Explicit Indirection
Ernie Lepore and Matthew Stone

Introduction

It is tempting to regard our interpretive judgments about so-called "indirect speech acts" as a straightforward reflection of pragmatic reasoning. Take (1), the now stock example from Searle (1975), on its most plausible interpretation:

1. Can you pass the salt?

Literally, it seems that the speaker has asked a question, and it's likely that the hearer will go on to answer that question (Bach and Harnish 1979; Clark 1979), perhaps as in either (2) or (3).

2. No, sorry, I can't reach it.
3. Of course. Here it is.

But (1) is not just a question on its most plausible interpretation. It is also a request. The speaker expects the addressee to pass the salt, and an addressee who fails to realize this—who merely answers the literal question, for example—has not understood the speaker's utterance. In short, the traditional view of utterances such as (1) is that there is a literal meaning, for (1) a question, and simultaneously a further meaning, for (1) a request, which becomes obvious when the utterance is used in a context where the request would be expected and appropriate. When we describe our interpretive judgments in the traditional way, it can look as though it's part of the very data about such cases that they are examples of conversational implicature, in the sense of Grice (1975). In particular, on the traditional view, the interpretation of (1) as a request is not part of its literal meaning but something additional that's derived or "calculated" from that meaning on the assumption that the speaker's use of (1) is intended to advance the established purposes of the conversation. For recent defenses of the traditional view, see for example Bezuidenhout (2016). But you have felt the force of this perspective already if the title of this chapter struck you as paradoxical.

As broadly accepted as it is, the view that the indirection involved in such utterances is a pragmatic phenomenon has far-reaching and, we think, problematic consequences. The fundamental difficulty is that there is overwhelming evidence that grammar governs indirect meanings. As Searle (1975) already noted, grammar distinguishes between utterances that regularly achieve indirect requests, such as (1), from apparently equivalent utterances that do not, such as (4):

4. Are you capable of passing the salt?

Indirect but grammatically licensed directives permit the modifier "please", just like imperatives, but unlike creative hints intended to prompt an action on the part of the addressee (Horn 1989, Lakoff 1973, Sadock 1974) This leads, for example, to the contrast between (5), which sounds natural, and (6), which sounds odd (the sentences are 94 and 95 from Lepore and Stone 2015).

5. I'd like a drink please.
6. #I'm thirsty please.

Finally, the expressions that are assigned such indirect meanings exhibit substantial cross-linguistic variation (Wierzbicka 1985). To reconcile pragmatic accounts of indirection with these facts requires postulating new categories of pragmatic rules in grammar—customs (or conventions in the loose sense of Millikan's 1998 self-perpetuating patterns) that constrain the general rational purposes for which speakers can use language in communication but are integrated into the very architecture of the language faculty. Such rules challenge the view of language—common to both Grice (1975) and Chomsky (e.g., 2005)—as a neutral representational system whose use is whatever people make of it.

Our goal in this chapter is to contest the traditional view of indirection in utterances such as (1) by developing a very different way of characterizing the interpretations involved. We argue that the felt "indirection" of such utterances reflects the kind of meaning the utterances have, rather than the way that meaning is derived. So understood, there is no presumption that indirect meanings involve the pragmatic derivation of enriched contents from a literal interpretation; rather, we argue that indirect meanings are explicitly encoded in grammar.

In particular, we argue that utterances such as (1) work by presenting a complex package of related meanings together as a single unit: on its usual interpretation, an utterance of (1) *first* raises the question whether the addressee can pass the salt and *second* expresses a preference, in the case that the answer is 'yes', that the addressee do so. This move feels indirect, we suggest, because of the weak commitment it imposes on the speaker and the flexible responses it affords the addressee. The move does not commit the speaker to a general preference for action on the part of the addressee; the preference is subject to the addressee's taking it on, by giving a 'yes' answer. Conversely, the preference itself does not even have to be addressed by the addressee for the request to be declined: a 'no' answer gives a coherent response to the open question, but simultaneously renders the speaker's conditional preference inert.

To make our ideas precise, we build on recent work on formalizing declarative, interrogative and imperative meanings as distinct but compatible kinds of content for utterances (Charlow 2011, Starr 2010, MS). In these frameworks, we can straightforwardly formalize our intuitive suggestion that (1) raises a question and then expresses a conditional preference. Moreover, we can formally analyze the information states that result from different potential responses on the part of the

addressee: a 'yes' answer resolves the interaction to a directive followed by compliance, while a 'no' answer leads to a conversational state where the speaker is *not* committed to the directive and the addressee has *not* rejected it. These calculations substantiate our explanation of what felt indirection involves: the indirect utterance is formulated obliquely, in a way that respects both the speaker's authority and the addressee's autonomy (key aspects of Brown and Levinson's 1987 theory of politeness).

The resulting picture allows for straightforward statements of the semantic rules associating utterances like (1) with complex "indirect" interpretations and for a straightforward meta-semantics where these rules amount to conventions for committing to content, broadly in line with the ideas of Lewis (1969, 1979). Thus it is compatible not only with our interpretive intuitions, but with the ample evidence that these interpretations have their origins in the ordinary rules of language.

Background

Our view, then, is that apparently indirect utterances combine multiple semantic contributions. Not surprisingly, philosophers and linguists have developed many similar views, over the years—going back at least to Sadock (1974) and perhaps even to Austin (1962). However, we find all the previous approaches problematic, and we find ourselves differing from them on key semantic and pragmatic issues. We therefore begin by acknowledging our debts to the past and highlighting what's distinctive about our view.

The key differences involve our take on meaning. We think of the content of a discourse in terms of changes to an abstract structure that records different kinds of contributions, following Lewis (1979) and Thomason (1990). We'll call this structure the conversational record. In order to capture the different contents of declarative, interrogative and imperative sentences, we model the conversational record as including not only the propositional information that is taken for granted in the discourse (Stalnaker 1978), but also the open issues that have been raised in the discourse (Ginzburg 2012, Roberts 2012) and the preferences that the discourse establishes (Charlow 2011, Portner 2005, Starr 2010, MS). When we give the meaning of an utterance, then, we need to specify how the utterance changes each of these components of the conversational record. We'll use 'contributing propositions', 'raising questions', and 'establishing preferences' as terms of art that describe the particular semantic effects characteristically associated with declaratives, interrogatives and imperatives.

We think of the conversational record as a level of meaning that is public, determined by language users' deference to semantic conventions and shared practices of meaning making. See Lepore and Stone (2015) for an extended development and defense of this characterization. This gives teeth to the idea that

updating the conversational record is a semantic notion. On our view, the updates associated with utterances are not simply a reflection of speakers' intentions for how the utterances should update the conversational record. Speakers can, for example, have mistaken assumptions about the meanings of utterances, and so they can inadvertently use utterances that contribute propositions, raise questions, or establish preferences that they did not intend. Cutting the familiar Gricean link between meaning and intention represents a deep disagreement with most research on speech acts (including Charlow 2011 and Harris 2014). Our view involves thinking of meaning not as psychological but as social, in the sense of Burge (1979), Kripke (1972), and Putnam (1975). We think this difference is important for making sense of our notion of explicit indirection. For example, we argue later that distinguishing meaning from intention recognition helps us explain the distinctive responses that interlocutors give to the contributions that utterances make explicitly, as distinguished from their engagement with information that other speakers merely make evident.

At the same time, the notion of the state of a discourse is underspecified in its import for the mental states of interlocutors. In particular, the commitments that speakers make in conversation cannot be reduced to such practical attitudes as belief and intention (Stalnaker 2002, Starr 2010, Thomason 1990). Contributing a proposition may commit the speaker to treat that proposition as true for the purposes of the conversation, but it does not commit the speaker either to believe the proposition herself or to intend her audience to believe it. Similarly, raising a question may commit the speaker to treat the question as open for the purposes of the conversation, but it does not commit her to need the answer or to expect one from her audience. An established preference, likewise, need not express the speaker's true desires or her actual intentions for her audience. Some philosophers think that this makes the idea of the conversational record counterintuitive (Harris 2014) or even empty (Bezuidenhout, to appear). By contrast, we argue below that our understanding of the conversational record dovetails with Brown and Levinson's (1987) influential theory of politeness. Politeness often seems to rely on serious utterances that nevertheless carry few consequences for interlocutors' attitudes—as illustrated by the ostensible but transparently insincere offers or the feigned and evidently bogus excuses that we use to smooth our interactions with one another. If this is the right way to think of politeness, then the flexibility of the conversational record must be something that speakers understand and exploit.

Finally, meaning itself is just the starting point for the effects speakers hope to bring about in using utterances. We see meaning as an input to a wide range of imaginative devices that exploit meaning but do not deliver meaning, including such practices as irony, sarcasm and humor (Lepore and Stone 2015). On our view, such practices offer productive and general ways for speakers to use utterances whose meanings establish preferences without undertaking the commitments that come when they use those utterances seriously. And of course, directive utterances, including those with conventional indirect interpretations, can be used ironically, sarcastically, humorously and so forth (Harris 2014).

Given the way we think of meaning, it should be clear that we do not think that the conventional meaning of an utterance can specify the speech act it performs, at least as speech acts are traditionally conceived. Searle (1969, 1975) characterizes speech acts in terms of preconditions and effects involving the mental states of the interlocutors. A request, for example, aims at getting the addressee to do something. In different ways, Charlow (2011) and Harris (2014) elaborate on this perspective to explain how philosophers might analyze directive meaning as a general constraint on the kinds of speech acts that utterances are generally used to perform. Harris in particular goes on to reject the notion of the conversational record and even the existence of conventions of meaning. This is not our view. We characterize the conventional meaning of utterances in terms of updates to the record, which do *not* entail specific preconditions or effects on the mental states of interlocutors and so do *not* accomplish specific *speech acts*, in Searle's or Harris's senses.

In particular, then, we do not suggest that (1) encodes a *request* as a matter of meaning. The directive content of (1) is simply to establish a preference. A speaker can sometimes make a request by using an utterance with such a meaning, but only when (among other conditions) the utterance is serious, the interlocutors have the right relationships, and they are interpreting the commitments of the conversational record in the right way. Since we think that indirection is a matter of semantics, we avoid talk of indirect speech acts or indirect requests, and will try to be explicit about the kinds of meanings we think are really involved. This may sometimes involve a certain amount of circumlocution.[1]

A corollary is that our view is not just a resurrection of the infamous performative hypothesis in grammar (Cresswell 1973, Lewis 1970, Ross 1970). The performative hypothesis is the idea that each main clause is embedded within a syntactically represented and semantically interpreted clause describing the speaker's speech act in using the utterance and involving a covert performative verb. Sadock (1974) is an extended development and defense of the performative hypothesis from a linguistic and philosophical point of view. Linguistically, it is controversial that the postulation of this covert structure explains the relevant syntactic and semantic phenomena as well as its proponents originally claimed (see McCawley 1985 for review). Philosophically, it leads to problems in answering which utterances are truth-evaluable and what their truth conditions are (Boer and Lycan 1980). At the same time, the performative hypothesis requires sentences to have logical forms that are inappropriately specific and that vary implausibly from one utterance of a sentence to the next (see Starr 2010 for discussion). These weaknesses of the

---

[1]Within semantics, there is a longstanding use of 'question' to name a semantic object, on a par with 'proposition', alongside uses naming a syntactic form and a kind of act. We think this terminology is unavoidable and will stick to it. Charitable readers should think of 'request', as used in Lepore and Stone's 2015 discussion of the meanings of indirect speech acts, in analogous terms, as a name for directive content, not a name for a kind of action.

performative hypothesis are well known (see Sadock 2004—we will not repeat the arguments). But it is also well-known that these weaknesses do not extend to current approaches to the semantics of mood, including those on which our account is based (see Harris 2014 and Starr 2010). The reason, of course, is that updates to the record may remain invariant even as the acts speakers accomplish with those updates vary.

Asher and Lascarides (2001) offer a different kind of formal development of the idea that utterances can conventionally combine a constellation of related speech acts. On their view, grammar assigns to conventionally indirect utterances semantic contents of mutually incompatible semantic types—for example, the content of a question and the content of a request. They assume that it is pragmatic reasoning that reconciles these incompatible meanings into a coherent whole, and that the result is in fact a series of related speech acts—for example, speech acts of questioning and of requesting. By contrast, we follow Starr's (2010) semantics, where declarative, interrogative and imperative utterances all denote updates to the conversational record; the different updates just happen to affect different attributes of the record. Thus, we think there is no conflict or coercion involved in first raising a question and then establishing a conditional preference, so there is no obstacle to specifying an indirect meaning explicitly in grammar.

Our differences with Asher and Lascarides (2001) are philosophical as well as formal. Asher and Lascarides insist that indirect readings are calculated and that cooperative reasoning is essential to this process. They therefore describe default rules that produce indirect readings and additional default rules that normally pre-empt the derivation of indirection outside of its conventional range. We deny that there are pragmatic processes of the sort that Asher and Lascarides envisage at work in conventionalized indirection. We think the generalizations that underwrite their default rules can be best explained, not as pragmatic principles, but rather as historical or meta-level generalizations about the kinds of meanings that a language tends to encode.

As we have already hinted, our technical approach follows Starr (2010, MS). His formalism is compatible with a range of philosophical interpretations: his models can characterize individuals' private takes on the conversation (as in Ginzburg 2012), or individuals' occurrent mutual suppositions about the conversation (as in Thomason 1990). We do not claim that Starr subscribes to or would necessarily endorse exactly our understanding of the conversational record as a public social construct.

Most importantly, Starr himself makes no claims about indirect speech acts or about indirection in utterance interpretation more generally. In demonstrating how his formalism could handle conventionalized indirection, we think we are making a philosophical contribution that attests to the strength of his framework. But of course, his framework in no way precludes the exploration of other accounts of indirection (including accounts based on pragmatics).

Formal Model

We now review Starr's (2010, MS) model and apply it to conventionalized indirection. We use the model, as formal semanticists commonly do, to make our empirical claims more precise. The formalism offers a specific realization that shows how propositional information, open issues, and established preferences can constitute an overarching record of the state of the conversation with substantive inferential and communicative dynamics. It lets us specify primitive meanings, compose them together into complex contributions, and explain what follows from them. In particular, we give a meaning for indirection that raises an issue and expresses a conditional preference; we show that the meaning allows for an answer, entails an ordinary directive if the answer is 'yes', and has no directive consequences if the answer is 'no'. (These inferences depend on some rather delicate definitions, so we proceed slowly, via worked examples.) Later, we will link these conversational properties to speakers' intuitions about the polite indirection of utterances such as (1). Thus, the model shows that our intuitive picture of conventionalized indirection is consistent in certain respects and allows us to substantiate our intuitive predictions about our view in a precise way.

Defining Content

We start by presenting the general picture of Starr's model, and giving (slightly streamlined versions) of the key definitions of the formal system. We refer the reader to Starr (2010, MS) for the complete definitions.

To get the logic off the ground, we have a set of possible worlds Ω describing ways the world might be, and an interpretation function I that specifies the extensions of predicates across worlds and the reference of terms. We also need a relation of accessibility A among possible worlds, to interpret 'can' sentences: wAw' just in case w' represents a possible alternative for w. Formulas are constructed from atoms, negation, and possibility operators. An atomic formula $P(n_1 \dots n_k)$ is true at world w if and only if $<I(n_1) \dots I(n_k)> \in I(P)(w)$. Conversely, if ρ is a formula, then $\neg\rho$ is true at w if and only ρ if is not true at w. Finally, if ρ is a formula, then $\Diamond\rho$ is true at w if and only if there is some world w' with wAw' such that ρ is true at w'.

Starr models propositional information using sets of possible worlds. (This idea has been a standard tool since Stalnaker 1978.) The basic operation of contributing information is to start from a set of worlds and narrow the set down to just those where the information is true. If c is a set of worlds and ρ is a formula, we use the notation $c[\rho]$ for $\{ w \in c \mid \rho$ is true at $w \}$.

The content of a discourse is not limited to propositional information, however. Discourse can also raise questions. Semantically, Starr proposes to model questions as sets of their possible answers; a question expresses an interest in establishing that one of the answer propositions is true, but which one, of course, remains to be determined. This is a familiar idea going back to Hamblin (1958). Now, by

representing discourse content as a set of alternatives, we can capture both declarative meaning and interrogative meaning; see Ciardelli, Groenendijk and Roelofsen (2015). The open questions of the discourse concern the differences among the alternative possibilities; its propositional content, meanwhile, consists in what all the alternatives have in common. We can record the contribution of new propositional information by selecting and refining the alternatives. We can record the raising of a new question by introducing additional alternatives to describe the possible ways the question could be resolved. Since many questions are normally on the table, these alternatives involve not only answering the question individually but also providing detailed answers that address this question in tandem with others.

Imperative meaning, for Starr, also involves alternatives. Imperatives involve a preference for one alternative over another. Starr models this with a binary relation R: given sets of possible worlds a and a' that are alternatives in the discourse, aRa' indicates that the content of the discourse establishes the preference that a is better than a'. This is naturally understood as a transitive, asymmetric relation (assuming that the preferences established by a discourse are consistent). It's also convenient to assume that the necessarily false proposition $\emptyset$ is always part of the domain of R and that consistent alternatives are always preferred to it. That way we can use R simply to identify a set of consistent alternatives, even when R establishes no substantive preferences among them.

In short, we can represent discourse content in terms of a binary preference relation R over propositions. The domain of R specifies the set of alternatives at issue in the discourse—encoding possible answers to the open questions raised in the discourse. The propositional content of the discourse is the disjunction of these propositions—giving the information that all the possible answers have in common.

Let's give a concrete example to illustrate the features of the formalism and motivate the definitions to follow. Consider a simple, idealized situation: there is an upcoming concert with four possible singers: Chris, Kim, Robin and Sandy. This gives a model with sixteen possible worlds, each of which we can indicate by listing who will sing (they are: ckrs, ckr, cks, ck, crs, cr, cs, c, krs, kr, ks, k, s, r, -). Suppose that a discourse says that Chris will sing, asks whether Kim will sing, commands that Robin sing, and is silent about Sandy. Let's compose the relation R_1 that we will use to represent the content of the discourse.

First, let's consider the alternatives. The discourse has explicitly raised the question whether Kim will sing. It also needs alternatives for whether Robin will sing, since it prefers that Robin sing. And it needs alternatives that give answers to both questions simultaneously. All the alternatives must reflect the fact that Chris will sing. And they should all make no commitment about Sandy one way or the other. That gives eight propositions, to which we add the information that Chris will sing and the necessarily false proposition to get the roster in (7):

7. {ckrs, ckr, cks, ck, crs, cr, cs, c } – Chris will sing.
   {ckrs, ckr, cks, ck} – Will Kim sing? Yes.
   {crs, cr, cs, c} – Will Kim sing? No.
   {ckrs, ckr, crs, cr} – Will Robin sing? Yes.
   {cks, ck, cs, c} – Will Robin sing? No.
   {ckrs, ckr} – Which of Kim and Robin will sing? Both.
   {cks, ck} – Which of Kim and Robin will sing? Kim.
   {crs, cr} – Which of Kim and Robin will sing? Robin.
   {cs, c} – Which of Kim and Robin will sing? Neither.
   ∅ – absurdity

We prefer that Robin sing, other things being equal: better Robin sings than not, better Robin sings with Kim if Kim sings, and better Robin sings without Kim if Kim does not sing. We also prefer all the live possibilities in (7) to the necessarily false proposition ∅. Thus, the relation R_1 that we get is given by the tuples in (8).

8. {⟨{ckrs, ckr, cks, ck, crs, cr, cs, c }, ∅ ⟩,
   ⟨{ckrs, ckr, cks, ck}, ∅ ⟩,
   ⟨{crs, cr, cs, c}, ∅ ⟩,
   ⟨{ckrs, ckr, crs, cr}, ∅ ⟩,
   ⟨{cks, ck, cs, c}, ∅ ⟩,
   ⟨{ckrs, ckr}, ∅ ⟩,
   ⟨{cks, ck}, ∅ ⟩,
   ⟨{crs, cr}, ∅ ⟩,
   ⟨{cs, c}, ∅ ⟩,
   ⟨{ckrs, ckr, crs, cr}, {cks, ck, cs, c}⟩,
   ⟨{ckrs, ckr}, {cks, ck}⟩,
   ⟨{crs, cr}, {cs, c}⟩}

This example might seem slightly involved, but it will allow us to give worked examples that walk through the definitions of updates precisely.

Capturing Dynamics

To describe discourse content, Starr moves from a level of formulas to a level of updates that transform the relationships that model discourse state. For example, an update that contributes propositional information requires us to incorporate that information into all the available alternatives. The information may answer an open question: in this case the information will also eliminate the alternatives that are incompatible with the information it provides. We formalize these effects using the definition in (9). The notation ▷ρ denotes an update that contributes the propositional information specified by the formula ρ.

9. $R[▷\rho] := \{ \langle a[\rho], a'[\rho] \rangle \mid aRa'$ and $a[\rho] \neq \emptyset \}$

According to the definition in (9), we update a context R with ▷ρ by taking each of the tuples aRa' and restricting the alternatives a and a' to reflect the information ρ –

as long as this is consistent with *a*. Thus, suppose we update R_1 with the information that Sandy will sing. In this case, we will simply refine the possibilities—eliminating all the worlds where Sandy will not sing from consideration. The relation we get for $R_1[\rhd\text{sing(Sandy)}]$ is therefore:

10. {⟨{ckrs, cks, crs, cs }, ∅ ⟩,
⟨{ckrs, cks}, ∅ ⟩,
⟨{crs, cs}, ∅ ⟩,
⟨{ckrs, crs}, ∅ ⟩,
⟨{cks, cs}, ∅ ⟩,
⟨{ckrs}, ∅ ⟩,
⟨{cks}, ∅ ⟩,
⟨{crs}, ∅ ⟩,
⟨{cs}, ∅ ⟩,
⟨{ckrs, crs}, {cks, cs}⟩,
⟨{ckrs}, {cks}⟩,
⟨{crs}, {cs}⟩}

On the other hand, suppose we learn that Kim will sing. The discourse is already structured to distinguish alternatives where Kim will sing from alternatives where Kim will not sing: those where Kim will not sing will now disappear. The relation we get for $R_1[\rhd\text{sing(Kim)}]$ is therefore simply:

11. {⟨{ckrs, ckr, cks, ck}, ∅ ⟩,
⟨{ckrs, ckr}, ∅ ⟩,
⟨{cks, ck}, ∅ ⟩,
⟨{ckrs, ckr}, {cks, ck}⟩}

Now imagine starting with a basic relation R_0 corresponding to complete ignorance about our example universe. We can specify R_0 as in (12).

12. { ⟨{ckrs, ckr, cks, ck, crs, cr, cs, c, krs, kr, ks, k, s, r, -}, ∅ ⟩}

Then $R_0[\rhd\text{sing(Chris)}]$ is simply:

13. { ⟨{ckrs, ckr, cks, ck, crs, cr, cs, c}, ∅ ⟩}

What about interrogatives? Contributing a new question just anticipates the possibility of a 'yes' answer or a 'no' answer as additional alternatives for the discourse. We formalize this in (14), where ?ρ denotes an update that raises the question whether ρ is true.

14. $R[?\rho] := R \cup R[\rhd\rho] \cup R[\rhd\neg\rho]$

It's best to start with a simple illustration: $R_0[\rhd\text{sing(Chris)}][?\text{sing(Kim)}]$. We combine the alternatives specified in (13) with the alternatives we get from the

update that Kim will sing and the alternatives we get from the update that Kim will not sing. That gives three possibilities:

15. $\{\langle\{\text{ckrs, ckr, cks, ck, crs, cr, cs, c}\}, \emptyset\rangle,$
   $\langle\{\text{ckrs, ckr, cks, ck}\}, \emptyset\rangle,$
   $\langle\{\text{crs, cr, cs, c}\}, \emptyset\rangle\}$

For a more complex example, consider constructing the relation $R_1[?\text{sing(Sandy)}]$. According to (14), we start with the relation R_1, as given in (8). Then we add the tuples in the relation $R_1[\triangleright\text{sing(Sandy)}]$, as given in (10). Finally, we add a complementary set of tuples derived from $R_1[\triangleright\neg\text{sing(Sandy)}]$. So the discourse allows for all the possible answers it did so far, allows any of those answers to be combined with the information that Sandy will sing if that's consistent, and also allows any of those answers to be combined with the information that Sandy will not sing if that's consistent. The preferences among these alternatives are inherited from the preferences already established in the discourse. The representation that results is rather cumbersome to write down, since the number of possible compound answers grows exponentially in the number of questions introduced. But the idea, as formalized in (14), should be clear.

What about establishing a preference for a proposition ρ? To start, this has to raise the issue of whether ρ is true; we will need to distinguish the ρ outcomes from the $\neg\rho$ outcomes. But we also have to relate the alternatives we have in the right way. We may have new alternatives of the form $a[\rho]$ that now should be preferred to corresponding alternatives of the form $a[\neg\rho]$. But, we also want to use our existing preferences of the form aRa' transitively, to encode our derived preference for consistent alternatives $a[\rho]$ over corresponding alternatives $a'[\neg\rho]$. This leads to the definition in (16), where $!\rho$ denotes the imperative update that establishes a preference for ρ over $\neg\rho$.

16. $R[!\rho] := R[?\rho] \cup \{\langle a[\rho], a[\neg\rho]\rangle \mid a \in \text{dom}(R[?\rho]) \text{ and } a[\rho] \neq \emptyset\} \cup$
   $\{\langle a[\rho], a'[\neg\rho]\rangle \mid aR[?\rho]a' \text{ and } a[\rho] \neq \emptyset\}$

We leave it to the reader to check using definition (16) and the key intermediate result in (15) that $R_0[\triangleright\text{sing(Chris)}][?\text{sing(Kim)}][!\text{sing(Robin)}]$ gives precisely the relation R_1 presented in (8). As it happens, in this case, the order of the updates does not matter: we derive the same relation R_1 from $R_0[!\text{sing(Robin)}][\triangleright\text{sing(Chris)}][?\text{sing(Kim)}]$, $R_0[?\text{sing(Kim)}][!\text{sing(Robin)}][\triangleright\text{sing(Chris)}]$, and so forth.

As a matter of notation, we can introduce an operation of sequencing as in (17).

17. $R[\varphi; \psi] := R[\varphi][\psi]$

This allows us to compose together meanings that convey complex constellations of content—for example, to describe R_1 as obtained directly from R_0 by an update with the content of an entire discourse: $R_0[\triangleright sing(Chris); ?sing(Kim); !sing(Robin)]$.

The last ingredient of the formalism is the conditional. The definition has two parts: a *test* describing the import of conditional information, and an *update* describing the import of conditional questions and preferences.

   18. $R[if\ (\varphi)\ (\psi)] := R \cup R[\varphi; \psi]$ if $\bigcup dom(R[\varphi]) = \bigcup dom(R[\varphi; \psi])$
                      $\emptyset$ otherwise.

The test makes sure that the information $\bigcup dom(R[\varphi; \psi])$ given by the antecedent and consequent together does not go beyond the information $\bigcup dom(R[\varphi])$ obtained simply from considering the antecedent itself. In other words, we pass the test if the consequent is informationally redundant given the content we get by taking on the antecedent.[2] If the test fails, the conditional is inconsistent with our present information, so the output state of the conditional is trivial. Of course, raising questions and expressing preferences don't give any new information; they have other effects. So conditional questions and conditional preferences will always pass this test.

If the conditional generalization is already implicit in the information that we have, the rule will construct a new output state in a certain way, designed with questions and preferences in mind. The output makes reference to an updated state $R[\varphi; \psi]$— this captures the contributions made by φ and ψ. Normally φ contributes information, so this involves taking on the information given by φ, and then incorporating the contributions of ψ (adding information, raising questions or expressing preferences) that apply just to those worlds where φ is true. This state is then *combined* with the initial state R by *set union*. In other words, the update of the conditional preserves all the open questions and preferences we started with, but adds some new ones: we're interested in the answer if φ is true, as well as a range of further questions and preferences that will come into play if the answer is 'yes'.

For our purposes in this paper, the key thing about definition (18) is in the way it handles conditional imperatives. Let's return to R_0 and consider $R_0[if\ (\triangleright sing(Chris))$

---

[2] We think it would be good to develop a version of Starr's system that gives conditionals content of their own. Such a system might be based on the work of Stojnic (to appear), who provides a dynamic semantics for modal discourse that gives 'if' statements the truth conditions of strict conditionals, but it would need to be extended to handle questions and imperatives. Such an extension remains a project for future work. For now, we note that the substance of our paper—explicit indirection—does not depend on the content of conditional assertions; Starr's formalism is therefore enough to demonstrate the consistency of our ideas and give a precise system that substantiates our intuitive predictions.

(!sing(Robin))]. According to (18), we first compute $R_0[\triangleright\text{sing}(\text{Chris})]$: recall that that's the relation given in (13). Then we compute $R_0[\triangleright\text{sing}(\text{Chris}); !\text{sing}(\text{Robin})]$. The reader can check that this is the relation given in (19).

19. {⟨{ckrs, ckr, cks, ck, crs, cr, cs, c}, ∅⟩,
⟨{ckrs, ckr, crs, cr}, ∅⟩,
⟨{cks, ck, cs, c}, ∅⟩,
⟨{ckrs, ckr, crs, cr}, {cks, ck, cs, c}⟩}

It's easy to see that the these relations satisfy the conditional side condition $\bigcup\text{dom}(R_0[\triangleright\text{sing}(\text{Chris})]) = \bigcup\text{dom}(R_0[\triangleright\text{sing}(\text{Chris}); !\text{sing}(\text{Robin})])$. In both cases, the relevant propositional information is just the information that Chris will sing. That means that $R_0[\text{if }(\triangleright\text{sing}(\text{Chris}))\ (!\text{sing}(\text{Robin}))]$ is the relation given in (20).

20. {⟨{ckrs, ckr, cks, ck, crs, cr, cs, c, krs, kr, ks, k, s, r, -}, ∅⟩,
⟨{ckrs, ckr, cks, ck, crs, cr, cs, c}, ∅⟩,
⟨{ckrs, ckr, crs, cr}, ∅⟩,
⟨{cks, ck, cs, c}, ∅⟩,
⟨{ckrs, ckr, crs, cr}, {cks, ck, cs, c}⟩}

Call this relation R_2. The key element of this relation is the final preference: better Chris and Robin both sing than Chris sings and Robin doesn't. This preference is part of the overall output state, but it's conditional. Our information doesn't say that Chris sings; Chris might or might not. If Chris doesn't sing, all bets are off. But if Chris does sing, Robin should sing too.

Interaction potential

Now we get the payoff. We can use this combination of raising questions, introducing preferences, conditionals and sequencing to capture the indirection of utterances such as (1).

Let's give the insight by contrasting two further updates based on our most recent example R_2: $R_2[\triangleright\text{sing}(\text{Chris})]$ and $R_2[\triangleright\neg\text{sing}(\text{Chris})]$.

$R_2[\triangleright\text{sing}(\text{Chris})]$ imposes the constraint, which most of the alternatives of R_2 already satisfy, that Chris will sing. That simply returns us to the relation given by (19). This illustrates the general fact about the system that if $R[\text{if }(\varphi)\ (\psi)]$ is nonempty then $R[\text{if }(\varphi)\ (\psi); \varphi]$ is exactly the same relation as $R[\varphi; \psi]$. The expected generalization of modus ponens holds. But in this case, it means that if we get the information specified by the antecedent, a conditional preference is not conditional any more. The preference holds generally.

$R_2[\triangleright\neg\text{sing}(\text{Chris})]$, meanwhile, imposes the constraint that Chris will not sing. This is compatible with the overall information in the discourse, but it's inconsistent with all of the other alternatives; they all involve Chris singing. That means that these alternatives, and the preferences over them, get eliminated. The final result is just

(21). In other words, if we get information ruling out the antecedent, then a conditional preference is again not conditional any more: The preference simply disappears.

    21. $\{\langle\ \{\text{krs, kr, ks, k, s, r, -}\}, \emptyset\ \rangle\}$

Now, finally, consider the qualitative behavior of a relation specified as in (22).

    22. R[?◊sing(Chris) ; if (▷◊sing(Chris) (!sing(Chris))]

This updates R in two ways. First it raises the question of whether it's possible for Chris to sing. The potential answers—that Chris can sing, that Chris can't—remain open issues in the discourse until at some point one or the other answer arrives. Then, it establishes a conditional preference: The worlds where it's possible for Chris to sing and Chris does sing are preferred to the worlds where it's possible for Chris to sing but Chris does not sing. Just as in the case we just considered, this preference comes into effect if we learn that Chris can sing. There's then nothing conditional about it any more. By contrast, if we learn that Chris cannot sing, the preference disappears. It's as though it was never there.

So where does (22) leave us, when we consider the results of the update as a whole? Well, it's still open whether Chris can sing or not. Nothing about the conditional changes that. So we'd expect the discourse to continue with an answer, 'yes' or 'no'. If the answer is 'yes', the discourse now involves an unrestricted preference for Chris singing. Thus, if the answer is 'yes' (and the relevant background is in place), we'd expect the interlocutors to work to bring it about that Chris sings. Again, however, if the answer is 'no', there's no such preference, and no such implications about the interlocutors' actions.

In general, then, we claim that an update of the form '?◊p ; if (▷◊p) (!p)' captures the *indirect meaning* found in the key reading of (1). This update captures both the intuitive content of this kind of indirect utterance and the intuitive follow-ups that those utterances have in conversation.[3] But it's a meaning: indirection, on our

---

[3]Given our formal theory, we note that there will be a range of possible ways to develop similar ideas. For one thing, we need not take '?◊p ; if (▷◊p) (!p)' as an *analysis* of indirect meaning. It's enough that whatever the meaning is of indirect utterances such as (1), that meaning has the entailments that we show that '?◊p ; if (▷◊p) (!p)' has. Thus, philosophers who insist that lexical meanings must be semantic primitives (e.g., Fodor 1970) can still accept the idea of primitive indirect meanings. Moreover, philosophers who contest Brown and Levinson's (1987) account of politeness might deny that the preference expressed by utterances such as (1) is in any sense conditional. They are free to base an account of indirection on something like '?◊p ; !p'. The formalism shows that it is possible to specify indirect meaning in these terms as well. However these choices are resolved, it won't challenge our key claim: that indirect meanings can be specified by the ordinary

analysis, describes the way the move is packaged together and bracketed so that it results in interactive effects that are different both from those of plain interrogatives and from those of unconditional expressions of preference. In particular, this move puts the preference forward only conditionally, and simultaneously raises an issue that the addressee can respond to in such a way as to silently neutralize that potentially problematic contribution. Thus, the move is indirect because its status (in some sense, even its very existence) depends on the addressee's answer. It should now be clear why we think indirection need not be a matter of pragmatic calculation. Our view is that indirect meanings such as '?◊p ; if (▷◊p) (!p)' are explicitly delivered by the grammar, whenever an utterance type is conventionally used to accomplish this indirection.

Discussion

In this section, we expand on the implications of the view that we have just articulated and formalized.

Politeness

Indirect requests, such as (1), are more polite than corresponding direct imperatives, such as (23).

    23. Pass the salt

Brown and Levinson (1987) offer an explanation. They see politeness as a reflection of a range of ideals that we aim to foster in our dealings with one another, including showing our good feelings for others and respecting their independence and autonomy. Politeness, for them, is a strategy that helps to defuse potential conflicts between these ideals and certain kinds of utterances. For example, to make a request is to threaten your interlocutor's autonomy: you are telling them what to do. Politeness demands that you make it easy for your interlocutor to opt out of the request. Explicit indirection offers a way to do this. It lets the addressee who's so inclined fib, saying that the request would be impossible to fulfill. Then it's as if the request was never made, and the addressee's autonomy, never challenged. Our formalism transparently implements this explanation.

---

semantic conventions of grammar. Finally, of course, it remains to derive indirect meanings for utterances compositionally from their parts and the way they are put together. To pursue this would require enriching Starr's formalism with the resources of the λ-calculus, so that we could factor meanings of the form '?◊p ; if (▷◊p) (!p)' into suitable lexical elements and the contributions of rules of combination. For example, one strategy might be to describe a new interrogative meaning for 'can', parameterized by individual x and property P, as in '$\lambda x\, \lambda P$: ?◊Px ; if (▷◊Px) (!Px)', and allow meaning to compose with a particular subject and predicate via function application. We might go further and represent this meaning as the output of a lexical rule, an option that Asher and Lascarides (2001) recommend.

You might think that Brown and Levinson's explanation of the politeness of indirect requests would be compatible with a wide range of views about the meanings of utterances and the contributions they make. We disagree. We think that making Brown and Levinson's proposal precise *depends* on our view that utterances merely update the conversational record, an abstract construct that is removed from speakers' actual beliefs and intentions. For example, you cannot give this kind of explanation, we think, if you assume that declarative utterances are conventionally used in attempts to produce beliefs on the part of the addressee. If speakers who make excuses are obligated to convince hearers that they are true, or if hearers must believe excuses to accept them, then making requests through indirection won't make declining them any smoother. It merely pushes the problems down the road. However, we think English speakers understand that excuses in these contexts are *pro forma,* and should neither be believed nor challenged. That's why indirection makes it easy to offer excuses. In short, if Brown and Levinson are right about the politeness of indirection, "putting information on the record" is the attitude that most happily reconciles the public commitments utterances are used for with interlocutors' private interests and perspectives.

Intention Recognition

Issues of autonomy and politeness also loom large when we consider the role of intention recognition in indirection. Let's start with an example. We're at a yoga studio, and the speaker is late to reserve a spot at a future class that often fills up. She might use either (24) or (25) to open a discussion with the staff at the registration desk about getting herself added to it.

24. Can I still get a spot in tomorrow's 6:30 class?
25. Is it still possible to reserve a spot in tomorrow's 6:30 class?

On our view, (24) appeals to conventional indirection to establish a conditional preference, but (25) does not. However, when we consider the implications of the utterances for the speaker's intentions, the difference between them seems very small indeed. In both cases, the speaker doesn't know if there is a free spot, but the speaker would like one if there is one. In both cases, the speaker expects her interlocutor to take this into account—for example, by checking availability through a reservation system that will allow her to be quickly added to the class if the opportunity arises. And of course, in both cases, it is the utterance that signals to the audience that this is what the speaker has in mind. A cooperative interlocutor will think the same thing about the speaker's plans no matter which utterance the speaker uses.

But a polite interlocutor may well respond differently to the two utterances, with (26) for (24) but with (27) for (25).

26. Sure. You're in.
27. It is. Would you like one?

In Brown and Levinson's framework, there's no problem with complying with explicitly encoded requests. However, autonomy demands that polite interlocutors ask before acting towards a speaker's inferred goals on her behalf.

Because we are skeptical that meaning can be reduced to a speaker's intentions for getting information across to her addressee, we think it's no surprise that there's a difference between encoding one's preferences and making those preferences obvious. And, we have argued, conventionalized indirection is a case of encoding.

Ambiguity

Our discussion so far has focused on the ability of our account to capture intuitive judgments about indirection. But the dissatisfaction that linguists and philosophers have with the idea of conventionalized indirection often depends on its relationships to broader issues. One of these is ambiguity.

Indirection is ambiguous. For example, sometimes 'can' questions just raise questions about ability. Sometimes they convey a richer meaning involving indirection. Lepore and Stone (2015) use (28, their 89) to make these different readings palpable.

    28. Can you play Chopin's E minor prelude?

With (1), it's normally obvious that the addressee can pass the salt, so it's hard to make sense of the utterance as a mere question. By contrast, (28) raises a substantive issue. (28) makes sense as a question about how good a pianist the addressee is; it also makes sense as a case of indirection, which also expresses a (conditional) preference for the addressee to play the piece now. As expected, we can formalize the two readings as in (29) and (30).

    29. $?\lozenge p$
    30. $?\lozenge p$; if ($\rhd \lozenge p$) (!p)

Our view, of course, is that the (29) interpretation and the (30) interpretation are alternatives that are both made available by the grammar, just as in the case of any other ambiguity. The formalism makes it clear that we predict the possibility of including a direct answer in response to the indirect meaning observed by Bach and Harnish (1979) and Clark (1979). Both (29) and (30) raise the issue of whether playing the E minor prelude is possible and so both afford an answer as a response.

Readers unused to thinking about lexical semantics may be surprised to find grammatically specified alternative meanings with such close affinities between them as (29) and (30) exhibit. In fact, this state of affairs is so common that it has its own name: "autohyponymy" (Horn 1984). We think 'can' is one of many verbs that have these kinds of overlapping senses. For an analogous example, consider 'climb', discussed extensively by Jackendoff (1990) and Hanks (2013). The action of climbing sometimes involves upward motion, sometimes involves hand-over-hand

clambering, and prototypically involves both. The availability of these different interpretations depends in complex ways not only on what does the climbing, but also on which other arguments are present and how they are realized syntactically. Thus, the different interpretations are clearly a matter of English grammar. But now imagine the conversation in (31) and (32), uttered by interlocutors on the middle deck of a New York City fire escape.

31. A: How did you get here?
32. B: I climbed.

B's response in (32) probably suggests that she climbed up from the ground. But it's also true, on the weaker but related sense of 'climb', if she climbed down from the roof. The different interpretations of 'climb' at play in (32) stand in exactly the same semantic relationships to one another as do the different interpretations (29) and (30) that our account of conventionalized indirection posits for (28).

Explanation in Semantics

Ambiguity, then, is a possible view, but many linguists and philosophers think it is an unattractive one. If we showed how to derive the indirect interpretation from the literal one by general principles, wouldn't we strike the indirect interpretation from the roster of grammatical meanings and so reduce the complexity of speakers' knowledge of language? Grice argued that we would: this general argument in favor of pragmatic accounts, his "modified Ockham's razor" (1975), has come to be known as Grice's razor.

In the case of conventionalized indirection, we think the question is moot. There are no general interpretive principles that can derive indirect meanings from literal ones exactly where they occur within and across languages. Languages vary in how they accomplish indirection; you have to learn how each one does it. You have to learn, for example, that (33) in Hungarian is an apology or that (34) in Japanese is a request.

33. Ne haragudjon.
    (Don't be angry.)
34. Kikasete itadakemasen ka?
    (Can't you do us the favor of having us listen?)

Lepore and Stone (2015) discuss these examples—(33) is their (99), taken from Suszcyńska (1999); (34) is their (101), taken from Horvat (2000)—and make the case for conventionality in detail. Despite frequent claims otherwise in the literature, going back to Searle (1975) and Morgan (1978), Lepore and Stone argue that the conventions involved are semantic in the most straightforward sense. English speakers use the formulation "I'm sorry" and not "Don't be angry" to express an apology indirectly. In so doing, they are coordinating on the update the utterance makes to the conversational record, and using only their learned mutual

expectations to do so. In short, this is a conventional meaning, in Lewis's (1969) sense of convention.

Nevertheless, the persistence of pragmatic accounts of indirection shows how convincing researchers find Grice's razor. We find the principle much less compelling, however. We think that researchers would often be led to better empirical theories, and a deeper and more nuanced understanding of linguistic meaning, if they approached the principle more skeptically.

Conventional meaning grows out of improvised meaning through a mixture of insight, luck, precedent and repetition. That's a key feature of Lewis's (1969) account of convention—one that later approaches, such as Millikan (2005) or Ludlow (2014) have only sought to strengthen and extend. The dynamic evolution of language through improvisation and conventionalization readily transmutes word meanings from their established senses into new related but distinct ones. Polysemy is thus ubiquitous in language. To attempt to explain it away synchronically through pragmatic principles is to miss the historical contingencies that really are at work in shaping speakers' linguistic knowledge, as evidenced in examples such as (33) and (34).

Conclusion

Indirection is a kind of meaning: meaning that combines multiple contributions—a first, explicit contribution that can be addressed unproblematically, and a second, more difficult contribution whose potential embarrassment can be neutralized by certain of the coherent responses that the first makes available. Indirection is not diagnostic of implicature or other kinds of pragmatic reasoning in interpretation. The indirect contribution is not in any sense derived from the literal one; both are specified by grammar. Speakers' linguistic judgments and their conversational interactions make sense only if indirection is explicitly encoded. Some ambiguities inevitably result, but they are attested, comparable to other semantic ambiguities, and to be expected given the nature of meaning in language.

We have presented a formal account that makes these claims precise, for the specific case of English 'can' questions as indirect expressions of preferences, based on the work of Starr (2010, MS). We have grounded our philosophical interpretation of the formalism in the perspective of Lepore and Stone (2015). This shows that our account of explicit indirection is consistent, and is consonant with the general approach to semantics that we favor.

We noted, however, that Starr's formalism can be interpreted in different ways. Moreover, we expect that a category of conventional indirect meaning can be formalized on similar lines using speech act approaches to meaning (Charlow 2011, Harris 2014) or truth-conditional reductions of imperative and interrogative meaning (as in Han 1998 or Kaufmann 2012). Thus, substantial opportunities

remain for future explorations of conventional indirect meaning. Those investigations will certainly have to address the full range of polite indirection, not only in English but across languages.

Nevertheless, we think our discussion has revealed some subtle philosophical and linguistic challenges that must be met to describe indirection as conventional. In retrospect, Lepore and Stone (2015) are too cavalier in making the claims about conventionalized indirection that they do, without specifying the formal and philosophical details. In explicating and formalizing the view, we hope to illustrate the kind of detailed development that remains to be done to get clear on the rules of language and their role in guiding people's communicative interactions with one another.

Acknowledgments

References

Asher, Nicholas, and Alex Lascarides. 2001. Indirect speech acts. Synthese, 128(1–2): 183–228.

Austin, J.L. 1962. How to Do Things With Words. Oxford: Oxford University Press.

Bach, Kent, and Robert M. Harnish. 1979. Linguistic Communication and Speech Acts. Cambridge, MA: MIT Press.

Bezuidenhout, Ann. 2016. What properly belongs to grammar? A response to Lepore and Stone. Inquiry 56(2):175-194.

Bezuidenhout, Ann. To appear. The conversational record. In G. Preyer (ed.), Semantics, Pragmatics and Interpretation.

Boër, Steven E., and William G. Lycan. 1980. A performadox in truth-conditional semantics. Linguistics and Philosophy 4:1-46.

Brown, Penelope and Stephen C. Levinson. 1987. Politeness: Some universals in language usage. Cambridge: Cambridge University Press.

Burge, Tyler, 1979. "Individualism and the Mental," in P. A. French, T. E. Uehling, and H. K. Wettstein (eds.) Midwest Studies in Philosophy, IV, 73–121. Minneapolis: University of Minnesota Press.

Charlow, Nathan A. 2011. Practical Language: Its Meaning and Use. PhD Dissertation, University of Michigan.

Chomasky, Noam. 2005. Three Factors in Language Design. Linguistic Inquiry 36(1): 1–22.

Ciardelli, Ivano, Jeroen Groenendijk, and Floris Roelofsen. 2015. On the semantics and logic of declaratives and interrogatives. Synthese, 192(6): 1689–728.

Clark, Herbert H. 1979. Responding to indirect speech acts. Cognitive Psychology 11: 430–477.

Cresswell, Max J. (1973). Logics and Languages. London: Methuen.

Fodor, J. A. (1970). Three Reasons for Not Deriving "Kill" from "Cause to Die". Linguistic Inquiry 1(4): 429–438.

Ginzburg, Jonathan. 2012. The Interactive Stance. Oxford: Oxford University Press.

Grice, H. Paul. 1975. Logic and Conversation. In P. Cole and J. Morgan (eds), Syntax and Semantics III: Speech Acts, 41–58. New York: Academic Press.

Han, Chung-Hye. 1998. The Structure and Interpretation of Imperatives: Mood and Force in Universal Grammar. PhD Dissertation, University of Pennsylvania.

Hamblin, C. L. 1958. Questions. Australasian Journal of Philosophy, 36:159–68.

Hanks, Patrick. 2013. Lexical Analysis: Norms and Expectations. Cambridge, MA: MIT Press.

Harris, Daniel. 2014. Speech Act Theoretic Semantics. PhD Dissertation, City University of New York.

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

Horn, Laurence R. 1989. A Natural History of Negation. Chicago: University of Chicago Press.

Horvat, Andrew. 2000. Japanese beyond Words: How to Walk and Talk like a Native Speaker. Berkeley, CA: Stone Bridge Press.

Jackendoff, Ray. 1990. Semantic Structures. Cambridge, MA: MIT Press.

Kaufmann, Magdalena. 2012. Interpreting Imperatives. Heidelberg: Springer.

Kripke, Saul. 1972. Naming and Necessity. In Gilbert Harman and Donald Davidson (eds), Semantics of Natural Language, 253–355. Dordrecht: Reidel.

Lakoff, George. 1973. Some Thoughts on Transderivational Constraints. In B. B. Kachru (ed.), Issues in Linguistics, 442–52. Champaign, IL: University of Illinois Press.

Lepore, Ernie and Matthew Stone. 2015. Imagination and Convention: Distinguishing Grammar from Inference in Language. Oxford: Oxford Univeristy Press.

Lewis, David K. 1969. Convention: A Philosophical Study. Cambridge, MA: Harvard University Press.

Lewis, David K. 1970. General Semantics. Synthese, 22(1/2): 18–67.

Lewis, David. 1979. Scorekeeping in a Language Game. Journal of Philosophical Logic, 8(3): 339–59.

Ludlow, Peter. 2014. Living Words: Meaning Underdetermination and the Dynamic Lexicon. Oxford: Oxford University Press.

McCawley, James D. 1985. What Price the Performative Analyisis? University of Chicago Working Papers in Linguistics, Vol. 1, 43-64. Chicago: Department of Linguistics, University of Chicago.

Millikan, Ruth G. 1998. Language Conventions Made Simple. Journal of Philosophy 95(4): 161–80.

Millikan, Ruth G. 2005. Semantics/Pragmatics: (Purposes and Cross-Purposes). In Language: A Biological Model. Oxford: Oxford University Press, 187-220.

Portner, Paul. 2005. The semantics of imperatives within a theory of clause types. In Kazuha Watanabe and Robert B. Young, eds., Proceedings of Semantics and Linguistic Theory 14. Ithaca, NY: CLC Publications.

Putnam, Hilary. 1975. The Meaning of "Meaning." In Mind, Language and Reality, 215–71. Cambridge: Cambridge University Press.

Roberts, Craige. 2012. Information Structure: Towards an Integrated Formal Theory of Pragmatics. Semantics and Pragmatics, 5(6): 1–69.

Ross, John Robert. 1970. On Declarative Sentences. In R. Jacobs & P.S. Rosenbaum (eds.), Readings in English Transformational Grammar, 222–272. Waltham, MA: Ginn.

Sadock, Jerrold. 1974. Towards a Linguistic Theory of Speech Acts. New York: Academic Press.

Sadock, Jerrold. 2004. Speech Acts. In Laurence R. Horn and Gregory Ward (eds), The Handbook of Pragmatics, 53–73. Oxford: Blackwell.

Searle, John R. 1969. Speech Acts. Cambridge: Cambridge University Press.

Searle, John R. 1975. Indirect Speech Acts. In P. Cole and J.Morgan (eds), Syntax and Semantics III. Speech Acts, 59–82. New York: Academic Press.

Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), Syntax and Semantics, IX. Pragmatics, 315–32. New York: Academic Press.

Stalnaker, Robert. 2002. Common Ground. Linguistics and Philosophy, 25(5–6): 701–21.

Starr, William B. 2010. Conditionals, Meaning and Mood. PhD Dissertation, Rutgers University.

Starr, William B. MS. Imperative Dynamics. Manuscript, Cornell University.

Stojnic, Una. to appear. One's Modus Ponens: Modality, Coherence and Logic. *Philosophy and Phenomenological Research*.

Suszczyńska, Małgorzata. 1999. Apologizing in English, Polish and Hungarian: Different Languages, Different Strategies. Journal of Pragmatics, 31(8): 1053–65.

Thomason, Richmond H. 1990. Accommodation, Meaning and Implicature. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack (eds), Intentions in Communication, 325–63. Cambridge, MA: MIT Press.

Wierzbicka, Anna. 1985. Different Cultures, Different Languages,Different Speech Acts: Polish vs. English. Journal of Pragmatics, 9(2–3): 145–78.