1998), and Carston (2002a, b) for related considerations.

³ The account I develop builds upon work by Bartsch (1984), Clark and Clark (1979), Clark (1996), and Ludlow (2014). It should be noted that my use of the term ‘dynamic’ differs from the sense in which it is used to denote descriptive approaches to semantics, such as those developed in Kamp (1981), Heim (1982), Groenendijk and Stokhof (1990), and others according to which linguistic meanings are specified in terms of “context change potentials.” As we shall see, while I utilize some aspects of these approaches to semantics (particularly the notion of conversational context), my account is neutral on the issue of whether semantic values should be identified with context change potentials.

comprehension process in cases of lexical innovation requires audience members to integrate distinctively structural information about linguistic form with distinctively social information about speakers' intentions. What distinguishes information that is semantically expressed from information that is merely pragmatically conveyed is thus not the fact that the latter depends on identifying speakers' intentions and the former does not. Rather, semantics and pragmatics are distinguished by the kinds of effects speakers intend their utterances to have, and the manner in which those effects are realized.

My plan for the paper is as follows: In Sect. 1 I will develop a version of conventionalism about semantics that I think is both widely accepted and *prima facie* plausible. In Sect. 2 I outline the problem that lexical innovations raise for conventionalism and I consider, and reject, two ways of explaining away the problem. In Sect. 3 I develop my dynamic account of linguistic conventions. In Sect. 4 I show how this dynamic account of linguistic conventions resolves the problems raised by lexical innovations. Finally, in Sect. 5 I elaborate some of the consequences of my view for standard debates in the philosophy of language.

2 Semantics and convention

The discussion to follow concerns the extent to which facts about linguistic meaning and communication can be illuminated by an appeal to social convention. I will begin by clarifying some of my terminology. I will then elaborate two core elements of the conventionalist account—what I call *the DETERMINATION thesis* and *the LINKING thesis*—and clarify their scope. This will then set the stage for the discussion of lexical innovation in Sect. 2.

2.1 Terminology

In what follows, I will assume that a semantic theory for a natural language is a *systematic* way of pairing meanings with expressions of that language. That is to say, a semantic theory for a language specifies both the meanings of the basic expressions in the lexicon of that language (be they individual morphemes, words, or idioms) as well as specifying a set of combinatorial principles that serve to generate the meanings of complex expressions (be they phrases, sentences, or even discourses) on the basis of the meanings of their parts and their linguistic structure.

For the purposes of my discussion, I shall remain as neutral as possible concerning the underlying nature of linguistic expressions and the nature of linguistic meanings. I will, however, assume that linguistic expressions (be they sounds or signs of some other kind) are individuated in a way that is sensitive to both their phonological properties and their lexical or grammatical category. I will also assume that a semantic theory for a natural language associates both *standing-meanings* and *occasion-meanings* with the expressions of that language. The standing meaning of an expression is a context-invariant constraint on the range of meanings to which occurrences of that expression can be assigned relative to contexts of use or, in an alternative idiom, the range of meanings that utterances of

that expression can be assigned.⁴ The occasion meaning of an expression is the particular meaning (or set of meanings) that an expression has relative to a particular context of use or a particular utterance event. Following Kaplan (1989a), I take it that this distinction can be fruitfully modeled in terms of *character* and *semantic content*: the character of an expression is a function from contexts of utterance to semantic contents, and semantic contents are in turn functions from circumstances of evaluation to extensions.

When I speak of “linguistic communication” I have in mind the intentional use of language for the purposes of interpersonal communication—communication which requires speakers to intentionally convey contents with their utterance, and which requires audience members to identify the same (or some sufficiently similar) content to the one the speaker intended to convey. Following much recent work, I will assume that linguistic communication can be described as a process in which speakers are intending to update the *common ground* or conversational scoreboard—a structured body of information that is presumed to be mutually accepted by the members of the conversation.⁵ Accordingly, successful linguistic communication requires audience members to correctly identify how a speaker intended to update the common ground with her linguistic utterance.

2.2 Determination and linking

At its core, conventionalism is the thesis that there is an important sense in which a descriptive semantic theory for a language, understood along the foregoing lines, depends on the attitudes and actions of a community of language users. In particular, conventionalism grows out of the idea that a semantic theory for a language is grounded in historically specific practices that perpetuate among members of a group of language users because of the shared interests and mutual understanding of those agents—that is to say, as a matter of social convention. Accordingly, a semantic theory for a language characterizes what competent speakers know, and expect other speakers to know, about the meanings of the expressions of their language.⁶

It is worth distinguishing two related claims made by the conventionalist. The first is what I’ll call the DETERMINATION thesis:

⁴ I do not intend to adjudicate the debate between those who, like Kaplan (1989a), take the bearers of linguistic content to be occurrences of expressions relative to contexts and those who, like Barwise and Perry (1983) take the bearers of linguistic content to be utterances. Rather I shall somewhat opportunistically move back and forth between these two formulations hoping that readers will not begrudge me a grain of salt and will translate my terminology into their favored framework accordingly.

⁵ See, for example, Stalnaker (1978) and Lewis (1979). In important recent work, Murray (2014) provides a generalization of this framework that is suitable for characterizing the way non-assertive and non-at-issue contents serve to update the common ground.

⁶ See Lewis (1969), Schiffer (1972) and Higginbotham (1992). The claim that a semantic theory is grounded in iterated mental states—whether the relevant mental states are taken to be states of common knowledge, states of common belief, or, simply common reason-to-accept—has been a controversial component of conventionalism. See Burge (1975), Millikan (1984, 1998), Sperber and Wilson (1986), Gilbert (1992), and Skyrms (1996) for critical discussion and Clark (1996) and Stalnaker (2002) for replies.

DETERMINATION: The standing meanings (or characters) of linguistic expressions are determined by social conventions that prevail among a group of language users.

This thesis has often been motivated by the claim that individual speakers do not endow expressions with standing meanings; rather, linguistic expressions become associated with meanings on the basis of the way those expressions have been used within a wider speech community.⁷ As Searle (1969) memorably put the point, the meaning of a word depends on more than a speaker's intentions.

Conventionalism also involves a claim about the relationship between the contents that speakers communicate with their linguistic utterances and the semantic content of those linguistic utterances. I will call this the LINKING thesis:

LINKING: All things being equal, the content communicated by an utterance of S in context C is the semantic content of S in C.⁸

This is to say, unless audience members have reason to think otherwise, they will take the speaker to have intended to update the common ground with the semantic content of her utterance. The LINKING thesis is typically motivated by desire to explain the fact that linguistic communication is typically both *efficient* and *productive*. To say that linguistic communication is efficient is to say that speakers can trust that audience members will be able to successfully identify the intended meanings of their linguistic utterances, and that audience members can expect that they will be able to successfully interpret a speaker's linguistic utterance as the speaker intended. To say that linguistic communication is productive is to say that it allows speakers to generate, and audience members to understand, the meaning of an indefinite number of sentences—including sentences that the members of a conversation have never previously encountered.

It is worth noting that the conventionalist has no reason to deny, and indeed good reason to accept, the claim that the ability of groups of agents to coordinate on a semantic theory in particular—or on a grammar in general—is made possible by the fact that they are endowed with a species-specific language faculty: an innate cognitive mechanism that substantially constrains the set of lexical categories and the combinatorial principles of syntax and semantics.⁹ Accordingly, the language faculty will serve to structure and restrict the kinds of coordination problems for which semantic conventions arise as solutions. Indeed, while conventionalism does require that members of a linguistic community utilize aligning principles of syntactic and semantic combination, the primary role of linguistic conventions is in

⁷ Kaplan (1989b, pp. 600–603).

⁸ Versions of this principle are defended in, among others, Stanley (2000, 2002, 2005), King and Stanley (2005), Devitt (2013), Lepore and Stone (2014); c.f. the principle of *Interpretive Economy* developed in Kennedy (2007).

⁹ See Guasti (2004), Isac and Reiss (2008), and Crain and Pietroski (2012) for instructive summaries of this evidence.

determining the meaning of the basic expressions in the lexicon.¹⁰ And there is reason to believe that conventions have a role to play here, for both the set of expressions in the lexicons and the meanings they are assigned are not universal among the world's languages.

Likewise, conventionalism does not entail that linguistic communication is simply a matter of identifying the semantic content of a speaker's utterance. It is, for example, compatible with the LINKING thesis that linguistic communication often requires audience members to identify contents over and above the semantic contents of speakers' uses of language. Indeed, it is compatible with the LINKING thesis that the semantic content of a speaker's utterance is not among the contents that the speaker is intending to communicate. But what the LINKING thesis does require is that these "pragmatically" generated contents are *principally motivated* departures from the semantic content of the utterance. It requires, in other words, that the semantic content of a linguistic utterance is the default content communicated—the content communicated by an utterance unless auxiliary principles of pragmatics would be violated or otherwise instruct audience members to find a content that the speaker intended to convey but did not semantically express.¹¹

3 The problem of lexical innovation

The forgoing account of linguistic meaning and communication has been the subject of much debate in recent literature in the philosophy of language and semantics. The primary focus of this debate has been on the relationship between the standing meanings of linguistic expressions and the contents of those expressions on occasions of use. A number of theorists have argued that the contents of many expressions vary in open-ended and unprincipled ways with background features of the contexts in which these expressions occur; as a consequence, it is maintained that the semantic contents of those expressions radically underdetermined what speakers express in using those expressions.¹² If successful, these arguments would

¹⁰ By allowing that the principles of syntactic and semantic combination are largely innate, and hence not fully established by convention, we can avoid the "meaning without use" problem raised by Schiffer (1993), Hawthorne (1990), and others. The problem is that if conventions primarily link sentences with their propositional contents—rather than atomic expressions and their standing meanings—the fact that speakers can generate and understand the meanings of a potentially infinite number of sentences seems inexplicable. While the move to subsentential conventions raises a number of issues, including issues concerning the relation between personal and subpersonal psychological states, a detailed discussion is beyond the scope of the present paper. See Loar (1981, pp. 256–260) and Davis (2003, Chap. 10) for relevant discussion.

¹¹ In this sense, the LINKING thesis is fully compatible with what Grice (1975) dubbed *implicatures* having a central role to play in episodes of linguistic communication.

¹² See Searle (1978), Travis (1985), Sperber and Wilson (1986), Atlas (1989, 2005), Recanati (1989), Bach (1994), Bezuidenhout (2002), Carston (2002a, 2b), and Soames (2009) for representative statements of this position; see Stanley (2000, 2002), Szabo (2001), King and Stanley (2005), and Kennedy and McNally (2010) for replies and critical evaluation.

show that shared semantic conventions are insufficient to explain straightforward cases of linguistic communication.

My focus in this section will be on a different kind of problem—one that challenges the claim that linguistic meanings are in any sense determined by linguistic convention and whether knowledge of conventional meaning is even *necessary* in order for successful linguistic communication to occur. After developing the problem in some detail, I will consider and reject a possible reply on behalf of semantic conventionalism as well as Davidson's own alternative explanation.

3.1 The core problem

In a series of papers, Davidson (1984, 1986, 1991) develops a powerful argument against the claim that linguistic conventions provide *any* explanatory purchase on an account of linguistic meaning and communication. The argument, as I shall develop it here, turns on basic considerations concerning the creative aspects of language use. Specifically, the argument turns on cases of *lexical innovation*—cases in which a linguistic utterance contains a pairing of a basic expression and a meaning that is novel to the speaker or hearers of that linguistic utterance.¹³

The basic phenomenon can be illustrated by cases like the following. Suppose that on a visit to your local post office you notice that the attendant has an oddly shaped ruler on the counter. If the attendant subsequently utters (1):

(1) Let me measure the volume of your package with a *koba*,

although you may well have never encountered the expression 'koba' before, you may be able to recognize that an utterance of (1) expressed the proposition that the attendant was going to measure your package with the oddly shaped ruler like the one on the counter.¹⁴ Likewise, if during a walk in the park a large dog jumps up and licks your face, and an exasperated teenager with a leash utters (2)

(2) *Mupsy* loves to greet strangers with a lick.

you can immediately identify that 'Mupsy' was used to pick out the dog in front of you, despite having never encountered that expression before. In each such case, we have an episode of successful linguistic communication without the speaker and audience members having prior shared knowledge of the semantic conventions governing the meanings of the expressions occurring in an uttered sentence.

Speakers can also use words in ways that do not accord with the semantic conventions obtaining among the members of any linguistic community. Suppose that Alf is a guard at a local prison complex who has been given the task of monitoring Bea, an inmate with a penchant for escaping from her cell. Alf returns from his lunch break and finds Bea's cell empty. After carefully inspecting the room

¹³ So understood, innovation is a *relational* property of *uses* of words: it is a feature that occurrences of words in linguistic utterances have for an agent or group of agents.

¹⁴ This example is adapted from Markson and Bloom (1997).

and finding no evidence of how Bea may have escaped, Alf runs to his superior Cal and utters (3):

- (3) Bea managed to *houdini* her way out of her cell.

The expression ‘houdini,’ as it occurs in (3) is a lexical innovation: for we can suppose that the semantic conventions shared by Alf and Cal did not fix the meaning of the verb ‘to houdini’ prior to Alf’s utterance.¹⁵ But it seems clear that Alf’s utterance has a semantic content—roughly, the utterance of (3) expressed the proposition that Bea managed to inexplicably escape from her cell. It also seems clear that Cal may readily identify the intended meaning of (3). This seems to conflict with the core claims of conventionalism, for here we have a case of linguistic meaning and communication without semantic conventions.

The innovation in (3) is not exceptional. It is an instance of a broad class of cases in which a novel linguistic expression is generated from an existing linguistic expression through a process of *zero-derivation*: a process that preserves the phonological shape, but not the lexical category, of the original linguistic expression.¹⁶ For example, the underlying problem could equally well be raised by utterances of (4)–(6):

- (4) A local resident expressed concern that incoming developers were going to *east village* her neighborhood in Brooklyn.
 (5) The delivery boy managed to *porch* the newspaper at every house on the block.
 (6) Pat made sure to *whisky* the punch before the teachers arrived.

But in each case, the sentences contain specific pairings of expressions and meanings that have not been established by prior convention. However, it is not difficult to imagine contexts in which these sentences could be used in episodes of successful linguistic communication.

A final class of cases, similar to the ones on which Davidson himself focused, concern speakers using words innovatively to adjust to the idiosyncratic needs of their audience.¹⁷ Suppose, for example, that a conversation is taking place involving

¹⁵ This example is adapted from Clark and Clark (1979), as is (5) below.

¹⁶ See Bauer (1983), Brinton and Traugott (2005), and Aronoff and Fudeman (2010) for relevant discussion. What makes zero-derivation especially apt for Davidson’s argument is that it is widely agreed that zero-derivation is not fully productive from the point of view of semantics: knowledge of the meaning of a linguistic expression produced through zero-derivation cannot be derived from knowledge of the meaning of the parent linguistic expression and knowledge of the underlying morphological operation alone; see Bauer (2001, Chaps. 3–4) and Jackendoff (2002, Chap. 6) for extensive discussion on this point.

¹⁷ Davidson took such cases to be on par with inadvertent malapropisms, slips of the tongue and other kinds of performance errors (1986, p. 95). I believe this was a mistake: in cases of inadvertent malapropisms and the like, speakers both intend to express a given content and they intend to do so by conventional means. While the former intentions are satisfied, the latter are not. For this reason, I think such cases should be explained in terms of a distinction between speaker meaning and semantic meaning; see Sect. 5.2 of the present paper for further discussion.

an audience member whom the speaker knows uses the word ‘erudite’ to express the property of being erratic. Given this knowledge, the speaker may utter

(7) You should avoid hiring Joe the plumber because he is very *erudite*.

with the intention to use ‘erudite’ to express the property of being erratic. As this case is described, it is clear that the audience member will take the speaker to have expressed the content that Joe is erratic. As before, we seem to have a case of linguistic meaning and communication that was not derived as a function of the conventional meanings of the expressions occurring in that sentence together with their manner of combination.

Lexical innovations raise a serious challenge to conventionalism. Insofar as they involve unprecedented uses of words with genuine semantic content, they seem to conflict with the thesis that word meanings are determined by linguistic convention. The semantic contribution of a lexical innovation does not, in other words, appear to be derived from applying an already given linguistic character (or standing meaning) in context. Likewise, we cannot explain how speakers are able to communicate successfully with their audiences in such cases simply by appealing to the fact that speaker and audiences had prior common knowledge of the convention of a shared language. In short, the phenomenon of lexical innovation appears to undermine both the DETERMINATION THESIS and the LINKING THESIS.

3.2 Conversational implicature

When faced with challenges of the sort raised by cases of lexical innovation, many conventionalists have replied by invoking a familiar distinction between what words in a language mean and what a speaker intends to communicate by using those words.¹⁸ According to this line of response, the meanings that speakers express by using sentences containing lexical innovations are not products of the semantic contents of their linguistic utterances; rather, the meanings that speakers express by using sentences containing lexical innovations are explained by appealing to uncontroversial pragmatic mechanisms such as conversational implicature.¹⁹ If cases of lexical innovation could be treated in such terms, they would pose no threat to conventionalism—for such cases would fall within the purview of a suitable pragmatic theory and hence provide no challenge to the claim that the semantic properties of a speaker’s words are determined by convention.

¹⁸ This general style of response is suggested by Stanley (2005), and more tentatively by Camp (2006), and is developed as a particular reply to Davidson in Reimer (2004).

¹⁹ There are a variety of other pragmatic mechanisms that might be invoked to explain cases of lexical innovation. For example, it could be argued that what’s communicated in cases of lexical innovation is the result of an *explicature*, in the terminology of Sperber and Wilson (1986), or the result of a conversational *implicature*, in the terminology of Bach (1994); indeed, Carston (2002a) and Recanati (2010) have developed accounts of phenomena that parallel cases of lexical innovation along precisely these lines. But these alternative pragmatic mechanisms are not available to the conventionalist who accepts the LINKING THESIS—for they require giving up on the claim that there is even a default presumption that the content communicated by a linguistic utterance is the semantic content of that utterance.

Unfortunately, the claim that lexical innovations should be treated as cases of conversational implicature suffers from a number of serious difficulties. First, lexical innovations differ from conversational implicatures in the way that they are “worked out” or “calculated.” According to Grice (1975), conversational implicatures involve cases in which a speaker intentionally expresses one content with the expectation that audience members will recognize that the content expressed is not the one the speaker was interested in communicating and will, therefore, search for an additional content that the speaker might have intended to convey. But lexical innovations are often not calculated in this way; speakers could, for example, utter (1) or (7) above without expecting their audience to infer some additional content that they were attempting to pragmatically convey but did not semantically express.

Second, the proposed response leads to odd predictions about semantic content—particularly in cases of lexical innovations involving novel denominal verbs. If innovative denominal verbs involved no changes at the level of semantic content—that is, if expressions such as ‘whisky’ or ‘porch,’ as they occur in utterances of (5) and (6), were associated with their standard conventional meanings—the sentences in which they occurred would be straightforwardly ungrammatical and hence not suitable for expressing a semantic content at all. But the sentences involved certainly *seem* to express semantic contents; barring a more principled reason to believe that the uttered sentences do not have semantic contents, this response appears to be no more than special pleading in an effort to save conventionalism.

Finally, and most importantly, lexical innovations appear to combine with other expressions in the sentences in which they occur in semantically significant ways. Their interaction with conditionals, as in (3’), negations, as in (6’), and quantifiers, as in (8) appear no different than non-innovative expressions:

- (3’) If Bea managed to *houdini* her way out of her cell, then Alf will be fired.
- (6’) Pat did not *whisky* the punch, though he did add some vodka.
- (8) At every physics department, there is someone who *einsteined* their way through grade school.

In this way, lexical innovations compositionally interact with the other expressions in the sentences they occur. This fact strongly suggests that lexical innovations are not the result of conversational implicature, for it is widely agreed that *conversational* implicatures—in contrast to other not-at-issue contents such as conventional implicatures or lexical presuppositions—do not exhibit such compositional interactions.²⁰

²⁰ C.f. what Robyn Carston calls the *Scope Principle*: if an element of meaning is shown to fall under the scope of propositional connectives, it cannot be part of an implicature but has to contribute to the truth-conditional content of the utterance; see Carston (2002a, pp. 191–197) for discussion.

3.3 Moving forward

If lexical innovations cannot be explained away by appealing to conversational implicature, then it seems to follow that conventionalism is mistaken. This is precisely what Davidson took cases of lexical innovation to show, and he consequently sought to develop an alternative conception that affords no foundational role to conventions in an account of linguistic meaning, and which “give[s] up the attempt to illuminate how we communicate by appeal to convention.”²¹

On Davidson’s alternative account, each conversational participant is equipped with a “prior” theory that specifies her private grammar, as well as general knowledge concerning the ways in which people tend to get their points across. In response to a speaker’s utterance, audience members develop a “passing” theory concerning how a speaker intended her words to be understood, making revisions to their prior theory as necessary. According to this account, successful linguistic communication requires the passing theories of both speaker and audience to assign the same meanings to the expressions occurring in the uttered sentence.

As Davidson insists, linguistic conventions need not play any deep explanatory roles within this framework.²² For the account he provides does not require the prior theories of the speaker and the audience to be shared before the interpretation of an utterance takes place; moreover, the passing theories that the speaker and the audience construct need not be expected to be utilized in any future communicative exchanges. In this sense, Davidson maintains that linguistic meanings are grounded in *momentary* agreements between speakers and their audiences, but nothing more.

However, Davidson’s rejection of any substantive role for semantic conventions comes at a serious cost: for the process by which lexical innovations are generated and understood is far more constrained than his theory predicts. Consider, for example, the following utterances of (9) and (10):

(9) #A student generously offered *to car* me home after the talk.

(10) #Oscar *glassed* the wine before serving dinner.

Even if audience members can figure out what a speaker would be trying to express in uttering (9) and (10), the denominal verbs occurring in these sentences are considerably more marked than the other cases involving denominal verbs discussed above. But there is nothing in Davidson’s framework that explains the contrast between the cases. In the next section, I will argue that the existence of background semantic conventions helps explain the unacceptability of these cases.

Furthermore, Davidson’s account implies that hearers can—and often do—correctly identify speakers’ intended meanings without relying on background semantic conventions at all.²³ But in each of the cases of lexical innovation that we

²¹ Davidson (1986, p. 107).

²² Davidson (1984, p. 16).

²³ Thanks to an anonymous referee at *Linguistics and Philosophy* for helping me to clarify the problem at issue here.

have considered, hearers seem to rely on prior linguistic conventions in order to identify the meanings associated with the novel word use. Again, contrast the cases of lexical innovation that we have considered with an utterance of (11), adapted from Lewis Carroll:

(11) All mimsy were the borogoves, and the momes rath outgrabe.

It seems preposterous to suggest—as Davidson (1986) himself does—that audience members have no trouble identifying the meaning of (11) on first encounter with it.²⁴ While the syntactic structure of (11) does provide the hearer with clues as to the intended meaning of the words, hearers have to resolve too many open variables, with far too few prior semantic precedents, to successfully recover the meaning of (11). What each of these problems shows is that Davidson’s full-stop rejection of linguistic conventions is unwarranted. By rejecting any explanatory role for semantic conventions, Davidson’s account misses important empirical generalizations about the phenomenon of lexical innovation.²⁵

I believe that cases of lexical innovation do provide a decisive challenge to accounts which maintain that linguistic conventions are relatively stable bodies of facts whose properties are determined independently of the particular agents engaged in a communicative exchange and prior to the interpretation of the utterance. But semantic conventions need not be understood in this way. Specifically, the mere fact that an expression-meaning pair was not established by convention prior to a speaker’s utterance should not, by itself, be taken as evidence that the expression-meaning pair is not conventional. For “a convention is so-called because of the way it persists, not because of the way it originated.”²⁶

In what follows, I develop an account of semantic conventions that incorporates these points. In particular, I will develop a *dynamic* account of conventions according to which lexical innovation can serve to update background conventions in such a way as to introduce new word meanings into the language or shift the meanings of words already in circulation. As I show, these semantic changes are grounded in the commitments and expectations of interacting agents, and, hence, are perfectly consistent with the nature of social conventions as Lewis understands them. Thus rather than motivating a rejection of conventionalism, cases of lexical innovation illustrate the power and flexibility of that framework.

²⁴ Reimer (2004) offers a related criticism concerning Davidson’s take on sentences like (11).

²⁵ In Armstrong (2016) I argue that while Davidson could resolve these problems by adding additional constraints on the process in which passing theories are constructed, the end result is a notational variant of conventionalism.

²⁶ Lewis, (1975 p. 181).

4 A dynamic account of semantic conventions

It is uncontroversial that semantic conventions change: new semantic conventions are introduced and existing semantic conventions are revised. Likewise, no one denies that semantic conventions vary among groups of language users: different groups of speakers utilize different lexicons and employ those lexicons for different practical purposes. Although obvious and uncontested, these claims have not played an important role in philosophical discussions of the nature of linguistic meaning and communication.²⁷

This lack of attention to the topics of semantic change and variation may be the product of a natural desire to abstract away from various complications for the purposes of theory building. But I suspect that the lack of attention is more often the result of substantive assumptions about the time scales in which semantic conventions change, and about the size of the relevant population within which semantic conventions prevail. In particular, I believe that it is widely assumed that semantic conventions are patterns of language use displayed a relatively large number of times among the members of a relatively large group of speakers.²⁸

Lewis' account of convention may seem to reinforce these assumptions. As we've seen, Lewis informally describes semantic conventions as regularities in language use that prevail among the members of a group of agents. In the slightly more technical formulation, semantic conventions are characterized as solutions to recurrent coordination problems that perpetuate among a group of agents because of a common interest in communication. This talk of "regularities" and "recurrent" problems has overwhelmingly been interpreted in *statistical* terms; that is, pairings of linguistic forms and meanings that reoccur among many speakers over many occasions of use.²⁹ As such, semantic conventions will be largely static across individual conversations and independent of the particular speakers and hearers involved.

This statistical construal of semantic conventions requires providing answers to a number of difficult questions. For example, we need to specify exactly how many times an expression-meaning pair must be used in order to be conventional, and exactly how many speakers must participate in that pattern of usage in order for a convention to be established. The problem is not merely that there appear to be no determinate answers to these questions, creating a vague boundary between the conventional and the non-conventional. Rather, the problem is that resolving these

²⁷ Although Evans (1973), Mercier (1994), and Tappenden (1999) provide notable exceptions.

²⁸ Rarely given any argument, this assumption is often taken for granted in many current discussions in the philosophy of language and semantics; see, however, Dummett (1978, 1993) and Wiggins (1997) for explicit statements and defenses of this "static" conception of linguistic conventions.

²⁹ It is not clear to me whether or not Lewis himself accepted this statistical construal of conventions. While his informal glosses suggest that he did, some of the paradigmatic instances of conventions that he cites do not. Indeed, in his discussion of Hume's boat rowers, he explicitly allows for conventions that hold "in a very small population for a very short time—between two people for a few minutes...." (1969, p. 44). Be that as it may, I will leave an exegesis of Lewis for another occasion.

statistical questions does not illuminate foundational issues concerning the nature of linguistic meaning or the process of linguistic communication.³⁰

Fortunately, the statistical construal is not required by Lewis' account of convention. The core element of Lewis' account is that the coordinated attitudes or actions of a group of agents give rise to various default commitments and expectations for future episodes of coordination among the members of that group. More exactly, a group of agents have established a convention if, having found an arbitrary solution to a coordination problem, the members of the group are committed (and expect that the other members of the group are likewise committed) to reverting to that solution should that problem re-occur. Crucially, this does not require that a pattern of attitudes or actions be displayed a relatively large number of times among a relatively large number of agents in order for that pattern to constitute a convention. Indeed, Lewis' account allows that a convention can be established after a single episode of coordination.

These points open up the possibility of a *dynamic* conception according to which semantic conventions can change within the course of a communicative exchange and across local groups of agents within a wider population. In what follows, I will develop each of these points in turn.

4.1 Adaptation

As before, I will assume that the members of a conversation have shared knowledge of a conventionally established semantic theory for a language. In particular I assume that the convention to use one language rather than another is sustained by a common interest in efficient communication among the members of that group. In this sense, it is because you and I wish to communicate efficiently that you and I use our words in accordance with our previously established linguistic conventions.

However, I maintain that the members of the group can adapt background semantic conventions rapidly in the course of a single communicative exchange. A novel semantic convention has been introduced among a group of agents if the members of that group coordinate on the form and meaning of a word, as it occurs in an uttered sentence, and are committed (and expect that the other members of the group are likewise committed) to using that form with that meaning in any future communicative exchanges. When these conditions obtain, the speakers' utterance serves to establish a new *lexical precedent* for the purposes of the conversation.

Lexical innovations can be fruitfully modeled in these terms. Consider, for example, the attendant's use of 'koba' in (1). Although this expression was not part of the communal lexicon shared by you and the attendant before your communicative exchange, the attendant's utterance served to update your communal lexicon with a new expression-meaning pair. Similarly, if audience members coordinate on the meaning of the novel verb 'to houdini,' as it occurs in (3), and form the relevant future directed expectations, then the lexicon shared by the members of the

³⁰ These points have been emphasized by Chomsky (1993, p. 17). But in contrast to Chomsky, I do not deny that there is theoretical interest in studying statistical regularities in language use. My claim is simply that a high degree of statistical regularity is not constitutive of linguistic conventions in general, or of semantic conventions in particular.

conversation is enriched with a new expression-meaning pair within the class of verbs.³¹

In these cases, we can say that the innovative use of an expression has resulted in a *lexical expansion* of the language.³² A language L' is a lexical expansion of language L iff L' is just like L , except that the lexicon of L' contains at least one additional linguistic expression not contained in the lexicon of L . More generally, lexical expansions involve cases in which the set of conventional meanings available for use between speaker and interlocutor has been enriched; the members of the conversation have, in other words, introduced a new character into the language.

Crucially, members of a conversation do not add new lexical meanings without reason. In many cases, the introduction of a new lexical meaning is prompted by the fact that there is no single expression already in the lexicon that has the meaning the speaker wants to convey. Moreover, agents generally attempt to respect what Clark (1993) has called the *principle of contrast*: all things being equal, differences in form should reflect differences in meaning. As Clark points out, the principle of contrast helps explain why the lexical innovations in (6) and (7) are marked:

- (6) #A student generously offered to *car* me home after the talk.
- (7) #Oscar *glassed* the wine before serving dinner.

For it seems that the occurrences of 'car' and 'glassed' are intended to express the properties of driving and pouring. But insofar as there are already conventionally established ways of expressing those meanings, the speaker who uttered sentences (5) or (6) would be violating the principle of contrast.³³ In cases such as these, prior conventions preempt the need for introducing new conventions.

The process of lexical expansion illustrates one of the ways in which background semantic conventions can change in the course of a communicative exchange. But background semantic conventions can change in other ways as well. In particular, members of a conversation may also adjust background conventions in order to revise the meanings of expressions already in the lexicon. Recall the case involving the speaker who has appropriated her lexical choices to match that of her

³¹ Some may worry that this treatment of denominal verbs runs counter to the "Modified Occam's Razor Principle" employed by Grice (1975), according to which linguistic meanings should not be modified beyond necessity. In reply, it should be noted that in Sect. 2 we saw a number of reasons for holding that denominal verbs have linguistic meanings, and hence all things are not equal in this case. Furthermore, there are many uncontroversial cases of denominal verbs which have conventionally established meanings such as 'lynch' and 'boycott.' If the Occam's razor principle applied to denominal uses of 'houdini,' it is not clear why the principle would not also apply to every denominal verb. See Devitt (2004) for more elaboration of these points and further discussion.

³² I am here adapting the terminology of Dowty (1979, pp. 298–300).

³³ So the claim here is not that the verbs 'to car' and 'glassed' are unaccepted in every possible discourse situation. There may well be a semantic convention according to which, for example, 'to glass' a bottle of wine means to break that bottle in preparation for a fight. What I am claiming is that the verbs in (6) and (7) are preempted or blocked from establishing new conventions when those expressions are intended to express word meanings already established by convention among those engaged in the conversation.

idiosyncratic audience member and used the expression ‘erudite’ to denote the property of being erratic.

In cases of this sort, we witness not a lexical expansion but a *lexical shift*. A language L' is a lexical shift of language L iff L' is just like L , except that the meaning of some linguistic expression in the lexicon of L' is different from the meaning of that same linguistic expression in the lexicon of L .³⁴ To put it a bit differently, prior linguistic conventions have been revised in a way that changes the semantic properties of some expression already in circulation among the members of a group.

In some cases of lexical shifts, like the one above, the new expression meaning is completely disjoined from the earlier expression meaning. But lexical shifts more commonly involve broadening or narrowing the extensions associated with the meaning of an expression. For example, a group of agents may come to associate a meaning with ‘bachelor’ that excludes some unmarried men from the extension of the expression and only correctly applies to unmarried men *looking for partners*. Similarly, a group of agents may come to associate a meaning with ‘martini’ that not only includes cocktails made with *gin* and *vermouth* but also cocktails made with *vodka* and *vermouth*.³⁵

Lexical expansions and lexical shifts are two standard ways in which groups of agents can dynamically adapt background semantic conventions because of a shared interest in linguistic communication. But it is worth emphasizing that a dynamic account of this sort does not entail that background semantic conventions change anytime an expression is used with an innovative meaning. After all, not all lexical innovations will result in coordination between speakers and their audiences; for example, although innovative, audience members may not be able to coordinate on the meanings of expressions such as ‘borogoves’ or ‘outgrabe’ as they occur in *Jabberwocky*. Moreover, not all cases of coordination on the form-meaning pair of an innovative word meaning will result in the relevant future-directed commitments and expectations among the members of the conversation. For example, an audience member may be able to identify the meaning of ‘erudite’ as it occurs in an utterance of (7):

(7) You should avoid hiring Joe the plumber, because he is very erudite.

but respond with “Joe is not erudite, but he is erratic” or with “hey, wait a minute, ‘erudite’ does not mean erratic!”

³⁴ This, again, roughly follows the terminology of Dowty (1979, pp. 298–300). The definition makes the controversial assumption that linguistic expressions do not have their meanings essentially, and that they can instead be individuated in terms of their phonological, morphological, and syntactic properties. We can remain neutral concerning this controversial assumption by reformulating the definition of lexical shifts in terms of the notion of a *lexical counterpart*. A disambiguated expression e^* in language L' is a lexical counterpart of disambiguated expression e in language L iff e^* and e have the same phonological, morphological, and syntactic properties. We can then say that L' is a lexical shift of L iff (i) each expression in the lexicon of L' has a lexical counterpart in the lexicon of L , and (ii) there is an expression e^* in lexicon of L' and an expression e in the lexicon of L , such that e^* and e are lexical counterparts and e^* differs in meaning from e . Thanks to Carlotta Pavese for very helpful discussion of these points.

³⁵ See Carston (2002a) and Ludlow (2014) for more discussion of these kinds of cases.

In this sense, the semantic changes in question require a kind of accommodation on the part of audience members. If audience members do not accommodate the lexical innovation—either because they cannot identify the meaning of the innovation or because they can identify the meaning but reject it—then no update to the background conventions will occur. Changes to prior conventions are thus not cheap, and they are not an inevitable consequence of an innovative word use.³⁶

4.2 Locality

Semantic conventions are not static: they can be updated in the course of communicative exchange with the result that new expression-meaning pairs are established in the lexicon. In some cases, these newly established conventions may well spread among the members of a wide population of language users and begin to be recorded in dictionaries. But, as I've noted, widespread use is not required in order for a semantic convention to be established. For while semantic conventions require future-directed commitments and expectations for language use among those *within* a group, they involve no such requirements among those outside the group.

It is in this sense that I claim that semantic conventions are *local*: they involve assignments of meanings to linguistic expressions suited for use among a specific group of agents. Accordingly, the semantic conventions that prevail among a group of agents need not correspond to the semantic conventions of institutionalized languages such as “standard” English, French, or Armenian. For this reason, it is useful to adopt the terminology of Ludlow (2014) and call the languages determined by local semantic conventions *micro-languages*: for they are a set of expression-meaning pairs that are suited for the needs and experiences of subgroups within a wider population.³⁷

Semantic conventions can, in this sense, be fruitfully compared with the common ground of a conversation. Consider, for example, the context set of a conversation—roughly, a set of possibilities compatible with what has been mutually accepted by those engaged in a conversation. Different context sets will capture the presuppositions of different groups of interlocutors, and individual agents may move between conversations involving different context sets. For example, while I will presuppose that humans have had a profound effect on global warming in a conversation among a group of environmental activist, I may not make this presupposition in a conversation with a group of oil tycoons. Similarly, distinct semantic conventions may prevail among different groups within a wider population and individual speakers may utilize different semantic conventions when engaging in conversations with different groups of speakers. So while I will assume one meaning for

³⁶ Indeed, Ludlow (2014) has recently drawn attention to the ways in which members of a conversation can engage in explicit debates concerning the appropriate meaning of a word. As Ludlow emphasizes, there are a variety of norms that language users hold each other responsible to when engaging in such debates.

³⁷ See also Clark (1998).

‘valid’ in a conversation with a group of logicians, I may not assume that ‘valid’ has the same meaning in a conversation with a group history professors.³⁸

The comparison with the context set also highlights the manner in which boundaries of the various subgroups are to be drawn. According to the present account, the semantic conventions that prevail among a group of agents are grounded in the fact that there is a set of expression-meaning pairs that are in the individual lexicons of each agent in that group, and which each agent expects to be in the individual lexicons of the other members of that group.³⁹ In this way, just as differences in the context sets of different conversations are grounded in the mutual attitudes of groups of agents concerning what the world is like, differences in semantic conventions are grounded in the mutual attitudes of groups of agents concerning how to interpret one another’s words.⁴⁰

5 Semantic coordination *on the fly*

I have argued that Davidson’s argument involving cases of lexical innovation turns on an erroneous conflation of what is conventional with what is established prior to the point of utterance. Semantic conventions can be dynamically updated within a local group of agents in the course of a single communicative exchange, resulting in expansions or shifts to the standing meanings of linguistic expressions. The claim that semantic conventions are dynamic in this way is compatible with a number of different proposals concerning the features that *determine* the semantic content of a lexical innovation on an occasion of use and with a number of different proposals concerning the *mechanisms* that enable audience members to identify the semantic content of a lexical innovation. In this section, I will embed the dynamic account of semantic conventions I developed within a particular account of the features that determine the semantic content of a lexical innovation in context and of the mechanisms that enable coordination to be achieved “on the fly.”

³⁸ This is not to deny that there are important disanalogies between the common ground and linguistic conventions. In contrast to semantic conventions, the common ground of a conversation is *constantly* being updated as a conversation proceeds—indeed, it is the primary point of an utterance to change the prior state of the common ground. But while linguistic utterances *can* serve to update background conventions, *they need not*; an audience member can, for example, update the common ground with the fact that Smith is in Boston, without thereby altering the meaning of ‘Smith’ or ‘Boston.’

³⁹ This is obviously not intended to be a sufficient condition on membership within a linguistic community, since the members of a linguistic community should presumably also be aligned with respect to the grammatical principles they utilize for generating complex expressions and their meanings from the lexicon.

⁴⁰ It has recently been suggest by von Fintel and Gillies (2011) that, rather than determining a single context set, agents’ factual attitudes and expectations determine *a range of admissible context sets*. Similar remarks are, I believe, apt for describing agents’ linguistic attitudes and expectations: semantic conventions serve to fix *a range of admissible (micro)languages*. Since it is not directly germane to the topic of lexical innovation, I will set this additional complication aside for the purposes of this paper.

5.1 Determination revisited

In recent years, a growing number of theorists have maintained that the semantic contents of many linguistic expressions—the occasion meanings of those expressions—constitutively depend on speakers' *referential* intentions: denotations for expressions that speakers have in mind and intend audience members to recognize as the denotations of the expressions they use.⁴¹ The invocation of speakers' intentions seems particularly natural in characterizing the semantic contents of lexical innovations, on occasions of use. For it is difficult to resist the claim that the semantic content of a denominal verb such as 'to houdini,' is intimately connected with what a speaker is intending to express in using that verb.

There is, however, a problem. It is typically maintained that speakers' intentions are semantically significant only to the extent that they are compatible with the conventions governing the standing meanings of the expressions involved. But insofar as a lexical innovation is not associated with a standing meaning by prior convention within a linguistic community, it would appear that speakers' intentions do all the work in determining the semantic contribution of a lexical innovation on an occasion of use—in such cases, an innovative word means whatever a speaker intends it to mean. This is not a happy result, from the point of view of conventionalism. For it is at odds with the claim that the semantic properties of a speaker's words depend in any sense on background conventions, and it eradicates the distinction between speaker meaning and semantic meaning.

Yet, the conventionalist can accept that speakers' intentions are semantically significant without being forced to accept the claim that an innovative word use means whatever a speaker intends it to mean. For the conventionalist can maintain that the semantic content of a lexical innovation depends on a speaker's referential intentions together with facts about the speakers' audience. In particular, I propose that the conventionalist should accept COORDINATED DETERMINATION:

COORDINATED DETERMINATION: The occurrence of a lexical innovation e in a sentence S has semantic content Φ , just in case the speaker intends Φ to be the content of e as it occurs in S , and members of the speaker's audience, with knowledge of the common ground, identify Φ as the content of e as it occurs in S .⁴²

According to this "coordinated intentions" account, speakers' intentions do indeed play a constitutive role in fixing the occasion meaning of a lexical innovation; however, those intentions are held to be semantically significant only insofar as

⁴¹ The move toward incorporating speakers' intentions into the features that determine the semantic contents of expressions in context was largely initiated by work by Donnellan (1966) on descriptions and by Kaplan (1989b) on demonstratives.

⁴² I am here building off a recent proposal by Jeffrey King (2014a, b), though I do not appeal to "idealized" audience members in the way suggested by King. Although developed in a different theoretical framework, see Neale (2005) and, in particular, Heck (2014) for related proposals.

members of the speakers' audience recover the content the speaker intended to convey on the occasion of use.⁴³

In this sense, innovative word meaning is an *achievement*: speakers must secure actual uptake with respect to their referential intentions in order to generate a semantic content for the lexical innovation in context. It is on the basis of successfully achieving coordination, and thereby generating a semantic content in context, that speakers' innovative word uses serve to update background semantic conventions. Specifically, it is in virtue of the fact that the members of the conversation have coordinated on the speaker's referential intentions, and formed the relevant future-directed expectations described in the last section, that background semantic conventions will either be expanded or shifted. But if no uptake is secured—if the value the speaker intended doesn't match the value the audience recovers—then the innovation will merely have speaker meaning.⁴⁴

5.2 Recoverability

The coordinated intentions account grounds the meaning of a lexical innovation in the way audience members successfully coordinate on the speaker's intended meaning.

Needless to say, audience members cannot look inside a speaker's head and identify her intentions. So what are the mechanisms that enable audience members to successfully coordinate with the speaker on the intended meaning of her words?

At a superficial level, the answer to this question is not difficult to discern: speakers must make their intended meaning *sufficiently salient* in order for coordination to be achieved (Lewis 1969). When speakers' utterances do not contain any lexical innovations, salience takes the form of previously established precedents that serve to guide the process of coordination. When speakers' utterances do contain lexical innovations, salience takes the form of a meaning that “stands out,” or is sufficiently distinguished, to allow audience members to identify that meaning as the one that the speaker intended. But at a deeper level, the appeal to salience does not help to resolve the central problem; for we would like an illuminating account of what makes an intended meaning sufficiently salient to audience members—particularly in cases of lexical innovation, where audience members cannot rely on prior precedents concerning the meaning of the word.

⁴³ There may, of course, be various mixed-cases where some, but not all, of the audience members recover the speaker's intended meaning of the lexical innovation. In such cases, the present account predicts that the speaker will have expressed semantic contents with respect to some, but not all, of the audience members. While it is beyond the scope of the present paper to tease out the implications of this prediction, I believe it to be a welcome result; see Cappelen (2008) and von Stechow and Gillies (2011) for relevant discussion.

⁴⁴ The claim here is not about what gets called “semantics.” The conventionalist need not deny that there may well be a perfectly good sense in which one could call the speaker's intended meaning “the semantic content of the word in the speaker's idiolect.” But the conventionalist will insist that there are important theoretical roles for a notion of content that is grounded in the coordinated attitudes and actions of social groups, and that there is a worthwhile project of semantics that systematically investigates the objects that play these underlying theoretical roles.

The question of the underlying mechanisms at work in enabling audience members to recover a speaker's intended meaning is, of course, an empirical matter—one whose full resolution is well beyond the scope of this paper to provide.⁴⁵ Nevertheless, I will tentatively suggest that audience members integrate distinctively linguistic knowledge concerning the lexical category of the innovative word and distinctively social knowledge concerning background semantic conventions and the common ground of the conversation in order to recover the intended meaning of a word.⁴⁶ Together, these language specific and language general constraints provide a kind of scaffolding on which coordination can be achieved “on the fly.”

According to the model I propose, audience members attempt to identify the intended content of a lexical innovation by first *identifying the lexical category of the innovation* as it occurs in the uttered sentence. For this will serve to constrain the search-space among the class of candidate meanings to those meanings that are appropriate for the category of the expression in question. Once the lexical category of the expression has been identified—for example, once the audience members identify whether the expression ‘whisky,’ as it occurs in an uttered sentence, is being used as a noun or a verb—the audience members *look to the common ground for an entity appropriate to an expression of that lexical category*.

As before, I will assume that the common ground of a conversation is a structured body of information that the discourse participants take to be available in interpreting one another's utterances. In particular, I will take the common ground to be a structure, consisting of a *context set*, a set of *questions under discussion*, and an *attention state*.⁴⁷ The context set of a conversation is the set of possible worlds compatible with everything that the discourse participants have mutually accepted as true; the questions under discussion divide these possible worlds into various subject matters or issues, reflecting a partially-ordered list of questions that the discourse participants are interested in answering; finally, the attention state specifies a class of entities in a shared perceptual space shared among the discourse participants.

In some cases of lexical innovation—particularly when the expressions at issue are nouns—the intended contents are made directly available by the attention state of the common ground. Consider, for example, the case of the postal worker's utterance of (1)

(1) Let me measure the volume of your package with a koba.

Your ability to recover the content of ‘koba’ is largely made possible by the fact the expression was used as a count noun and the fact that there is an unfamiliar type of

⁴⁵ See, Garrod and Anderson (1987), Garrod and Doherty (1994), Brennan and Clark (1996), and Brennan (1996) for relevant discussion.

⁴⁶ The model has been developed in recent work on lexical acquisition; see, for example, Gleitman et al. (2005), Trueswell et al. (2013).

⁴⁷ This roughly follows the proposal in Roberts (1996, 2004), Ginzburg (1995, 2012), Ginzburg and Cooper (2004).

ruler in the attentional space shared by you and the attendant. This is perhaps even more vivid in the scenario in which an utterance of (2) was made:

(2) Mopsy loves to greet strangers with a lick.

and there is an unfamiliar dog in front of your face. Again, recovery here turns on the fact that ‘Mopsy’ is being used as a proper name and that there is an individual in the attentional space shared by you and the speaker that is not already associated with a proper name.⁴⁸

But in many other cases, recoverability turns on the interaction of commonly presumed factual information together with the questions under discussion. Suppose that Pat and Sue are in Paris for a trip with their history class and are discussing what each of them did that day; while sharing various memorable encounters while taking picture of the city, Pat utters (12):

(12) The tour guide laughed when I did a napoleon for the camera.

Once Sue identifies that ‘napoleon’ is being used as a verb, she will look for information in the common ground that is mutually presupposed about the man Napoleon. Given the current questions under discussion, and given what Pat and Sue both know about Napoleon, Sue may identify that ‘napoleon,’ as it occurs in (12), was used to denote the act of putting one’s hand inside one’s shirt at the stomach. However, suppose that Pat and Sue had not been talking about what they did around Paris that day but instead had been talking about what they each learned about military history. Suppose now that in this alternative discourse context Pat utters (13):

(13) I was surprised to learn that Hitler did a napoleon in Russia in 1941.

Again, once she identifies that ‘napoleon’ is being used as a verb, Sue will look for information in the common ground that is mutually presuppose about the man Napoleon. But given the questions that are under discussion, Sue may identify that ‘napoleon,’ as it occurs in (13), was used to denote the act of marching one’s land troops through Russia in the middle of winter.

The process by which semantic coordination is achieved can also be shaped by interactions between specific features of the common ground and general expectations concerning the cooperative norms of conversation, discussed by Grice (1975) and further elaborated by Craige Roberts (1996) and (ms).⁴⁹ Consider a

⁴⁸ Again, the fact that the speaker is not violating the principle of contrast is important. If, for example, it is common knowledge that the dog currently licking your face is named ‘Fido’ it would be natural for you to take ‘Mopsy,’ as it occurs in (2), to refer to some other dog or else to take the speaker to have been confused about the discourse situation.

⁴⁹ Following Roberts, we might also characterize these norms in terms of the following default expectations that audience members bring to the communicative exchange: audience members expect that a speaker’s assertion is *informative*, in that the content of that assertion should eliminate some possible worlds from the context set; that a speaker’s utterance is *relevant*, in that the content of that utterance serves to answer one of the questions under discussion; that a speaker’s utterance is expected to be

conversation among a pair of housemates about what drinks to serve the guests at their dinner party, given that they forgot to pick up wine and only currently have vodka and vermouth in the house. If one of the housemates subsequently utters (14):

(14) It's no problem; I'll make a martini for anyone that wants one.

It would be natural for the other housemate to take the speaker to have intended a meaning for 'martini' that includes drinks made with vodka and vermouth. For otherwise the speaker's utterance would not answer the current question under discussion, and hence be uncooperative given the goals of the conversation.⁵⁰

In short, audience members utilize information about the lexical category of an expression together with information provided by both background conventions and the common ground in order to bootstrap their way onto the speakers' referential intentions concerning lexical innovations. Of course, the process of recovery sometimes fails: audience members are not always able to identify either the form or the intended content of a speaker's utterance. A speaker may not have made it possible for the audience to recover her intended meaning, or audience members may have only partially heard what she uttered or were not paying sufficient attention. When coordination fails, the discursive techniques of clarification and feedback play a crucial role. An audience member may say, "What did you say?", "What do you mean by the word 'e'?", or even simply "'e'?" In so doing, they put a new question under discussion.⁵¹ Given the goals of conversation, a speaker must respond to such clarificatory questions either by elaborating on what she intended to express or by using an altogether different expression.

5.3 Linking revisited

Generalizing on the discussion in the last two subsections, we can say that the semantic content of a sentence *S* in context *C* is determined on the basis of the semantic contents in *C* of the basic expressions occurring in *S*, and their manner of semantic combination in accordance with the syntactic structure of *S*. In some cases, the semantic contents in *C* of the basic expressions occurring in *S* will be secured by prior episodes of coordination among the members of a linguistic community. But in other cases, like those involving lexical innovations, the semantic contents of the

Footnote 49 continued

orderly, in that the utterance is formed grammatically and in a way that satisfies principles like contrast; finally, a speaker's utterance is expected to be *accurate*, in that the content of that utterance is taken to hold in the circumstances under consideration.

⁵⁰ As Neale (1992) and Bach (2004) have pointed out, there is little reason to believe that Grice took the cooperative norms to only govern the process by which speakers' conversationally implicate information to their audiences. More generally, the fact that cooperative norms are utilized by audience members to coordinate on the speaker's intended meaning is not itself reason to believe that the product of the coordination is not semantic. See Sect. 5.1 for more discussion of this point.

⁵¹ See Clark (1996), and, in particular, Ginzburg and Cooper (2004), for more discussion.

basic expressions occurring in S in C will be secured by the occurrent coordination between speakers and audience members.⁵²

An account of this kind straightforwardly preserves the LINKING thesis; for linguistic communication does seem to require that speakers and their audiences coordinate on the semantic contents of the basic expressions occurring in an utterance.

In particular, while successful linguistic communication does not require speakers and their audiences to have shared knowledge of the semantic properties of the expressions occurring in a sentence *before* the time of utterance, it does require them to be able to coordinate on the semantic properties of the expressions occurring in an uttered sentence *as a dialogue unfolds*.

6 Consequences

The phenomenon of lexical innovation does not undermine an explanatory role for semantic conventions in an account of the meanings of the expressions in a language or in an account of the process of linguistic communication. Indeed, I have argued that semantic conventions play an essential role in elucidating the manner in which the members of a conversation coordinate on the meanings of a lexical innovation and in characterizing the significant constraints on this process. In this final section, I will consider some of the implications that the present account has for standard ways of understanding the semantics-pragmatics interface, the distinction between speaker reference and semantic reference, and the relationship between literal and non-literal meaning.

6.1 The autonomy of semantics

The dynamic account I have developed assigns a central role to the shared attitudes and expectations of language users in determining the semantic contents of sentences, relative to the contexts in which they are used. For some, the claim that semantics depends on language users' intentions and expectations is something of a contradiction in terms. For a wide variety of theorists have maintained that a semantic theory for a language operates completely autonomously from speakers' intentions and expectations and from the way that language is employed to achieve interpersonal coordination. Indeed, the claim that semantics is autonomous in this sense has been commonly accepted by both those who have sought to defend explanatory roles for semantics and those who have downplayed explanatory roles for semantics.⁵³

⁵² In this respect, I believe that the semantic contents of lexical innovations have much in common with the semantic contents of expression that depend on context as a matter of their lexical meanings. Again, see King (2014a, b) and Heck (2014) for discussions of the meta-semantics of context-sensitive expressions.

⁵³ See Borg (2004, 2012) for a version of the former, and Sperber and Wilson (1986/1995), Carston (2002a, b, 2008), Recanati (2004, 2010), and, with some important caveats, Bach (2004), for developments of the latter.

I believe that the claim that semantics is autonomous from speakers' intentions and expectations has received spurious support from failing to distinguish between two senses in which language, construed as a distinctive cognitive capacity, is autonomous. Let's distinguish between the cognitive faculty of language in a *narrow sense* and the cognitive faculty of language in a *broad sense*. The faculty of language in a narrow sense will pertain to those cognitive capacities that concern combinatorial principles unique to language, whether those combinatorial principles reside exclusively in syntax (as suggested by Hauser et al. 2002) or in other areas of language as well (as suggested by Pinker and Jackendoff 2005). The faculty of language in the broad sense will concern all those cognitive capacities that are utilized in the production and processing of language, including those which are distinctively linguistic or combinatorial. While it is debated whether or not the faculty of language in the narrow sense is autonomous from speakers' intentions and expectations and from the way that language is employed to achieve interpersonal coordination, it is definitional that the faculty of language in the broad sense is not autonomous in this way.

With respect to semantics, it may be plausible to suggest the principles of semantic composition are autonomous—for they may well be part of the language faculty in the narrow sense. But it is highly implausible to suggest that the principles that serve to associate semantic properties with the basic expressions in the lexicon are autonomous—for they must interface with cognitive capacities outside of the faculty of language in the narrow sense, particularly with an agent's general theory of mind. For this reason, it is not surprising that explaining agents' abilities to generate and understand lexical innovations has turned critically on the ability to anticipate and recognize speakers' intentions. More generally, the fact that a process turns on speakers' intentions and their recognition does not by itself entail that the process is not semantic.

How then, if at all, should the line between semantics and pragmatics be drawn? Semantics, as I have understood it here, is dictated by linguistic structure in a way that pragmatics is not. Speakers—both individual and collectively—are constrained by the structural features of the language they use when they associate semantic contents with linguistic expressions. At the lexical level, the semantic content of an expression must be licensed by the phonological, morphological and grammatical properties of that expression in addition to background features of the common ground; likewise the semantic contents of complex phrases and sentences will, in general, be wholly determined as a function of the semantic contents of their parts and their manner of semantic combination as licensed by the syntactic structure of those complex expressions. Conversely, the process of semantic interpretation involves audience members seeking to identify the content expressed by a linguistic utterance on the basis of the expressions occurring in the uttered sentence and their manner of combination relative to the context of utterance. As we've seen, audience members may well use information about the common ground and about the cooperative norms of conversation in the process of semantic interpretation; but this general-purpose reasoning is in the service of identifying the intended meanings of an utterance in a structurally sensitive manner—specifically, in a way that is sensitive to the phonological and morphological forms of the expressions occurring

in an uttered sentence, and how they are to be combined together to give the meaning of the sentence given the particular syntactic structure of that sentence.⁵⁴

In contrast, there is no similarly direct link between meaning and linguistic structures within pragmatics. In the case of pragmatics, speakers are intending to convey information over and above, or even in spite of, the semantic contents of their utterance. While speakers do utilize the structural features of the sentences they utter to do this, the messages they intend to pragmatically convey will not, in general, be wholly determined as a function of the semantic contents of the sentences parts and their manner of semantic combination as licensed by their syntactic structure. What a speaker intends to pragmatically convey *by* an utterance need not, in this sense, correspond to any structural element *in* that utterance. Conversely, the process of pragmatic interpretation involves audience members seeking to identify not the intended meanings of the expressions occurring in a sentence or how they are to be combined together to give the meaning of the sentence as whole, but rather additional collateral information that the speaker meant. The principles that audience members use in this process are not specific to any particular language but are held to follow wholly from general principles of rationality and from general principles of cooperative action.⁵⁵

This division of labor between semantics and pragmatics is, I believe, very much in the spirit of Grice's original proposal. For Grice (1957, 1975), what distinguishes "what is said" from what is "merely implicated" is not the fact that the former depends on speakers' intentions and the latter does not; rather, the two ways of communicating information are distinguished by the kinds of effects speakers intend their utterances to have and the manner in which those effects are intended to be realized.

6.2 Speaker reference and semantic reference

Consider the following case from Kripke (1977): Al and Bo are having a conversation, and in the distance they see a man working in his yard. Al asks, "What is Jones doing over there?" Bo answers: "Jones is raking the leaves." But Al and Bo have made a mistake; the man they were looking at was Smith, not Jones. Furthermore, Al and Bo both know that Smith is called 'Smith' and Jones is called 'Jones.' Kripke proposes the following diagnosis: while Al and Bo have both *semantically* expressed something false about Jones with their uses of 'Jones,' they have *speaker-meant* something true about Smith.

It may seem that Kripke's diagnosis is incompatible with the dynamic account of conventional meaning I have developed. After all, Al and Bo coordinated on Smith

⁵⁴ In broad outline, this is very much in the spirit of the account of semantics developed in King and Stanley (2005).

⁵⁵ C.f. Kripke (1977, p. 263): "The notion of what words can mean, in the language, is semantical: it is given by the conventions of our language. What they mean, on a given occasion, is determined, on a given occasion, by these conventions, together with the intentions of the speaker and various contextual features. Finally what the speaker meant, on a given occasion, in saying certain words, derives from various further special intentions of the speaker, together with various general principles, applicable to all human languages regardless of their special conventions."

as the value for the use of the name ‘Jones’ (i.e., they both took Smith to be the referent of ‘Jones’).

It may appear, then, that the present account would predict that this coordination would result in Al and Bo semantically expressing information about Smith with their uses of ‘Jones,’ thus controverting the semantic reference/speaker reference distinction.

However, the dynamic account that I have developed does not predict that Al and Bo have semantically expressed something about Smith. As the case is described, Al and Bo are mutually committed to referring to Smith with their use of ‘Smith’ and referring to Jones with their use of ‘Jones’ and they each intend to refer to Jones, on the occasion of their utterances, in accordance with these commitments. But Al and Bo unwittingly failed to speak in accordance with their own commitments. In particular, due to their confusion, Al and Bo spoke with *conflicting intentions*: they both intended to refer to the man they see in the distance and they also intended to do so in accordance with their background commitments concerning the name ‘Jones.’ It is in virtue of the fact that Al and Bo have failed to speak in accordance with their own commitments, and are instead guided by conflicting intentions, that Smith is not the semantic content of their uses of ‘Jones.’ If we follow Kripke’s diagnosis of the case, this is exactly as it should be.

In this respect, Kripke’s case is quite disanalogous with the cases of lexical innovation that we have been considering. In cases of lexical innovation, speakers are not violating their own semantic commitments and they are not guided by conflicting intentions—their innovative word uses are either not governed by prior semantic conventions, or they are intentionally using words in a way that departs from those conventions. When coordination is achieved in cases under these conditions, it is plausible that semantic reference has been secured. Consider, for example, a variant of Kripke’s case in which all other details are held fixed except the fact that Al and Bo had prearranged to use ‘Jones’ to refer to Smith for the purposes of their conversation. In this situation, I submit it is plausible to maintain that Bo’s utterance of “Jones is ranking the leaves” semantically expresses a proposition about Smith. Unlike this variant of Kripke’s case, the semantic contents of lexical innovations are not established by prearranged agreement but, rather, are established on the fly. But like this variant of Kripke’s case, cases of lexical innovation do not involve speakers’ violating their own background commitments or speaking with conflicting intentions.

6.3 Metaphor and the limits of conventions

It is natural to wonder about the extent to which a dynamic account of convention blurs the boundary between literal and non-literal uses of words. This is particularly the case with metaphor. After all, metaphorical uses of language seem to have many of the same properties that I have attributed to cases of lexical innovation: they would appear to involve innovative ways of pairing meanings with linguistic expressions, and those meanings seem to be able to interact compositionally with the meanings of other expressions in the sentences in which they occur. It therefore appears that if lexical innovations should be analyzed in semantic terms, so too

should metaphor. But some may think that this is surely incorrect, and that metaphor is, rather, a paradigm of the pragmatic.

In response, I deny that all metaphorical interpretation should be treated using the same theoretical tools. I propose that we distinguish between at least two kinds of metaphor, namely “shallow” and “deep.” We will call a metaphor shallow if the metaphorical interpretation of the sentence can be localized to a single expression occurring in the sentence uttered. For example, (15) and (16) are examples of shallow metaphors:

(15) Pat is a bulldozer.

(16) Many Americans are hawks.

In contrast, we will call a metaphor deep if its interpretation cannot be localized to a single expression occurring in the sentence. For example (17) and (18) would be examples of deep metaphors:

(17) I am a little world made cunningly.

(18) We are drifting back and forth/between each other like a tree breathing through its spectacles.

The metaphorical interpretation of both (16) and (17) is due to a complicated interaction of each of the expressions occurring within these sentences. While the dynamic account of linguistic conventions I have developed could be straightforwardly applied in the case of shallow metaphors, it does not obviously apply in the case of deep metaphors. When it comes to shallow metaphors, the speaker is using a language in which the expressions ‘bulldozer’ and ‘hawk’ are assigned meanings that depart from the meaning of those expressions in the language established by prior convention; here, we essentially have an instance of a lexical shift. In cases of deep metaphor, however, there is no principled way to shift the language that the speaker is using in order to arrive at the relevant metaphorical interpretation.

The fact that the account I have developed readily applies to shallow metaphors but does not apply to deep metaphors is a welcome result. First, shallow metaphors are part of the regular process of language change: they rapidly “die” or dwindle into highly conventionalized forms, including colloquialism or cliché. Second, shallow metaphors are often used to communicate information that directly contributes to the practical deliberations of those engaged in a conversation. For example, in a debate over who should be nominated for the chair of the Parent–Teacher Association, a speaker may utter (15) with the intention to convey that Pat would make an excellent chair because she removes obstacles in her path.⁵⁶

Neither of these features generally hold true for deep metaphors. Deep metaphors do not rapidly solidify into conventional meanings, and, when they do, they are preserved as more or less whole phrases. Moreover, deep metaphors are often open-ended in ways ill-suited for communicating information that directly contributes to the practical deliberations of those engaged in a conversation. In cases of deep

⁵⁶ This example is adapted from Bezuidenhout (2001).

metaphor, there does not seem to be a single content that a speaker intends to express, or which she expects audience members to recognize her as intending to express. This feature of deep metaphors seems to be, at least in part, what gives them their power, whether that power is affective or aesthetic: there is a varied and open-ended number of contents a speaker could be expressing, and audience members must imaginatively consider which (if any) were intended to be expressed.⁵⁷ It seems to me a mistake to try to explain the vast potential of deep metaphor using a semantic theory, and so I happily acknowledge that such cases fall outside the purview of the framework I have developed.

7 Conclusion

I have shown that cases of lexical innovation do not undermine an essential role for linguistic conventions in explaining the process of linguistic communication. Against Davidson, I have argued that a commitment to linguistic conventions helps explain why certain cases of innovation fail and how audience members identify the meaning a speaker intended to semantically express with her use of a word. Linguistic conventions arise as *local* solutions to coordination problems, and it is these local conventions that do the work in elucidating the process of successful linguistic communication.

Acknowledgements I am grateful for the feedback I received on this paper from Endre Begby, Liz Camp, Adam Crager, Sam Cumming, Daniela Dover, Katie Elliot, Lila Gleitman, Peter Godfrey-Smith, Gabe Greenberg, Daniel Harris, Jeffrey King, Ernie Lepore, Peter Ludlow, Eliot Michaelson, Sarah Murray, Anahid Nersessian, Carlotta Pavese, Jessica Rett, Michael Schweiger, Jason Stanley, Will Starr, and two anonymous referees from *Linguistics and Philosophy*. I am also grateful to audiences at the Department of Philosophy at Rutgers University and the Institute of Research in Cognitive Sciences at University of Pennsylvania.

References

- Armstrong, J. (2016). Coordination, triangulation, and language use. *Inquiry*, 59(1), 80–112. doi:10.1080/0020174X.2015.1115270.
- Aronoff, M., & Fudeman, K. (2010). *What is morphology?* (2nd ed.). Hoboken: Wiley-Blackwell.
- Atlas, J. D. (1989). *Philosophy without ambiguity*. Oxford: Oxford University Press.
- Atlas, J. D. (2005). *Logic, meaning, and conversation: semantic underdeterminacy, implicature, and their interface*. Oxford: Oxford University Press.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9, 124–162.
- Bach, K. (2004). Minding the gap. In C. Bianchi (Ed.), *The semantics/pragmatics distinction* (pp. 27–43). Stanford: CSLI Publications, Stanford University.
- Bartsch, R. (1984). Norms, tolerance, lexical change, and context-dependent meaning. *Journal of Pragmatics*, 8, 367–393.
- Barwise, J., & Perry, J. (1983). *Situations and attitudes*. Cambridge: MIT Press.
- Bauer, L. (1983). *English word formation*. Cambridge: Cambridge University Press.
- Bauer, L. (2001). *Morphological productivity*. Cambridge: Cambridge University Press.
- Bezuidenhout, A. (2001). Metaphor and what is said: A defense of a direct expression view of metaphor. *Midwest Studies in Philosophy*, 25, 156–186.

⁵⁷ See Lepore and Stone (2010, 2014) for a development of this point.

- Bezuidenhout, A. (2002). Truth-conditional pragmatics. *Noûs*, 36(16), 105–134.
- Borg, E. (2004). *Minimal semantics*. Oxford: Oxford University Press.
- Borg, E. (2012). *Pursing meaning*. Oxford: Oxford University Press.
- Brennan, S. (1996). Lexical entrainment in spontaneous dialog. In *Proceedings of international symposium on spoken dialogue* (pp. 41–44). Philadelphia, PA: ISSD-96.
- Brennan, S., & Clark, H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Meaning, and Cognition*, 22(6), 1482–1493.
- Brinton, L., & Traugott, E. C. (2005). *Lexicalization and language change*. Cambridge: Cambridge University Press.
- Burge, T. (1975). On knowledge and convention. *Philosophical Review*, 84, 249–255.
- Camp, E. (2006). Contextualism, metaphor, and what is said. *Mind and Language*, 21(3), 280–309.
- Cappelen, H. (2008). The creative interpreter: content relativism and assertion. *Philosophical Perspectives*, 22, 23–46.
- Carston, R. (2002a). *Thoughts and utterances. The pragmatics of explicit communication*. Hoboken: Wiley-Blackwell.
- Carston, R. (2002b). Linguistic meaning, communicated meaning, and cognitive pragmatics. *Mind and Language*, 17(1&2), 127–148.
- Carston, R. (2008). Linguistic communication and the semantics/pragmatics distinction. *Synthese*, 165(3), 321–345.
- Chomsky, N. (1993). *Language and thought*. London: Moyer Bell.
- Clark, E. (1993). *The lexicon in acquisition*. Cambridge: Cambridge University Press.
- Clark, H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Clark, H. (1998). Communal lexicons. In K. Malmkjaer & J. Williams (Eds.), *Context in language learning and language understanding* (pp. 63–87). Cambridge: Cambridge University Press.
- Clark, E., & Clark, H. (1979). When nouns surface as verbs. *Language*, 55, 767–811.
- Crain, S., & Pietroski, P. (2012). The language faculty. In S. Stich, R. Samuels, & E. Margolis (Eds.), *The Oxford handbook of philosophy of cognitive science* (pp. 361–381). Oxford: Oxford University Press.
- Davidson, D. (1984). Convention and communication. *Synthese*, 59(1), 3–17.
- Davidson, D. (1986). A nice derangement of epitaphs. In R. Grandy & R. Warner (Eds.), *The philosophical grounds of rationality* (pp. 443–446). Oxford: Oxford University Press.
- Davidson, D. (1991). James Joyce and Humpty Dumpty. *Midwest Studies In Philosophy*, 16, 1–12.
- Davis, W. (2003). *Meaning, expression, and thought*. Cambridge: Cambridge University Press.
- Devitt, M. (2004). The case for referential descriptions. In A. Bezuidenhout & M. Reimer (Eds.), *Descriptions and beyond* (pp. 280–305). Oxford: Oxford University Press.
- Devitt, M. (2013). Is there a place for truth-conditional pragmatics? *Teorema*, XXXII, 85–102.
- Donnellan, K. (1966). Reference and definite descriptions. *Philosophical Review*, 75, 281–304.
- Dowty, D. (1979). *Word meaning and Montague grammar*. Boston, MA: D. Reidel Publishing Company.
- Dummett, M. (1978). The social character of meaning. In M. Dummett (Ed.), *Truth and other Enigmas* (pp. 420–430). London: Duckworth.
- Dummett, M. (1993). Language and communication. In M. Dummett (Ed.), *The seas of language* (pp. 166–187). Oxford: Clarendon Press.
- Evans, G. (1973). The causal theory of names. *Proceedings of the Aristotelian Society, Supplementary Volume 47*, 187–208.
- Garrod, S., & Anderson, A. (1987). Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition*, 27, 126–181.
- Garrod, S., & Doherty, G. (1994). Conversation, co-ordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, 53, 181–215.
- Gilbert, M. (1992). *On social facts*. Princeton: Princeton University Press.
- Ginzburg, J. (1995). Resolving questions, I and II. *Linguistics and Philosophy*, 18(5), 459–527, and 18(6), 567–609.
- Ginzburg, J. (2012). *The interactive stance*. Oxford: Oxford University Press.
- Ginzburg, J., & Cooper, R. (2004). Clarification, ellipsis, and the nature of contextual updates. *Linguistics and Philosophy*, 27(3), 297–366.
- Gleitman, L., & Gleitman, H. (1971). *Phrase and paraphrase: Some innovative uses of language*. New York: Norton Press.
- Gleitman, L., Cassidy, K., Nappa, R., Papafragou, A., & Trueswell, J. (2005). Hard words. *Language Learning and Development*, 1(1), 23–64.

- Grice, H. P. (1957). Meaning. *Philosophical Review*, 66: 377–388. Reprinted in Grice (1989), pp. 213–223.
- Grice, H. P. (1975). Logic and conversation. In D. Davison and G. Harman (Eds.), *The Logic of Grammar*, Encino: Dickenson (pp. 64–75). Reprinted in Grice, 1989, pp. 22–40.
- Groenendijk, J., & Stokhof, M. (1990). Dynamic Montague grammar. In L. Kalman & L. Polos (Eds.), *Papers from the second symposium on logic and language* (pp. 3–48). Budapest: Akademiai Kiado.
- Guasti, M. (2004). *Language acquisition: The growth of grammar*. Cambridge: MIT Press.
- Hauser, M., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve. *Science*, 298, 1569–1579.
- Hawthorne, J. (1990). A note on 'languages and language'. *Australasian Journal of Philosophy*, 68(1), 116–118.
- Heck, R. (2014). Semantics and context-dependence: Towards a Strawsonian account. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 327–364). Oxford: Oxford University Press.
- Heim, I. (1982). File change semantics and the familiarity theory of definiteness. In R. Bäuerle, C. Schwarze, & A. von Stechow (Eds.), *Meaning, use and interpretation of language* (pp. 164–189). Berlin: De Gruyter.
- Higginbotham, J. (1992). Truth and understanding. *Philosophical Studies*, 65(1–2), 3–16.
- Isac, D., & Reiss, C. (2008). *I-Language: An introduction to linguistics as cognitive science*. Oxford: Oxford University Press.
- Jackendoff, R. (2002). *Foundations of language*. Oxford: Oxford University Press.
- Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, & M. Stokhof (Eds.), *Formal methods in the study of language* (pp. 277–322). Amsterdam: Mathematisch Centrum.
- Kaplan, D. (1989a). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). New York: Oxford University Press.
- Kaplan, D. (1989b). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 565–614). New York: Oxford University Press.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjective. *Linguistics and Philosophy*, 30(1), 1–45.
- Kennedy, C., & McNally, L. (2010). Color, context, and compositionality. *Synthese*, 174(1), 79–98.
- King, J. (2014a). Speaker intentions in context. *Nous*, 48(2), 219–237.
- King, J. (2014b). The metasemantics of contextual sensitivity. In A. Burgess, B. Sherman, & H. Wettstein (Eds.), *Metasemantics: New essays on the foundations of meaning* (pp. 97–118). Oxford: Oxford University Press.
- King, J., & Stanley, J. (2005). Semantics, pragmatics, and the role of semantic content. In Z. Szabo (Ed.), *Semantics versus pragmatics* (pp. 111–164). Oxford: Clarendon Press.
- Kripke, S. (1977). Speaker's reference and semantic reference. *Midwest Studies in Philosophy*, 2, 255–276.
- Lepore, E., & Stone, M. (2010). Against metaphorical meaning. *Topoi*, 29(2), 165–180.
- Lepore, E., & Stone, M. (2014). *Imagination and convention distinguishing grammar and inference in language*. Oxford: Oxford University Press.
- 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, pp. 3–35). Minneapolis: University of Minnesota Press.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1), 339–359.
- Loar, B. (1981). *Meaning and mind*. Cambridge: Cambridge University Press.
- Ludlow, P. (2014). *Living words: Meaning underdetermination and the dynamic lexicon*. Oxford: Oxford University Press.
- Markson, L., & Bloom, P. (1997). Evidence against a dedicated system of word learning in children. *Nature*, 385(6619), 813–815.
- Mercier, A. (1994). Consumerism and language acquisition. *Linguistics and Philosophy*, 17(5), 499–519.
- Millikan, R. (1984). *Language, thought and other biological categories: New foundations for realism*. Cambridge: MIT Press.
- Millikan, R. (1998). Language conventions made simple. *Journal of Philosophy*, 95(4), 161–180.
- Moravcsik, J. (1998). *Meaning, creativity, and the partial inscrutability of the human mind*. Stanford: Center for the Study of Language and Information.
- Murray, S. (2014). Varieties of update. *Semantics and Pragmatics*, 7, 1–53.

- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy*, 15, 509–559.
- Neale, S. (2005). Pragmatism and binding. In Z. Szabo (Ed.), *Semantics versus pragmatics* (pp. 165–285). Oxford: Clarendon Press.
- Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy*, 3(2), 143–184.
- Pinker, S., & Jackendoff, R. (2005). The faculty of language: What's special about it? *Cognition*, 95, 201–236.
- Recanati, F. (1989). The pragmatics of what is said. *Mind and Language*, 4, 295–329.
- Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.
- Reimer, M. (2004). What malapropisms mean: A reply to Donald Davidson. *Erkenntnis*, 60, 317–334.
- Roberts, C. (1996). Information structure: Towards an integrated formal theory of pragmatics. In J. H. Yoon and A. Kathol (Eds.), *OSUWPL volume 49: Papers in Semantics*, 1996 (pp. 1–53). The Ohio State University Department of Linguistics.
- Roberts, C. (2004). Context in dynamic interpretation. In L. Horn & G. Ward (Eds.), *The handbook of pragmatics*. Oxford: Blackwell.
- Roberts, C. (ms). Solving for interpretation. Columbus: Ohio State University.
- Schiffer, S. (1972). *Meaning*. Oxford: Oxford University Press.
- Schiffer, S. (1993). Actual-language relations. *Philosophical Perspectives*, 7, 231–258.
- Searle, J. (1969). *Speech act*. Cambridge: Cambridge University Press.
- Searle, J. (1978). Literal meaning. *Erkenntnis*, 13(1), 207–224.
- Skyrms, B. (1996). *The evolution of the social contract*. Cambridge: Cambridge University Press.
- Soames, S. (2009). The gap between meaning and assertion. In *Philosophical essays: Natural language: What it means and how we use it* (pp. 278–298). Princeton, NJ: Princeton University Press.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Oxford: Blackwell. (2nd revised edition, 1995).
- Stalnaker, R. (1978). “Assertion” reprinted in *Context and content*. Oxford: Oxford University Press, 1999.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–725.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23, 391–424.
- Stanley, J. (2002). Making it articulated. *Mind and Language*, 17(1&2), 149–168.
- Stanley, J. (2005). Semantics in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning, and truth* (pp. 221–254). Oxford: Oxford University Press.
- Szabo, Z. (2001). Adjectives in context. In I. Kenesei & R. M. Harnish (Eds.), *Perspectives on semantics, pragmatics, and discourse* (pp. 119–146). Amsterdam: John Benjamins Publishing Company.
- Tappenden, J. (1999). Negation, denial and language change in philosophical logic. In D. M. Gabbay & H. Wansing (Eds.), *What is negation?* (pp. 261–298). Berlin: Kluwer Academic Publishers.
- Travis, C. (1985). On what is strictly speaking true. *Canadian Journal of Philosophy*, 15(2), 187–229.
- Trueswell, J., Medina, T., Hafri, A., & Gleitman, L. (2013). Propose but verify: Fast mapping meets cross-situational word learning. *Cognitive Psychology*, 66, 126–156.
- von Fintel, K., & Gillies, A. (2011). ‘Might’ Made Right. In Andy Egan & Brian Weatherson (Eds.), *Epistemic Modality* (pp. 108–130). Oxford: Oxford University Press.
- Wiggins, D. (1997). Languages as social objects. *Philosophy*, 72, 499–524.