Indirect Speech Acts

Nicholas Asher
Department of Philosophy,
University of Texas at Austin
nasher@bertie.la.utexas.edu

Alex Lascarides
Division of Informatics,
University of Edinburgh
alex@cogsci.ed.ac.uk

May 26, 2006

Abstract

In this paper, we address several puzzles concerning speech acts, particularly indirect speech acts. We show how a formal semantic theory of discourse interpretation can be used to define speech acts and to avoid murky issues concerning the metaphysics of action. We provide a formally precise definition of indirect speech acts, including the subclass of so-called conventionalized indirect speech acts. This analysis draws heavily on parallels between phenomena at the speech act level and the lexical level. First, we argue that, just as co-predication shows that some words can behave linguistically as if they're 'simultaneously' of incompatible semantic types, certain speech acts behave this way too. Secondly, as Horn and Bayer (1984) and others have suggested, both the lexicon and speech acts are subject to a principle of blocking or "preemption by synonymy": Conventionalised indirect speech acts can block their 'paraphrases' from being interpreted as indirect speech acts, even if this interpretation is calculable from Gricean-style principles. We provide a formal model of this blocking, and compare it with existing accounts of lexical blocking.

1 Introduction

Understanding the motives behind utterances is often crucial to successful communication. But the relationship between the surface form of an utterance and its underlying purpose isn't always straightforward, as Searle (1975) shows:

(1) Can you pass the salt?

Sentence (1) is an interrogative and so expresses a question. Usually, the speaker's goal in asking a question is to get an answer. But (1) plausibly has a different purpose: it's a request, where the speaker's goal is for the interpreter to pass the salt. This is an *indirect speech act* (ISA), which Searle defines to be an utterance in which one speech act is performed indirectly by performing another. With (1), requesting the hearer to pass the salt is performed indirectly by performing another communicative act—asking about the hearer's ability to pass the salt.

ISAs are puzzling in several respects. First what exactly does it mean for one action to be performed *by* performing another? Are there in fact two acts here? One act under several descriptions? Or one act with several distinct purposes?

Secondly, in as much as the semantic type of an utterance determines its linguistic behaviour, the semantic types of some ISAs are puzzling. On the one hand, there's evidence that some ISAs, such as (1), behave *linguistically* as if they're requests: e.g., like other requests, they can be modified with *please* where simple questions can't (Sadock 1974, Horn 1989, Lakoff 1973).

(2) a. Could you please pass the salt?
 b. Please pass the salt.
 c. ??Where did you please put the broom?

On the other hand, (1) also behaves linguistically like a question, in that a direct answer is a felicitous response (Bach and Harnish, 1979):

(3) a. A: Could you please pass the salt?
 B: Yes
 b. A: Please pass the salt.
 B: ??Yes.

Linguists consider questions and requests as not only distinct, but *incompatible* semantic types (e.g., Ginzburg and Sag, 2000). Questions denote a set of propositions (i.e., its direct answers; e.g., Groenendijk and Stokhof, 1991), whereas requests denote a relation between worlds and actions (e.g., Segerberg, 1990). As (1) has the linguistic hallmarks of both questions and requests, its semantic type must be distinct from, but related to, both of these; but it's unclear what its semantic type is. Many other ISAs also behave linguistically as if they're 'two speech acts in one', involving other kinds of incompatible semantic types. An adequate account of ISAs must model this.

Finally, the necessary and sufficient conditions for interpreting an utterance as an ISA are puzzling. On the one hand, there's evidence that ISAs arise from general principles of rationality and cooperativity. Gordon and Lakoff (1975), for instance, use Gricean maxims to analyse (1) and (4); Searle classifies (4b) as an ISA, since *B* rejects the proposal in (4a) by asserting something.[1]

(4) a. A: Let's go to the movies tonight.
 b. B: I have to study for an exam.

But paraphrases to some ISAs suggest that Gricean inferences aren't always sufficient. Paraphrases typically have the same calculable implicatures. But although replacing (4b) with a paraphrase leaves the communicative act unchanged, (5) shows this is not the case for (1):

(5) Do you have the physical ability to pass the salt?

[1] Principles of rationality and cooperativity are also used within the AI community to recognize speech acts (e.g., Grosz and Sidner 1990, Litman and Allen 1990).

Unlike (1), (5) is infelicitous as a request—at least in the 'null' context presented here—even though this interpretation is calculable.

Morgan (1978), Searle (1975) and others conclude from these observations that the calculable ISA for (1) is also conventionalized. Horn (1989) and Levinson (2000) go further, and argue that the difference between (1) and (5) is evidence that speech acts are subject to a principle of blocking or *preemption by synonymy*: conventionalised ISAs can block their paraphrases from being interpreted as ISAs, even if this interpretation is calculable. We agree and argue that this view is correct, but this raises puzzles anew: how can a precise model of the interaction between conventional and calculable information account for this blocking?

Our paper addresses these puzzles by giving a new interpretation of what a speech act *type* is that abstracts away from the language of actions. We use a formal model of discourse interpretation (Asher 1993, Lascarides and Asher 1993) in which rhetorical structure and lexical and compositional semantics interact to determine truth conditions and for which we have developed or adapted various technical tools that turn out to be useful here: nonmonotonic or defeasible reasoning, semantic types and type shifting rules (Montague 1974, Sag and Wasow 1999, Pustejovsky 1995, Asher and Pustejovsky 2000). Many speech act types turn out from this perspective to be rhetorical relations.

We will solve the "two speech acts in one" puzzle of ISAs by analyzing ISAs as a particular sort of complex semantic type that's formed from two incompatible types. The constituent types of the complex reflect the dual communicative role of an ISA like (1). When this dual communicative role is conventionalized (e.g., (1)), an extended notion of grammar assigns the ISA a complex semantic type; this then serves to explain why the ISA bears the *linguistic* hallmarks of both questions and requests. Other ISAs get assigned a complex type 'on the fly' rather than by convention, as a result of pragmatic reasoning and the demands of discourse interpretation. Roughly, an utterance u is conventionally assigned one semantic type, but a second semantic type is accommodated in the sense of Lewis (1969), in order to preserve assumptions that the discourse is 'well-formed' or coherent, and that the speaker was rational.

We'll use our formal tools to analyze blocking at the speech act level (e.g., that (1) blocks (5) from being interpreted as an ISA, even though this is calculable). Many researchers argue that blocking is a general principle of lexical organisation too, in morphology (e.g., Evans and Gazdar, 1989) and semantics (e.g., Bauer, 1983). Horn and Bayer (1984) use neg raising verbs as evidence for blocking in sentential semantics as well. At all these levels, blocking isn't absolute: a linguistic form can be assigned an interpretation that's normally blocked in sufficiently rich contexts. We'll offer precise conditions for both blocking and unblocking.

2 Defining ISAs via a Formal Theory of Interpretation

We begin by examining Searle's definition of speech acts, and refining it by linking speech act types to rhetorical relations. We'll then introduce the theory of discourse interpretation that forms the basis for formalizing these ideas, including the analysis of ISAs as complex types.

2.1 Searle's Speech Acts and ISAs

Searle argues that an ISA carries "two illocutionary forces" (the illocutionary force is the meaning the speaker intended to convey in performing the illocutionary act); he also says that "one illocutionary act is performed indirectly by means of performing another." Illocutionary acts (Austin, 1962) include things like informing, promising, asking, ordering, warning etc; and they are realised in performing locutionary acts (i.e., making utterances).

To understand how an utterance might have two illocutionary forces, one must understand how it has an illocutionary force at all. The connections between an utterance and its illocutionary force are typically a matter of linguistic convention (Searle, 1969). For instance, they're encoded within sentence mood: interrogatives (e.g., (6)) express questions; declaratives (e.g, (7)) express assertions; and imperatives (e.g., (8)) express requests. For Searle and others, illocutionary forces are also encoded in the lexicon: sentences like those in (9) express warnings, promises etc, as determined by the performative verb.

(6) Is your name Anakin?

(7) Your name is Anakin.

(8) Avoid the dark side of the force!

(9) a. I warn you that Anakin will turn to the dark side of the force.
 b. I promise I'll train Anakin to be a Jedi knight.

ISAs arise once one has individuated the different illocutionary acts in a taxonomy, and established a theory which aligns linguistic forms to these acts, as we did for (6–9). An utterance exemplifies an ISA if the illocutionary act that's predicted by the linguistic form, according to the theory of alignment, is a distinct act in the taxonomy from a further illocutionary act that has been performed. For example, (1) is an interrogative, and so like (6), the theory of alignment predicts the illocutionary act of asking a question. But (1) is also a request, and this is distinct in the taxonomy from asking a question.

The above rules for linking speech acts to linguistic form, however, look as though they accord two illocutionary acts to utterances (9): they are assertions in virtue of their mood, but also warnings etc. in virtue of the verb. So if one distinguishes warnings from assertions, as Searle (1976) does, then (9a) is plausibly an ISA, as Bach and Harnish (1992) argue.[2]

We take Searle's remark about two distinct illocutionary forces seriously. While Morgan argues that (1) is really a request and Bach and Harnish argue that it is a question, we think that all they show is that (1) is neither simply a question nor simply a request. We think it's both, in spite of the fact that as semantic types go, these are distinct. In fact, there's evidence that these types are incompatible. First, the linguistic behaviours of questions and of requests are incompatible (e.g., (2), (3)). Secondly, by the compositional semantics of questions (e.g. Groenendijk and Stokhof 1991, Asher and Lascarides 1998a) and of requests (e.g., Segerberg 1990, Hintikka 1983), they denote distinct types of semantic objects. And finally, these types

[2]Sentences (1) and (9) have some distinct properties: the additional illocutionary force of (1) can be 'cancelled' in some contexts; while (9) cannot fail to have a performative side when uttered sincerely.

cannot undergo 'mixed' quantification; contrast the acceptable quantifications over requests in (10a) and over questions in (10b), with the unacceptable quantification over both requests and questions in (10c):

(10) a. It is difficult to do everything the King demands.
b. Everything the professor will ask will be difficult to answer.
c. ??Everything the Queen asks that is difficult to answer the King will demand.

How can an utterance be assigned incompatible types? Of course, one way is through ambiguity. But (1) isn't ambiguous. It's not that the illocutionary force of (1) is indeterminate; rather, it's overdetermined and has *both* illocutionary forces, as Morgan's (1975) and Bach and Harnish's (1979) observations attest.

We have just used model-theoretic semantics and linguistic behaviour as the basis for distinguishing the speech act types of questioning and requesting. This contrasts with Searle (1976) and the more recent work on speech acts (e.g., Vanderveken 1990, Grosz and Sidner 1990), who use illocutionary points and direction of fit—i.e., the speaker's intentions—to distinguish among illocutionary acts. A corollary of Searle's view is that the performance of an ISA rests on certain facts holding of the speaker's cognitive state. But it's unclear, given this dependence of speech acts on cognitive states, whether there are one or two acts. For to have two intentions is logically equivalent to having one complex intention. More formally, where $\mathcal{I}_A\phi$ means that A intends ϕ, the following is an axiom:

$$\mathcal{I}_A(\phi \wedge \psi) \leftrightarrow (\mathcal{I}_A\phi \wedge \mathcal{I}_A\psi)$$

Thus intentions on their own fail to determine whether one (complex) act has been performed, or whether two acts have been performed (at the same time),[3] or whether we have two different descriptions of the same act (one in terms of two intentions and one in terms of one complex intention).

Speech act theory traditionally thus leads to murky issues in the ontology of action, with probably no clear resolution on the basis of linguistic data. Without denying the importance of Austin's insight that people do things with words, we don't see any easy route through that thicket. So we will try to avoid it altogether. For us, intentions are important to speech acts, but not the sole or even primary basis for distinguishing various speech act types. By using semantic typing and other model-theoretic terminology, we aim to account for ISAs while abstracting away from the attendant psychologism that threatens to be part of a detailed discussion of the metaphysics of action.

2.2 Speech Acts and Rhetorical Functions

Many types of speech acts must be understood relationally, because successfully performing them is logically dependent on the content of an antecedent utterance. So technically speaking,

[3]We of course believe that an agent can do two things at once. Looking at actions about which one has stable intuitions, it's clear that one can listen to a pianist play a sonata and watch her play at the same time and that these are distinct actions. On the other hand, if the only thing distinguishing two actions are the intentions behind them, it's hard to see how the intention that ϕ and the intention that ψ, when linked to the same utterance, can in and of themselves define two distinct acts.

the type must be (at least) a two place relation. For example, if one uses an utterance to conclude something, then that conclusion must be relative to some antecedent hypothesis or argument. Answering is also inherently relational: an answer is an answer to some prior question. Searle's example (4) also requires relational speech acts. Understanding that (4b) is a rejection is essential to understanding the content of the dialogue and why it's coherent. But a rejection is a *relation* between the utterance and some antecedent proposal (in this case, (4a)), rather than a property of the utterance (4b) itself, because successful performance of this speech act is logically dependent on this prior contribution.[4]

These relational speech acts mirror aspects of the rhetorical function of an utterance in the context. Since we use rhetorical relations to capture rhetorical function, types of relational speech acts are in fact rhetorical relations. Thus, a theory of rhetorical relations can help us understand different types of speech acts.

The rhetorical relations, which are used in Segmented Discourse Representation Theory (SDRT; Asher 1993, Lascarides and Asher 1993) and other theories (e.g., Mann and Thompson 1987, Hobbs 1985) to model various discourse phenomena, constitute distinct types of illocutionary force. Explanations, elaborations, giving backgrounds or describing results are all things that speakers *do* with utterances. Moreover, in rhetorically-based theories of discourse, these illocutionary contributions are all defined via not only an individual utterance, but also an antecedent utterance in the context.

To illustrate these ideas, let's consider the discourses (11) to (15):

(11) Max fell. John pushed him.

(12) A: Who's coming to the party?
B: John and Mary.

(13) A: Did John fail his exams?
B: He got 60%.

(14) A: John failed his exams.
B: No he didn't, he got 60%.
A: I meant John Smith.

(15) A: Let's meet next weekend.
B: How about Saturday?

The illocutionary contribution of an utterance to the overall discourse is *anaphorically dependent* (or in the case of initial utterances in a conversation, cataphorically dependent) on some antecedent contribution. In (11), the speaker doesn't simply assert that John pushed Max, he *explains* the antecedent that John fell. Providing this explanation must be an intentional act by the speaker; if it were not, then one cannot understand why he juxtaposed the sentences, or why the discourse is coherent. This explanation is relational: explanations explain some prior contribution, and the constraints on successfully providing an explanation are dependent on the content of that prior contribution. Dialogue (12) is similar to (11). It

[4]This relational view is also implicit in the definitions of some speech acts in Carberry and Lambert (1999); e.g., the act of expressing doubt.

features a question-answer pattern, where the act of answering must be viewed as relational, since the successful performance of answering is dependent on the semantics of the question. Dialogue (13) is similar to (12), except the utterance isn't a *direct* answer (which would have to be *yes* or *no*, given the form of the question), but rather, proviso certain assumptions about the participants' cognitive states, it provides sufficient information that the questioner can compute a direct answer to his question, a relation that we represent as *IQAP* in SDRT, standing for *indirect question answer pair* (Asher and Lascarides, 1998a). Dialogue (14) is an example where A and B *correct* perceived misunderstandings, and corrections must be relational for the same reasons as answers are.

In dialogue (15), B's utterance being an interrogative is a clue that his goal is to know an answer to the question. But in fact, B's intentions are slightly more specific than this. He utters this question with the intention of gaining information that will help achieve the goal behind A's utterance: namely, knowing a time in next weekend when they can meet. In fact, recognizing this contribution of B's utterance is an important part of what makes the dialogue coherent. Now, in SDRT, we call this particular speech act *Q-Elab*, standing for Question Elaboration (Asher, 1999), and its semantics is defined (in contrast to *Explanation*, for example) in terms of goals, as well as content. First, we define a *speech act related goal* (SARG) to be a goal that is either conventionally associated with a particular type of utterance or is recoverable by the interpreter from the discourse context (Asher and Lascarides, 1998a); this distinguishes the goals that interact with linguistic knowledge from goals in general. Then, *Q-Elab*(α, β) holds only if β is a question and all possible answers to β (according to the compositional semantics of questions and answers) elaborate a plan to achieve the SARG that prompted the utterance of α. In all of these cases, the illocutionary act relates the utterance to an antecedent utterance.

So, the rhetorical relations in SDRT are all *speech act types*: the second term of the relation is a speech act of the appropriate type relative to its discourse context, which of course includes the first term in the relation. For example, in (11), the speech act or content expressed by the second utterance in this context stands in an *Explanation* relation to the speech act expressed by the first utterance. We write this as *Explanation*(α, β), where α and β label the results of semantic processing of the utterances—i.e., with them are associated logical forms for bits of discourse with particular semantic values. We'll also assume that each such label has not only an associated logical form, but also a speaker, *Agent*(α). This means that by saying β, *Agent*(β) performs the speech act of providing an *Explanation* relative to the discourse context α. This has truth conditional effects: the second clause doesn't merely convey that John pushed Max; rather, since it explains why Max fell, it also conveys the *cause* of Max falling.

By defining a theory of speech acts in terms of a theory of rhetorical relations, we gain a means for addressing problems we raised earlier. First, SDRT is a formal semantic theory, and so it distinguishes the contributions made by utterances at the level of *semantic value*, although intentions/SARGs can also play a role. For instance, questions are distinguished from requests because they are distinct types of semantic objects (sets of propositions vs. relations between worlds and actions); and for each of these, there's a distinct set of possible rhetorical relations to which these semantic objects can be arguments. Thus, formal semantics makes precise Searle's ideas about direction of fit, which he uses to define speech acts. Second, we get an account of the relational nature of many speech acts in virtue of their analysis as rhetorical relations. Thirdly, formal semantics will provide the basis for analyzing ISAs as complex semantic objects, while bypassing the need to make assumptions about how many actions were performed, as we'll shortly see.

Rhetorical relations yield a finer grained typology of speech acts than is traditionally construed. For example, the speech act of asserting in traditional speech act theory gets divided into several (sub)types of speech act in the SDRT theory: explanations, narrations, corrections and so on (Asher and Lascarides, 1998a). As a result, it provides richer mechanisms for computing the implicatures that follow from speech acts (e.g., the causal relation in (11)), while often bypassing complex reasoning about intentions; we'll see this in section 3.1.

The typology is also more robust, because the various speech act types are individuated on the basis of their truth conditional effects, rather than just intentions and other propositional attitudes. A rhetorical relation—or equivalently, a speech act type—is included in SDRT only if it has empirical consequences within truth conditional semantics. That is, the rhetorical relation must affect some aspect of the context change potential of an utterance within dynamic semantics, where this effect can't be explained straightforwardly by other means. For example, it might impose constraints on antecedents to anaphora that can't be expressed otherwise. On this basis, a *Correction* like that in (14) is distinguished from an *Explanation* like that in (11), because *Correction*(α, β) doesn't entail that the proposition labelled by α is true, while *Explanation*(α, β) does, making their contributions to discourse interpretation different. But neither *Correction* nor *Explanation* are distinguished from assertions. In this way, we have abstracted away from the need to individuate actions solely on the basis of intentions and other attitudes.

Finally, SDRT provides a detailed and formally precise theory of alignment between speech acts and linguistic form. While this is recognized as an important part of speech act theory (e.g., Searle 1969, Bach and Harnish 1979, Vanderveken 1990), the logical tools of SDRT lead to a better theory of alignment than these theories, by capturing the *defeasibility* of the link between illocutionary force and linguistic form (Hobbs *et al.* 1993, Lascarides and Asher 1993, Perrault 1990). For instance, uttering a sentence with indicative mood isn't an assertion in some contexts. In SDRT, rhetorical relations are inferred from linguistic form via axioms in a default logic. We'll say more about this in section 3.1.

2.3 ISAs as Complex Types

With this view of speech act types as rhetorical relations, let's return to ISAs. As we've seen, there's evidence that ISAs behave as if they are assigned incompatible semantic types. We observed this with (1), where observations about *please* and direct answers to the question suggested it behaves linguistically like both a request and a question. Moreover, no prior knowledge of the speaker's intentions are needed for interpreting (1): in a typical context it's interpreted as both a question and a request, and in the absence of information to the contrary, the requisite speaker intentions are assumed. For us, it is a matter of *interpretation* that utterances give rise to ISAs. Thus, we can abstract away from talk of actions and simply speak of the communicative or semantic value of an utterance.

But how should one represent the 'overdetermined' semantic value of an utterance like (1), typing it as both a question and a request? We will resort to a device that is used within the lexicon to represent an item that has an overdetermined semantic type; namely, dot types. That is, sentence (1) will be assigned a *dot type* of semantic object, with a question and a request as constituents of this dot type. The logic that computes the rhetorical connections between utterances has access to the fact that the semantic value of (1) is a dot type, and so it can exploit either its question value or its request value.

To get a better feel for dot types, let's look at their lexical use. Lexical and compositional semantics has focused lately on underspecified semantic values and types (e.g., Copestake and Briscoe 1995, Fernando 1997), but overdetermined semantic values also appear to have a role. For instance, Pustejovsky (1995) uses dot types to analyze copredication, where coordinated VPs select for arguments of incompatible semantic types. Consider the word *book*. The type of object that this word denotes can be understood from several perspectives: as a physical object (e.g, (16a)); or as an abstract object of information content (e.g., (16b)). And indeed, it can be understood as both at the same time, as in (16c), which is an example of copredication.

(16) a. The book has a purple cover.

b. The book is hard to understand.

c. The book is 500 pages long and documents in detail the theories of Freud.

Physical objects and abstract objects are typically incompatible types in lexical hierarchies. And yet, a book partakes of both, though this combination is not to be understood as the *meet* of the two types in a type lattice. Rather, it's to be understood at a 'meta' level—a combination of two possible *perspectives*.

Formally, we assume an operation on types which takes two distinct types and combines them into a complex type: if t_1 and t_2 are types, then so is $t_1 \bullet t_2$ (even if the meet of t_1 and t_2 isn't defined). This complex type $t_1 \bullet t_2$ can then be exploited in various ways in the composition of an utterance meaning. Pustejovsky (1995) demonstrates that by assigning the lexical entry *book* the complex type $phys\_object \bullet abst\_object$, semantic composition in the grammar can exploit this typing to predict that predicates in the logical form can modify *phys_object*, as in (16a), *abst_object*, as in (16b), or both constituent types, as in (16c). The formal properties of these dot types are explained in Asher and Pustejovsky (forthcoming). Their main property for our purposes is that if a term v is typed as $t_1 \bullet t_2$ and v is predicated over by a predicate P that selects an argument of type t_1 (or t_2), then we may introduce a new term w of type t_1 (t_2) that's linked to v, and P predicates over w. The rule that encapsulates this is known as `Dot Exploitation` (we'll specify this rule formally in Section 3.1). And unlike other rules for manipulating semantic types, it is *ampliative* not *destructive*: the complex type remains in context for further predication. On this basis, there are many lexical items that benefit from a dot type analysis, including nominalizations (Pustejovsky, 1995) and causative verbs (Asher and Pustejovsky, forthcoming).

We will use dot types to define ISAs. Consider for instance conventionalized ISAs like (1), that are 'combinations' of distinct—and as we've argued, incompatible—semantic objects (in this case, questions and requests). Just as as (16c) involves *simultaneous* modification of both the physical and the informational types of *book*, the following dialogue exchange simultaneously attends to both the question and request types of (1) (via *please* on the one hand, and the direct answer on the other):

(17) a. A: Can you please pass the salt?
b. B: Yes [(*uttered as B passes the salt*)]

To explain how these two predications in (17) on the same speech act object can nevertheless have conflicting type requirements, we assume that the grammar assigns (1) the dot type *question•request*. The rule `Dot Exploitation`, which we describe shortly, will then introduce a label of type *request* that is the argument of *please*, and that is linked with *O-Elab* (standing for Dot Elaboration) to the original speech act of complex type. Another application of `Dot Exploitation` will generate the label of type *question* to which B's verbal response (17b) is connected, and this is also linked with *O-Elab* to the original speech act.[5] So for us, the hallmark of conventionalized ISAs is that the grammar assigns to them a dot type involving incompatible constituent types, which allow us to explain their 'dual' linguistic behaviours.

While in principle any combination of independent types can define a dot type, natural language does not seem to function that way. There always seems to be some natural relationship between the constituent types. Consider the book: its physical and abstract perspectives are naturally related in that the physical is a realization of the abstract information content—a relation familiar to metaphysics since the time of Plato. With regard to ISAs, the natural connection between the constituent types in a dot type comes from the Gricean reasoning that Searle remarked on. Reasoning about the speaker's goals from assumed mutual beliefs yields the coherent link between the question and the request in (1), for example.

Gricean reasoning can also make a question into an answer to another question (e.g., (18b)) and a question into a negative commentary—which we refer to as a *correction*—of some prior assertion (e.g., (19b)):

(18) a. A: Do you want another piece of chocolate cake?
b. B: Is the pope Catholic?

(19) a. A: Reagan was the best US president of the 20th century.
b. B (to C): Is he really such an idiot as to believe that?
c. C: Yes, he is.
c′. A: Well, maybe you're right. Maybe Reagan was mediocre.

C's response (19c) to B's utterance (19b) indicates that (19b) still functions as a question, in that C can provide a direct answer to it. The alternative 'continuation' of the dialogue (19ab) given in (19c′) indicates that (19b) also functions as a commentary by B on A's view. In this sense, the utterance (19b) contributes to the meaning of the dialogue as both a question and as a correction, even though these types of speech act are incompatible (since corrections are a kind of assertion, which are of an incompatible semantic type to questions).

The hallmark of all ISAs, then, seems to be that Gricean reasoning provides a connection between one kind of speech act and an incompatible kind of speech act. For conventionalized ISAs like (1), Gricean reasoning links the incompatible constituent types of the dot type that's

[5] In SDRT we assume that to attach new information to the discourse context, we must attach it to a particular, available label (Asher, 1993). Here we'll assume that *O-Elab*, like other elaborations, makes both labels it relates available for further links.

assigned to the utterance by the grammar. For what we'll call *unconventionalized* ISAs (e.g., (19b)), where the grammar doesn't assign the utterance a dot type, Gricean reasoning connects the speech act that's assigned by the grammar (e.g., *question* in the case of (19b)) to the 'implict', incompatible type of speech act (*Correction*), which is essentially 'accommodated' into the interpretation because of demands on content that are imposed by factors such as discourse coherence and speaker rationality.[6]

In this respect, (4b) appears different: Gricean reasoning connects the act of assertion to the act of rejection, but there isn't any *semantic* incompatibility between an assertion and a rejection in the way that there is between questions and assertions, for example. So (4b) doesn't involve two incompatible types of speech acts. Perhaps one could make the case that there is some other incompatibility, but we know of no such argument in the literature.

This discussion thus yields an alternative definition of ISAs to those given by speech act theory (e.g. Searle, Bach and Harnish, and the like). An utterance is a conventionalized ISA if (a) the grammar assigns it a complex speech act type of the form $s_1 \bullet s_2$, such that s_1 and s_2 are distinct (incompatible) types of semantic objects; and (b) Gricean-style principles of rationality and cooperativity link the constituent type s_1 to the type s_2. This Gricean link means there's a sense in which the complex type $s_1 \bullet s_2$ is asymmetrical, characterized by the Gricean information flow from s_1 to s_2.[7] For example, (1) is a conventionalized ISA because (a) there's linguistic evidence (e.g., (17)) that the grammar assigns (1) the complex dot type *question•request*, and (b) as Gordon and Lakoff (1975) show, Gricean-style reasoning links the question to the request. We'll shortly provide evidence that (18b) is also a conventionalized ISA, with dot type *question • IQAP*.

An utterance is an unconventionalized ISA, if similar Gricean style reasoning leads to the inference of an implicit speech act—we'll make precise the notion of *implicit* shortly, but for now it can be taken to mean that this speech act isn't derivable from the theory of alignment—and this implicit type is semantically incompatible with that inferred for the utterance itself (by the theory of alignment). So (19b) is an unconventionalized ISA, but (4b) is not. This makes unconventionalized ISAs a special case of a general pragmatic process during interpretation. Generally, interpreters accommodate content in order to preserve assumptions that the speaker was rational and that the discourse is semantically well-formed (Lewis, 1969). An unconventionalized ISA arises when what's accommodated into the semantic representation of the discourse is a speech act s_2 that's distinct from the speech act s_1 that's predicted by the grammar. We'll say more about such semantic representations shortly.[8]

In appropriate contexts, Grice's example (20) is interpreted as giving rise to a request.

(20) I'm out of gas.

[6]Note that utterance (19b) doesn't behave *linguistically* like a correction, since it cannot be felicitously preceded by *no*, in the way that corrections can; compare *??No, is he such an idiot as to believe that?* with *No, Reagan was mediocre.*

[7]Copestake and Briscoe (1995) show, via data involving quantification, that the complex types assigned to lexical entries like *book* are also asymmetrical.

[8]Note that these defining criteria for ISAs distinguish them from idioms like *John kicked the bucket.* There's no evidence, analogous to the modifications of (16c) in (17), that such idioms should be assigned complex types. So the grammar provides distinct representations of the literal and idiomatic meanings (Nunberg *et al.* 1994, Copestake and Briscoe 1995), and these are both assigned 'simple' semantic objects.

While (20) is not grammatically functioning as a request as well as an assertion (e.g., it can't be modified with *please*), its discourse function connects it to an implicit request: the proposition expressed by the utterance in (20) explains *why* the speaker needs help and what sort of help he needs. An interpreter needs to recognize this in order to respond in a competent way (i.e., help the speaker obtain gas). And in order to recognize that (20) explains the request for help, he must reason about the agent's cognitive state (Asher, 1999). This Gricean reasoning therefore links two incompatible types just as it does for (1), except the grammar doesn't assign (20) the dot object. So we would classify (20) as an unconventionalized ISA.

The second sentence of (11) isn't an ISA, however: although (11b) conveys a causal relation between the pushing and the falling that isn't implied by the compositional semantics, assertions are a supertype of explanations and so no implicit object of an incompatible semantic type is inferred. Furthermore, in some cases, an utterance can have several rhetorical functions; for example, consider (21):

(21) a. We bought the house.
b. But then we rented it out.

Given the cues *but* and *then*, a competent language user must infer that the relations *Contrast* and *Narration* connect (21b) to (21a). These speech acts trigger further inferences about intended content: the renting out occurs after the buying (from *Narration*) and renting out is construed as cancelling an expectation that would be inferred from the buying (from *Contrast*, see Asher 1993, Asher, Busquets and Hardt 1997). But *Narration* and *Contrast* are compatible types; they're both assertions, and so they both correspond semantically to a proposition. So (21b) is not an ISA: it's neither assigned a complex dot type by the grammar, nor do we infer via Gricean reasoning an implicit object of an incompatible type. Rather, the interpretation of (21b) involves two independent and compatible simple speech act types.

2.4 Conventionalization

The link between the two illocutionary forces in an ISA like (1) is calculable. Indeed, we've claimed that this is the basis for the coherence of the dot type that's assigned to (1) by the grammar. Furthermore, these calculable inferences, which apply to (1) and to (18b), apply also to their near-synonyms and paraphrases. But contrary to these calculable predictions, the paraphrases don't always have the same illocutionary force. The fact that B's responses in (22) are odd as ISAs shows that a question with a mutually known positive answer is insufficient (though necessary) for the ISA interpretation of (18):

(22) a. A: Do you want some chocolate cake?
b. B: ??Is it really the case that the Pope is Catholic?
b′. B: ??Am I self-identical?
b″. B: ??Does two times three equal six?

These differences between (1) vs. (5) and (18) vs. (22) are evidence that calculable implicatures, though necessary, aren't always sufficient for an ISA interpretation.

In line with earlier research (e.g., Horn, 1989), we claim that *conventionalization* and *blocking* or preemption by synonymy underlies the differences in the interpretations of these 'paraphrases'. The argument in favour of this claim rests on independent evidence that some ISAs are conventionalized. Once one accepts conventional ISAs, the differences between (1) and (5) and between (18) and (22) show that blocking must guide inferences about communicative goals.

Bach and Harnish (1979, 1992, 1995) suggest that conventionalization isn't the right view of ISAs like (1), arguing against Searle's idea of illocutionary conventions. Their arguments against Searle aren't relevant to us, since our conventions are not illocutionary; they're at the level of semantic types or logical form, though pragmatics affects what those types are.[9] Bach and Harnish propose instead a notion of *standardization* to account for what we call conventionalized ISAs, which they define in terms of mutual beliefs about what the illocutionary intent is of an utterance. We think this leads to odd consequences. For instance, they claim that (1) is literally a question but "insincerely" meant and thus mutually believed to have another illocutionary force. We find this psychological judgement odd, and agree with Searle that at least some ISAs like *Can you reach the suitcase?* can be sincerely meant as a question and also as a request. Furthermore, Bach and Harnish are forced to reject sentences like (1b), in which the question cooccurs with *please*, as ungrammatical. Our account doesn't force us to this conclusion, which is all to the good, since we (and most English speakers) judge such examples to be perfectly acceptable, and hence—according to the usual empirical linguistic standards—grammatical. For us, linguistic data like (17) is evidence that (1) behaves linguistically like a request and a question.

We now turn to some more evidence that some ISAs are conventionalized dot objects—evidence that we find hard to account for with Bach and Harnish's standardization thesis. There is in fact a sharp difference between the behavior of (1) or (23a) and, for instance, (23c). (23c) can be used to make a request, but it cannot be modified with *please* and so doesn't appear to function like a request conventionally (e.g., (23d)) (Sadock 1974, Lakoff 1973).[10]

(23) a. Can you shut the door?
b. Can you shut the door please?
c. I can't reach the door.
d. ???I can't reach the door please.

Furthermore, those speech acts that can be modified with *please* can also be coherently followed by propositions which *explain* why the ISA was performed (e.g., (24a)); and those that can't, can't (e.g., (24b)) (see Davison (1975) and Horn and Bayer (1984) for discussion).

[9] In this we agree with Levinson (2000), who gives persuasive evidence that pragmatics can intrude on the construction of logical form.

[10] Nevertheless, we think that it can be a mutual belief that (23c), like (23d), is typically used to make a request. Thus, Bach and Harnish can't account for this difference in grammatical behaviour.

(24) a. $\left.\begin{array}{l}\text{Shut the door}\\ \text{Can you shut the door}\\ \text{Could you shut the door}\\ \text{Would you shut the door}\\ \text{I would like you to shut the door}\\ \text{I was wondering if you could shut the door}\end{array}\right\}$ because it's freezing in here

b. I can't reach the door, ???because it's freezing in here

Similarly, the ISAs in (24a) can be preceded with the conditional *If you would be so kind...*, while (24b) can't. Indeed, *If you would be so kind, I would like you to shut the door* doesn't convey a (logical) dependency between your kindness and my desire, but rather it's interpreted as a conditional request, indicating that the consequent clause is acting semantically like a request.

Lakoff (1973) also observes that the implicatures of certain forms affect their linguistic behaviour, providing further evidence for some degree of conventionalisation (this view is also supported in Levinson (2000), who advocates pragmatic intrusion into syntactic and semantic analysis). For example, he observes the sentence amalgam in (25a), where a whole phrase which implicates a modifier (something like the adjective *unusual*) can felicitously slot into the position where the modifier would be, in contrast to other phrases which don't carry this impilcature (e.g., (25b)):

(25) a. John invited you'll never guess how many people to the party.

b. ??John invited Bill will never guess how many people to the party.

Gordon and Lakoff (1975) also discuss the phonological and morphosyntactic effects of the 'conventionalized' indirect suggestion in (26a) which doesn't apply when the suggestion isn't implicated (26b):

(26) a. Why don't you move to California?
Why dontcha move to California?
Whyntcha move to California?

b. Why don't you resemble your father?
??Why dontcha resemble your father?
??Whyntcha resemble your father?

Similarly, Brown and Levinson (1987) observe that honorifics don't collocate with ISAs that implicate disrespect, even if the linguistic form would allow such a collocation.

The second source of evidence that some ISAs are conventionalized rests on simple intuition. Morgan (1975) notes that although the implicature that (1) is a request is calculable, it's counterintuitive to assume that people actually calculate this implicature 'on the fly' from Gricean maxims. People don't need to do this, because it's part of a competent user's knowledge of the English *language* and its use that it should be interpreted this way. Morgan refers to this as implicatures being *short circuited*. The hypothesis is that synchronically, the connection between (1) and the request is knowledge of language (or, perhaps more accurately, language *use*) rather than of conversational implicature, although diachronically, the implicatures were no doubt involved in the process of conventionalization. So for us, dot types reflect short circuited implicatures: in $s_1 \bullet s_2$, s_1 is the input and s_2 the output to some calculable logical consequence relation, but $s_1 \bullet s_2$ on its own doesn't reveal the *details* of the (calculable) proof.

Finally, cross-linguistic variation provides evidence of conventionalization (Cole, 1975). For example, Cole argues that diachronically, the conversational implicatures that *let's X* conveys a request have, by convention, been assimilated to the literal meaning of *let's*. But although this process of lexicalization has happened in English, it hasn't in Hebrew, for example. Searle (1969) himself notes similar variations for English and Czech.

Given the linguistic evidence, we think that ISAs like (1) and (18b) are conventionalized and so are assigned dot types by the grammar. The difference in interpretation between (1) vs. (5) and (18b) vs. (22) indicate that a principle of blocking applies: if one speech form is conventionally associated with a particular ISA, then speech forms with similar meanings are blocked (by default) from being interpreted this way, even if this ISA interpretation is calculable from Gricean-style maxims. This is analogous to lexical blocking: if the lexicon includes a word form w_1 with a particular interpretation s, then a different word form w_2 is (by default) blocked from being interpreted as having this sense s, even if this interpretation is derivable from some lexical generalization. So, *pig* is blocked from denoting the edible substance because of *pork* (e.g., *I ate pig* is highly marked), even though, following Copestake and Briscoe (1995) and others, there's a default generalization that animal nouns can be interpreted this way (e.g., *haddock*, *salmon*, *chicken*, *turkey*, and the less frequent *snake*, *mole*, *badger*).

Horn and Bayer (1984) argue that neg raising verbs are also subject to blocking, and if this is true, then blocking also occurs at the level of sentential semantics, since neg raising phenomena focus on the relative semantic scope of negation to other scope bearing elements in the sentence. For example, (27a) implicates the stronger (27b).

(27) a. I don't think Mary will come.
b. I think Mary won't come.

Lakoff (1973) and Horn and Bayer (1984) observe that neg raising interpretations are subject to semantically unmotivated lexical exceptions. For example, *suppose* is neg-raising on its paranthetical reading where the semantically similar *guess* is not:

(28) I don't suppose/??guess Max will arrive until midnight.

And *wish* neg raises where *desire* does so only with difficulty; similarly for *expect* (neg raising) vs. *anticipate* (not neg raising). This can be explained via a principle of blocking: if the implicatures that a verb has a neg raising interpretation are short-circuited/conventionalized, then this blocks its paraphrases from having this interpretation.

At the speech act level, blocking can explain why (1) is interpreted as an ISA where (5) is not, even though this interpretation is calculable. In a similar vein, Levinson (2000) argues that the implicatures that the utterance (29a) is an ISA (for it's an assertion that acts as an apology) is short-cicuited, because its paraphrase (29b) is blocked from having this ISA interpretation; similarly (30a) blocks (30b) from having an ISA interpretation as a greeting, and (31a) blocks its paraphrases (31bc) from having the ISA interpretation as requests:

(29) a. I'm sorry.
b. That saddens me.

(30) a. I'm delighted to meet you.
b. I'm gratified to meet you.

(31) a. I would like the door closed.
b. I would admire the door closed.
c. I would desire the door closed.

Note also that (29a) behaves *linguistically* like an apology, in that it can be followed by the response *apology accepted*. This contrasts with the utterance *I acted in an inappropriate manner*, where *apology accepted* is an awkward response, even if it's interpreted as an *isa* (i.e., as an apology).

A crucial part of our hypothesis is that blocking is default. This is for two reasons. First, there are cases where a paraphrase of a conventionalized ISA is itself conventionalized. For example, *Can you X?* is similar in (literal) meaning to *Could you X?*, but the ISAs are conventionalized in both cases (at least, for certain kinds of X), and so neither one blocks the other. Second, blocking isn't absolute: the context provided below 'unblocks' the ISA interpretation of (32a), and (33) is attested to convey a similar ISA to the conventionalized (18b), even though normally it's blocked by it:

(32) [Context: No one except perhaps for B can get to the meeting because they're not in town, and B is sick but in town]
a. A (addressed to B): Are you physically able to get to the meeting?
b. B: OK, I'll go.

(33) a. Do I like the XK8?
b. Do rabbits dig up my lawn?
(*Caledonia Magazine*, July 1999, in a review of the Jaguar XK8)

Analogously, at the lexical level, blocking is rarely absolute (see Briscoe *et al.*, 1996), making blocking a *default* principle of lexical organization that's overridable in sufficiently rich discourse contexts. For example, Briscoe and Copestake (1999) attest the following use of *cow* to mean meat, even though this interpretation is usually blocked by *beef*:

(34) In the case of at least one county primary school…they were offered (with perfect timing) sauté potatoes, carrots, runner beans and roast cow.
(*The Guardian* newspaper, May 16th 1990, in a story about BSE).

We provide below a formal model of ISAs and their blocking. We will first briefly describe a formal theory of alignment for speech act types such as *Explanation* and *Q-Elab*. We will then briefly describe how principles of cognitive reasoning can be integrated with a formal theory of discourse interpretation, so that we can predict calculable speech acts. And we will demonstrate that a formal model of speech act blocking requires a framework with a high degree of *modularity*: in particular, conventional information must be separate from, but interact with, information pertinent to the cognitive modelling of dialogue agents.

3 The Formalization within SDRT

The framework for our analysis is Segmented Discourse Representation Theory (SDRT, e.g., Asher 1993).[11] SDRT attempts to represent the content which any competent language user would take to be what the speaker intended to convey. So this content is both semantically and pragmatically determined. The discourse is represented as an SDRS, which is a recursive structure of labelled DRSs (Kamp and Reyle, 1993), with rhetorical relations (or in our new terminology, speech act types) like *Elaboration* and *Contrast* between the labels. In previous work, we have argued that an adequate logical form for discourse must feature these rhetorical relations for (at least) two reasons. First, they provide a natural definition of discourse coherence: a discourse is semantically and pragmatically coherent just in case in its logical form, every bit of information is rhetorically connected to some other bit of information. Secondly, rhetorical relations evoke truth conditional effects on the content of the discourse that can't be explained straightforwardly by other means (for a full discussion see Asher and Lascarides 1995, 1998a, 1998b, 1998c).

Since the logical forms feature rhetorical relations, building these logical forms dynamically involves computing a rhetorical connection between the new information and its discourse context. We review in Section 3.1 how this is done in SDRT.

3.1 The Theory of Alignment for Relational Speech Acts

Austin (1962) and Searle (1969) align sentence mood and other grammatical features to illocutionary acts and goals. For example, indicative sentences align with the act of asserting and the goal of belief transfer; interrogatives align with questioning and a goal of knowing an answer; and imperatives align with requesting and a goal that the action be performed.

In SDRT, the act of assertion is a 'supertype' to to 'relational' speech acts like *Explanation*, *Correction* and *Contrast*. Inferring these relational acts requires information about the prior utterance that the speech act is connected to; knowing that the current utterance is indicative is insufficient. But what information does one need about the prior utterance? AI approaches to dialogue interpretation rightly point out that the goals that are introduced in the dialogue context affect inferences about the speech act of the current utterance (e.g., Grosz and Sidner 1986, 1990, Litman and Allen, 1990). But they use *only* the goals of the antecedent utterance, rather than its compositional or lexical semantics directly, to constrain the recognition of the

[11] In principle one could use other formal theories of discourse interpretation (e.g. Hobbs et al. 1993). We believe that SDRT has a number of advantages and have argued for it at length elsewhere. We forego this here.

current speech act. We believe that this strategy won't always work.

There are at least two reasons why. First, the successful performance of the current speech act is often dependent on the logical structure of the antecedent utterances, and goals don't reflect this logical structure; rather compositional semantics does (following DRT, Kamp and Reyle 1993). Dialogue (35) demonstrates this:

(35) a. A: How about meeting next weekend?
b. B: That sounds good.
c. How about Saturday afternoon?
d. A: I'm afraid I'm busy then.
e. ??How about 2pm?

Given the context, a SARG of (35d) is to find a time to meet that's next weekend but not on Saturday afternoon. So computing (35e)'s speech act solely on the basis of the prior goals and the current linguistic form would predict that *2pm* successfully refers to 2pm on Sunday and the speech act *Q-Elab*(35d,35e) is performed. The fact that (35e) is odd indicates that recognizing (35e)'s speech act is constrained by something else. On our approach, the logical and rhetorical structure of (35a-d) plays a central role, for according to *linguistic* constraints defined within dynamic semantics (e.g., Kamp and Reyle 1993, Asher 1993), (35a-d) make Sunday inaccessible, thereby forcing *2pm* to denote 2pm on Saturday. And therefore, one cannot consistently connect (35e) to (35d) with *Q-Elab*, resulting in incoherence. This *direct* logical dependence between the logical form of the context and the speech act *Q-Elab* must be encoded in the theory of alignment.

Secondly, inferring the speech act *Explanation* for (11) doesn't so much depend on the communicative goal of the first sentence, but on its lexical semantics; in particular, the fact that *fall* is a movement verb, while *push* is a verb that describes a force that causes movement.

(11) Max fell. John pushed him.

It may be possible in principle to infer that the speaker performed an *Explanation* via reasoning about the SARG (presumably, of belief transfer) that underlies his utterance *Max fell*, rather than by reasoning with the semantics of *Max fell* directly. But inferring the current speech act this way would be cumbersome and inefficient; it would be computationally and conceptually more elegant to use linguistic content rather than goals whenever it is possible to do so.

In light of these observations, SDRT enriches the plan recognition approaches to interpretation, by allowing the linguistic form of a prior utterance to affect the recognition of the current speech act *directly* (see also Carberry and Lambert, 1999). Examples such as (11–15) show that, in line with Searle and Austin's views, the rhetorical connection often follows by *default* from the linguistic clues present, because monotonic clues (such as the cue phrase *because* in the case of (11)) are absent. SDRT axiomatizes these default links from linguistic form to (relational) speech acts via a collection of axioms known as DICE (Discourse in Commonsense Entailment; Lascarides and Asher 1993, Asher and Lascarides 1995, 1998a). We present only brief highlights of DICE here.

DICE is specified in a modal propositional logic with a conditional operator $>$ to represent defaults: $A > B$ means *If A then normally B*. Rules in DICE have the form (36):

(36) $(\langle \tau, \alpha, \beta \rangle \wedge \mathit{Info}(\tau, \beta)) > R(\alpha, \beta)$

$\langle \tau, \alpha, \beta \rangle$ means β is to be attached to α with a rhetorical relation (α and β label SDRSs) where α is part of the discourse context τ; $\mathit{Info}(\tau, \beta)$ is a gloss for information about the content that τ and β label (in particular, their compositional and lexical semantics); and R is a rhetorical relation. We make one innovation to DICE here: we will be explicit about semantic types. So we'll extend the DICE language with type declarations in the style of Asher and Pustejovsky (forthcoming). α :? means that α has question type; α :! that α is a request and $\alpha : |$ that α is an assertion. Some DICE axioms exploit these semantic types. And all predicates in DICE bring along with them a *type context*, in which the types for the arguments of the predicates are given. As we add or infer new formulas in DICE, we'll be extending and merging type contexts. For example, we'll assume that $\langle \tau, \alpha, \beta \rangle$ requires both α and β to be of simple types. An independent type checking mechanism like that used in Asher and Pustejovsky (forthcoming) will control for and resolve any type clashes in the type context, as new information is added or inferred. Thus, for example, when β is a dot type and we have $\langle \tau, \alpha, \beta \rangle$, we will get a type clash, at which point the rule of `Dot Exploitation` will be used within the DICE language to introduce a new label γ of a type that is a constituent of β's type, and the ill-typed $\langle \tau, \alpha, \beta \rangle$ will be replaced with the well-typed $\langle \tau, \alpha, \gamma \rangle$ and $\mathit{O\text{-}Elab}(\beta, \gamma)$.

For the record, the rule of `Dot Exploitation` is given below. `Dot Exploitation` deals with a situation where the term x has conflicting type restrictions imposed on it by the premises. The rule resolves this conflict by introducing a new term v, in DRT terms a new discourse referent, with a specific relation (*O-Elab*) to the original term x, where v is of the appropriate type.

- `Dot Exploitation`
  If ϕ and ψ both take a discourse referent x as an argument, where ϕ's types x as of type t_1 and ψ restricts x to type $t_1 \bullet t_2$, then a new discourse referent v of type t_1 is introduced, such that $\mathit{O\text{-}Elab}(x, v)$ and x is uniformly replaced by v in ϕ hold.

This rule exploits the "left" constituent t_1 of the dot. A similar rule exploits the right constituent t_2.[12]

Finally we'll assume axioms on types that allow us to recover the logical forms for labels which are introduced by the `Dot Exploitation` rule. In the case of (1) for example, the logical form (LF) of the (new) label that's of type *request* is constructed by adding an imperative operator to the LF for the VP of the interrogative. The LF of the question, on the other hand, is identical to that for the label with the complex type *question•request*, that's assigned to (1) by the grammar.

We turn now briefly to some examples of DICE rules. In Asher and Lascarides (1995) we provide an axiom of the form (36) which captures the following: if (a) $\langle \tau, \alpha, \beta \rangle$ holds; (b) according to lexical semantics (in particular of the verb in α), α describes a change along

[12] For the technical details, see Asher and Pustejovsky (forthcoming).

some dimension; and (c) lexical semantics also determines that β describes a cause for change along that same dimension; then normally (d) *Explanation*(α, β) holds. Clauses (b) and (c) reflect a particular value of *Info*(τ, β) and clause (d) a particular value for R in (36). The words *fall* and *push* satisfy the conditions (b) and (c), because *fall* describes a change in spatial location and *push* describes a cause for change in spatial location. So *Explanation*(α, β) is nonmonotonically inferred to be part of the representation of the semantic content of (11). The thing to stress about this is that the inference (d) about β's speech act is logically dependent on the premise (b) about the antecedent utterance α's lexical and compositional semantics. So the theory of alignment for relational speech acts involves reasoning about the semantics of both the current utterance and the antecedent utterance.

The examples (12), (13) and (15) above demonstrate that, like indicatives, the typology of relational speech acts associated with questions is also richer than the traditional speech act view. The relational speech act *Q-Elab*, for example, connects a question to some antecedent utterance, and can be thought of as a *subtype* of the speech act of questioning. Default axioms for inferring *IQAP* and *Q-Elab* are given below:

- `IQAP`: $(\langle \tau, \alpha, \beta \rangle \wedge \alpha :? \wedge \beta : |) > IQAP(\alpha, \beta)$
- `Q-Elab`: $(\langle \tau, \alpha, \beta \rangle \wedge \beta :?) > \textit{Q-Elab}(\alpha, \beta)$

In words, `IQAP` says that the default contribution of a response to a question is to provide sufficient information that the questioner can infer an answer to it. And `Q-Elab` says that the default role of a question is to help achieve a SARG of an antecedent utterance. These axioms are used in DICE to infer that B's speech act is *IQAP* in (12) and (13), and *Q-Elab* in (15). These axioms, although derived from more general Gricean principles (see Asher and Lascarides 1998a, Asher 1999), allow one to infer the speech acts on the basis of just the sentence moods. They must be defaults, because obviously further information could block the inference.

In a similar way, imperatives can appear in a variety of rhetorical relationships. For example in (37), the action described in the imperative is part of a plan to achieve the SARG of the antecedent utterance, a relation we represent as *R-Elab* (for Request Elaboration):

(37) a. A: I want to catch the 10:20 train to London.
b. B: Go to platform 1.

R-Elab is a subtype of the speech act of requesting, and it can be inferred by default from the sentence moods in a manner that's similar to `Q-Elab` above. In (37), inferring that *R-Elab* connects (37b) to (37a) leads to additional inferences: the SARG of (37b) is not only that the action of going to platform 1 is performed, but that it's performed as part of a plan to catch the 10:20 train. This imposes additional constraints on the requested action; for example, on the time it's performed.

The speech acts associated with questions and requests feature an important connection to the speaker's mental states, specifically his SARGs in our terminology. This is central to speech act theory and to theories of discourse interpretation that utilize plan recognition techniques (e.g., Litman and Allen 1990, Carberry and Lambert 1999). But so far, we have largely abstracted away from issues about speakers' cognitive states. This is partly because many rhetorical relations don't impose specific constraints on SARGs: when someone utters (11), why did he do it? Because he wanted to tell a story? In order to kill time? The interpreter doesn't necessarily know, nor does he have to know in order to understand its meaning.

But we don't want to deny the importance of cognitive states in interpretation. Many relational speech acts, like *Q-Elab* and *R-Elab*, impose constraints on the underlying SARGs of its arguments. And although information about an agent's attitudes (including his SARGs) isn't identical with the rhetorical and logical structure, these types of information interact. Such interactions, for example, are essential to analyzing (35). SDRT is able to predict it's incoherent via reasoning with a *combination* of compositional semantics, logical structure and goals. Using any subset of these distinct sources of information would have been insufficient.

In fact the only aspect of the connection between dialogue and underlying intentions that we will formalize here is the link between particular speech act types and their SARGs. For example, following Searle and others, the rule `Question Related Goals (QRG)` below specifies that the SARG of an interrogative is that the questioner believe an answer:

- `Question Related Goals (QRG):`
  $QAP(\alpha, \beta) > \mathcal{I}_{Agent(\alpha)} stit(\mathcal{B}_{Agent(\alpha)} K_\beta)$

K_β is the proposition labelled by β; $QAP(\alpha, \beta)$ means that β labels a direct answer to α; $\mathcal{B}_{Agent(\alpha)}(\phi)$ means that the agent who uttered α believes that ϕ; and *stit* (*see to it that*) is an operator from propositions into plans, such that *stit*(p) denotes the plans whose outcome necessarily include that p is true. So asking a question can lead to an inference about at least one of its SARGs. In the future, we will write SARG(α, ϕ) to mean that the speaker who uttered α has the SARG of making ϕ true—so the consequent of `QRG` can be written as SARG$(\alpha, \mathcal{B}_{Agent(\alpha)} K_\beta)$. Similarly, `Request Related Goals` formalizes that an imperative has the default SARG that the action it describes be performed:

- `Request Related Goals (RRG):`
  $\alpha :! > \mathcal{I}_{Agent(\alpha)} stit(K_\alpha)$

All relational speech acts R generate particular SARGs; e.g., for *Explanation*, the SARG is that the interpreter believe both the explanans and that it is an explanans. But here we'll only make use of more general SARGs associated with assertions, which we can derive via Gricean reasoning that we've formalized elsewhere (Asher and Lascarides 1998a, Asher 1999): the SARG of an assertion is to convey the contents thereof.

The cognitive modelling component also includes general axioms of rationality and cooperativity, which link belief, intention and action. And we also assume that detailed world knowledge can be imported into the cognitive modelling component of SDRT, together with knowledge about agents' general desires, and so on (for discussion see Asher et al. 1995, Asher and Fernando 1999). We won't reiterate the discussion here (see Asher 1999 for details), but simply assume that we have adequate cognitive axioms to treat our examples.

4 Conventionalized and Unconventionalized ISAs

On the basis of cognitive modelling, an interpreter may infer from one speech act and its SARG, as defined by axioms like `QRG` and `RRG`, another intention that is typically associated, again via rules like `QRG` and `RRG`, with some other speech act of incompatible type. This kind of reasoning is the trademark of unconventionalized ISAs. The inference to this second SARG typically triggers the inference to an additional implicit speech act; implicit in the sense that it wasn't assigned to the utterance via its linguistic form. We'll encapsulate this by allowing semantic and pragmatic reasoning to lead to the construction of a dot type. The dot type in turn requires that semantic and pragmatic information is sufficient to infer a rhetorical relation between the two constituent speech acts. We'll then assume, as with conventionalized ISAs, that a type conflict with the predication for attachment—e.g., $\langle\tau, \alpha, \beta\rangle$—leads to an application of `Dot Exploitation`.

We use various information sources to infer the speech acts of an utterance, as differences in these information sources form the basis for distinguishing conventionalized from unconventionalized ISAs. The framework SDRT already assumes a high degree of modularity, in order to ensure that the competence model of dialogue interpretation is computable (Lascarides and Asher, 1999). Here, the separation of conventional information about interpretation from other information (e.g., principles of rationality and cooperativity) forms the basis of (a) distinguishing conventionalized from unconventionalized ISAs; and (b) modelling speech act blocking, where there's conflict resolution between conventional clues about speech acts on the one hand, and calculable clues about speech acts on the other (details of blocking are given in Section 6).

Since we need to keep conventional information separate from other information, we assume a set $\mathcal{C}$ (standing for Convention) of facts and rules, which encapsulate conventional information. $\mathcal{C}$ contains, among other things, the DICE axioms, information about linguistic form, compositional and lexical semantics, and the semantics of speech acts (e.g., the fact that a question denotes a set of propositions, and the fact that $\mathit{Explanation}(\alpha, \beta)$ entails that α and β are true). This is the purely linguistic part of the theory. $\mathcal{R}$, on the other hand, contains the axioms relevant to cognitive modelling (e.g., `QRG`, `RRG` and Gricean-style axioms of rationality and cooperativity) and also other nonlinguistic elements that are transferred, when needed, from world knowledge.

With this in mind, we can now formulate our coherence constraint on complex speech act types. If a speech act β is of complex type $t_1 \bullet t_2$, its constituents γ_1 of type t_1 and γ_2 of type t_2 must be related by a rhetorical relation, that can be inferred (by default) from conventional information, contextual information and pragmatic reasoning. This coherence constraint applies prior to `Dot Exploitation` in our derivations.

- `Coherence Constraint on Complex Speech Act Types`
  Suppose that
  - $\langle\tau, \alpha, \beta\rangle$
  - $\beta : t_1 \bullet t_2$
  - $\mathit{O\text{-}Elab}(\beta, \gamma_1) \wedge \mathit{O\text{-}Elab}(\beta, \gamma_2)$
  - $\gamma_1 : t_1 \wedge \gamma_2 : t_2$

Then: $\mathcal{R}, \mathcal{C}, \langle \tau, \alpha, \gamma_1 \rangle, \langle \tau', \gamma_1, \gamma_2 \rangle, \mathit{Info}(\gamma_1, \gamma_2) \mathrel{|\!\sim} R(\gamma_1, \gamma_2)$,
where τ' labels an SDRS that results from attaching γ_1 to α in the SDRS labelled by τ.

This constraint singles out γ_1 for attachment to α, but of course γ_2 could be chosen instead; we'll forego writing the constraint with the roles of γ_1 and γ_2 reversed. Recall that we assumed that $\langle \tau, \alpha, \beta \rangle$ is well typed only if α and β are assigned simple speech act types. This rule spells out precisely the conditions under which a type conflict that's generated by $\langle \tau, \alpha, \beta \rangle$ (when β's a dot type) can be resolved: if this coherence constraint is satisfied, we can use `Dot Exploitation`; if it isn't, we can't.

Let's turn now to conventionalized ISAs. Conventionalization *transfers* Gricean reasoning into (complex) semantic type assignments within $\mathcal{C}$. More formally, we have the situation described in `Conventionalized ISAs` below, where $dot(\beta)$ means that β is assigned a dot type:[13]

- `Conventionalized ISAs`
  Suppose that β is uttered in the context τ, of which α is part. Then β is a conventionalized ISA (written *conv*-ISA(β)) iff:

$$\mathcal{C}, \langle \tau, \alpha, \beta \rangle \vdash dot(\beta)$$

In words: β is a conventionalized ISA if the extended notion of grammar (which includes, for instance, DICE axioms) assigns the speech act β a complex type. Note that this generates a type conflict, because $\langle \tau, \alpha, \beta \rangle$ demands that β be a simple type. But $\mathcal{C}$ isn't sensitive to this. `Dot Exploitation` will take care of the type conflict independently.

For the sake of space, we gloss over details of exactly how the (extended) grammar would assign these complex types. But such assignments are entirely compatible with current constraint-based grammars (e.g., HPSG, Pollard and Sag 1994). First, these grammars advocate a seamless transition among syntax, semantics and certain contextual factors. So in principle, the context component of phrasal signs in the grammar can be used to constrain the semantic type of the utterance (e.g., it can be used to ensure that *Is the Pope Catholic?* isn't of complex semantic type when there's no *yes/no* question in the context). Second, in order to handle the varying degrees of compositionality of idioms (see Nunberg *et al.* (1994) for details), such grammars allow one to treat phrasal signs as primitive entries within the lexicon. Like idioms, conventionalized ISAs are subject to varying degrees of compositionality: *Is the Pope Catholic?* cannot be assigned its complex type in a compositional manner; but *Can you X?* must be assigned the complex type *question*•*request* for many values of X, indicating that a degree of compositionality is possible, though not wholly so because of lexical exceptions (e.g., *Can you hear at the back?*). Since phrases as well as words are specified within the lexicon, such grammars can account for this semi-productivity of conventionalized ISAs.

The seamless interface between syntax and semantics in such grammars allows the acceptability of modifiers like *please* to be determined by the *semantic type* of its argument (i.e., it must be a *request*), instead of its syntactic category (e.g., that it be imperative). Grammars like HPSG already use semantic typing to specify the subcategorization frames for verbs.

[13] Note that for cases like (1), τ and α may be $\top$, but for cases like (18b), α will need to be a yes/no question.

Adopting similar, semantic-based well-formedness constraints on *please* allows for a uniform analysis across both imperatives and conventionalized ISAs like (1), which have *request* as a constituent type. A similar semantic-based account would account for the linguistic data we observed in Section 2.4.

The *process* by which ISAs become conventionalized, although an interesting issue, is also beyond the scope of this paper. Frequency effects are almost certainly involved, just as they are at the lexical level (Briscoe and Copestake, 1999). But we're concerned with the *synchronic* knowledge that a competent language user has, rather than the diachronic process of conventionalization. So we'll say nothing more on how $\mathcal{C}$ changes and comes to validate `Conventionalized ISAs`.

Unconventionalized ISAs are speech acts that aren't conventionalized, but thanks to axioms in $\mathcal{C}$ and $\mathcal{R}$, are nevertheless assigned a dotted type (in that particular discourse context):

- `Unconventionalized ISAs`
  Suppose that β is uttered in the context τ, of which α is part. Then β is an unconventionalized ISA iff:
  (a) $\mathcal{R}, \mathcal{C}, \langle\tau, \alpha, \beta\rangle, \mathit{Info}(\tau, \beta) \mathrel{|\!\sim} \mathit{dot}(\beta)$;
  (b) $\neg \mathit{conv}\text{-}\mathrm{ISA}(\beta)$

Clause (a) means that $\mathcal{R}$ (i.e., principles like cooperativity and `QRG`), conventional information $\mathcal{C}$, plus the information $\langle\tau, \alpha, \beta\rangle \wedge \mathit{Info}(\tau, \beta)$ about the context β is related to, all nonmonotonically yield an inference that β has a complex type. We analyse unconventionalized ISAs in terms of complex types, so that as with conventionalized ISAs, `Dot Exploitation` can be used to link responses to the utterance to either component; e.g., in (38), B's acknowledgement in (38c) is linked to the request that's part of the meaning of (38b) in this context, and a label that's typed as *request* is introduced via `Dot Exploitation`:

(38) a. A: It's really cold outside, and the door's not closed.
b. A [after having made some effort to get the door]: I can't reach the door.
c. B: I'll get it.

As we suggested earlier, this complex type for (38b) is typically inferred because a a SARG ϕ (close the door) is inferred, which isn't conventionally associated with (38b), but which *is* associated with some incompatible speech act type γ (i.e., request). So reasoning about the SARGs of (38b) typically justifies the inference given in (a) of `Unconventionalized ISAs`. The rule `Inferring Dot Types` spells this out in more detail:

- `Inferring Dot Types`
  Suppose for some γ that:
  (a) $\mathcal{R}, \mathcal{C}, \langle\tau, \alpha, \beta\rangle, \mathit{Info}(\tau, \beta) \mathrel{|\!\sim} \mathrm{SARG}(\beta, \psi)$;
  (b) $\beta : t_1, \neg(\mathit{Info}(\beta) > \mathrm{SARG}(\beta, \psi)$;
  (c) $\gamma : t_2, (\mathit{Info}(\gamma) > \mathrm{SARG}(\gamma, \psi)) \wedge \mathit{incompatible}(\gamma, \beta)$.
  Then:
  (d) $\beta : t_1 \bullet t_2$

`Inferring Dot Types` says that if you can derive via pragmatic reasoning a SARG ψ for a speech act β that, by (b), you don't associate with it conventionally via some axiom like `QRG` or `RRG`, but, by (c), you *do* associated with some speech act γ that's of an incompatible type to β, then that suffices to make β have a complex type that consists of γ's type and β's original and conventional type.

5 Calculating Indirect Speech Acts

Let's now apply the above framework to the analysis of ISAs. We start with an example where a **question also acts like an assertion**:

(15) a. A: Let's meet next weekend.

b. B: How about Saturday?

The default axiom `Q-Elab` applies when interpreting (15b), and its consequences are consistent with the premises. So according to the underlying logic for $>$, *Q-Elab*(α, β) is inferred. Now, the speech act *Q-Elab*(α, β) imposes constraints on the intended content of β: all potential answers γ to β must help specify a plan to achieve the SARG of α. Assuming Cole's (1975) analysis of *let's*, the SARG of α is to meet next weekend. So all answers to β must help narrow the search for the temporal parameter in the plan to meet.

This constraint on β plus facts about *cognitive modelling*—i.e., knowledge of planning strategies for finding the temporal parameter in a plan—together provide more information about what β meant in this context. In particular, the search for the temporal parameter is narrowed only if (a) the Saturday B mentions in (15b) is the Saturday in next weekend,[14] and (b) B can meet A on Saturday (next weekend). If he can't, then regardless of whether A's response to the question is positive or negative, cognitive principles of planning will entail that A and B will fail to mutually know further information that narrows the search for a time to meet. Therefore, *Q-Elab* plus cognitive modelling lead to an inference about the communicative goal behind β: B intended that A believe that he can meet A on Saturday. And B will know that A will deduce this, because the cognitive principles and semantics of *Q-Elab* are mutually known.

But this goal is normally associated with an assertion (being, as it is, a goal of belief transfer). It's an additional goal to the one that follows from the axiom `QRG`; i.e., to know an answer to the question. Furthermore, an assertion is a distinct and incompatible speech act type from a question (because of their compositional semantics). Thus clauses (a)–(c) of `Inferring Dot Topes` are all satisfied, with the value of γ being the assertion that B can meet A on Saturday (for this conventionally has the goal that we inferred for β). So by `Inferring Dot Types`, we infer that β is a dot type. But this complex type clashes with the predications on β; in particular, $\langle \alpha, \alpha, \beta \rangle$. So by `Dot Exploitation`, this triggers the introduction of a new label δ of type question that's related to α via *Q-Elab*.[15] δ is related to β via *O-Elab*, making both

[14] In fact, this interpretation of the anaphoric expression *Saturday* follows from linguistic form too, because according to DRT's notion of accessibility, *next weekend* is the only accessible antecedent.

[15] Notice crucially that we haven't said in DICE that attaching and dot typing are inconsistent; that's only available to the typing system and is resolved by `Dot Exploitation`.

β and δ available for further attachment.

To check the `Coherence Constraint on Complex Types`, we have to compute a rhetorical link between the constituent types of β: i.e., the question δ and the assertion γ that B can meet A on Saturday. As γ is stative, it attaches with the relation *Background* to δ, reflecting that it provides background information (details of the DICE axioms for inferring *Background* are in Lascarides and Asher, 1993). So uttering (15b) is a coherent, unconventionalized ISA.

Now let's consider an example where **assertions act also like requests**. We mentioned earlier that (20) can be used to signal a request, as well as to *explain* why that request is being made.

(20) I'm out of gas.

Asher (1999) shows how the speaker and hearer mutually infer a relation *Explanation** that relates the assertion to the request for help. *Explanation**(α, β) means that β explains why the speech act of uttering α was performed (this contrasts with *Explanation*(α, β), where the explanation is at the content level rather than the speech act level). The reasoning that leads to this interpretation of (20) exploits cognitive principles of rationality and cooperativity, as well as particular assumptions about the speaker's cognitive state. Let ϕ be the proposition that the speaker A is out of gas. Gricean reasoning suitably formalized yields the default conclusion that A believes that ϕ. Let's assume, as seems reasonable, that being out of gas is normally undesirable. Then further Gricean reasoning should lead to the speaker's intention to get out of this state. This SARG follows via `RRG` from the request $!stit(\neg\phi)$. And requests and assertions are incompatible speech act types. So `Inferring Dot Types` applies as before. The `Coherence Constraint on Complex Types` is also satisfied because of the rhetorical relation *Explanation** that's inferrable between the assertion and the request components of the dot type. So (20) is a coherent unconventionalized ISA.

Now let's consider cases where **questions act as requests**:

(39) Can you close the window?

Like (1), this is a conventionalized ISA of type *question* $\bullet$ *request*, but this interpretation is also calculable (Bach and Harnish, 1979). We focus on this here. Roughly, the conventional goals of questions (as specified by `QRG`) don't apply here because the answer is mutually known; and script-like information about preconditions therefore leads to an inference that the underlying intention is for the action of closing the window to be performed. So by `Inferring Dot Types`, (39) is assigned a dot type. In fact, Asher (1999) demonstrates that an apparent violation of a maxim (in this case, `QRG`), isn't necessary for inferring a request; rather, if the answer to the question isn't known (as it might not be for a question like *Can you reach that suitcase?*), cognitive reasoning leads to a conditional request interpretation.

(40) paraphrases (39), and so similar reasoning about cognitive states applies, thereby predicting that the interpretation of (40) as an ISA is calculable. The same reasoning makes the ISA interpretation of (5) calculable:

(40) Do you have the physical ability to close the window?

(5) Do you have the physical ability to pass the salt?

Clearly, these calculable implicatures must get blocked. We'll now provide the means to do this via a conventional blocking principle.

6 Blocking and Unblocking

We've claimed that conventionalized ISAs block inferences from paraphrases to implicit speech acts even when they're calculable. We'll now describe a formal model of this blocking.

In an attempt to do justice to the complexity of interaction between the different information sources that contribute to dialogue interpretation—both conventional and non-conventional—many researchers in AI have assumed a radically unmodular framework, so that a single reasoning process can access the different kinds of information at any time (e.g., Hobbs *et al*, 1993). In such a framework, calling one bit of reasoning conventionalized and another not doesn't have logical bite; it's simply a way of labelling what justifies the default rule's existence within the model. But we believe that this lack of modularity is unsatisfactory, because its treatment of conflict resolution among the different knowledge resources during interpretation is unsatisfactory.

To see why, consider the requirements that speech act blocking imposes on a *logical* model of interpretation. Blocking involves two different sources of information—the conventionalized ISA on the one hand, and the cognitive axioms on the other—giving conflicting default clues about the speech act of the paraphrases. And normally, the conventional clues win and the ISA is blocked. Now, nonmonotonic logic provides two options for modelling one default clue winning over another conflicting default clue. First, one can resolve conflict *within* the logic when one clue is strictly more specific than the other. This is the way conflicts are normally resolved in DICE (Lascarides and Asher, 1993). But it won't work here, because there are no grounds for saying that conventional information is more specific than cognitive information; they are simply different knowledge sources with different domains of application. The second option is to (i) *separate* one kind of default reasoning from the other—in other words separate the conventional clues from the cognitive ones—and (ii) ensure that the conventional module takes priority over the cognitive one, by ensuring that the defaults arising from that module are worked out *first*, and then treated as indefeasible when one considers the affects of rationality and cooperativity on interpretation. That way, conventional clues will win because indefeasible information always overrides conflicting default information.[16] Our earlier definitions of ISAs separated conventional information $\mathcal{C}$ from the non-conventional information $\mathcal{R}$. We will exploit this to model blocking.

6.1 The Speech Act Blocking Principle

We now define the ISA blocking rule in $\mathcal{C}$. Let $\alpha \approx \beta$ mean that α and β have similar compositional and lexical semantics. We gloss over a precise definition of similarity, but for

[16] We talked earlier of conventional clues winning over conflicting clues by default because of unblocking. We'll see shortly that we can assume that conventional clues always win, so long as facts about the discourse context are taken to be part of conventional information.

the purposes of this paper, assume that (41a-e) are all in $\approx$ relations to each other:

(41) a. Can you pass the salt?
 b. Are you able to pass the salt?
 c. Do you have the physical ability to pass the salt?
 d. Could you pass the salt?

The blocking principle then stipulates the following: in words, if (a) an utterance β is a conventionalized ISA, (b) β' is not a conventionalized ISA, but (c) β' is similar in content to β (e.g, (41a) is similar to (41b–d)), then normally, updating the context τ with β' does *not* yield an ISA interpretation of β'. More formally, the blocking rule for ISAs is stipulated in `Speech Act Blocking (SAB)` below; so `SAB` is one of the axioms in $\mathcal{C}$:

- `Speech Act Blocking (SAB)`
  If
  (a) $conv\text{-ISA}(\beta)$; and
  (b) $\neg conv\text{-ISA}(\beta')$
  Then
  (c) $\vdash \beta \approx \beta' >$
  $\neg((\langle \tau, \alpha, \beta' \rangle \wedge Info(\tau, \beta')) > dot(\beta'))$

`SAB` is a monotonic rule. Given our assumptions about the complex semantic type that $\mathcal{C}$ assigns to (1) (in the null context), clause (a) would be satisfied if τ and α are $\top$ and β is (1). Clause (b) would be satisfied if β' is (5), since unlike (1), the grammar doesn't assign this the dot type $question \bullet request$. Moreover, $\beta \approx \beta'$ (i.e., the antecedent of the default rule (c) that's the consequent of `SAB`) is verified by these values for β and β'. The consequent of the default rule (c) is a negated default (i.e., something of the form $\neg(A > B)$). This negated default stipulates that it's *not* the case that if we were updating the discourse context with β' (in our example, if we were updating the null context $\top$ with (5)), then normally, we would infer that β' has a dot type. In other words, β' is not interpreted as an ISA. And in our example, (5) isn't normally interpreted as a request. However, `SAB` doesn't propose an alternative interpretation of (5); that is a matter for other clues in context to decide.

This shows that `SAB`, if it interacts correctly with other knowledge resources such as $\mathcal{R}$, could model what we want; i.e., that the conventionalized ISA (1) normally blocks (5) from being an ISA. Moreover, `SAB` correctly doesn't allow (41a) to block (41d), because even though their contents are similar, clause (b) of `SAB` won't be verified.

Note that although `SAB` is a monotonic rule, its consequent is a default axiom; it's of the form $\beta \approx \beta' > \neg\Delta(\beta')$, where $\neg\Delta(\beta')$ is a gloss for the negated default. This is default because blocking isn't absolute (e.g., (32)). If clause (c) were replaced with a monotonic rule $\beta \approx \beta' \rightarrow \neg\Delta(\beta')$, then it would be *logically impossible* to infer $\Delta(\beta')$ when clauses (a), (b) and $\beta \approx \beta'$ are true, *whatever* else we add to the premises. In other words, unblocking would be logically impossible, contrary to the evidence in (32). We'll show how `SAB` allows unblocking in section 6.2.

We must make sure that information from $\mathcal{C}$ and from $\mathcal{R}$ interact in such a way that conventionalized ISAs block similar sentences from having similar ISAs, *even if* that ISA is calculable.

To achieve this we need two things. First, we need to establish this relative priority between $\mathcal{C}$ and $\mathcal{R}$. Second, we need to extend the nonmonotonic logic, so that we can use negated defaults, as found in clause (c) of `SAB`, to block certain conclusions that would follow if those negated defaults were absent from the premises.

As we've already mentioned, since we want $\mathcal{C}$ to take priority over $\mathcal{R}$, all defaults that are derivable from $\mathcal{C}$ must be treated as indefeasible when we come to consider the defaults that are derivable from $\mathcal{R}$ as well. How can this interaction between $\mathcal{C}$ and $\mathcal{R}$ be represented? Note that if $\mathcal{C}|\!\sim\phi$, then according to the logic Commonsense Entailment (CE), ϕ will *monotonically* follow from the intersection $\underline{\mathcal{C}}$ of all nonmonotonic extensions of $\mathcal{C}$ (Asher and Morreau 1991, Morreau 1995). That is, $\mathcal{C}|\!\sim\phi$ if and only if $\underline{\mathcal{C}} \vdash \phi$. So we get the interaction between $\mathcal{C}$ and $\mathcal{R}$ that we want, if we assume that the premises for computing discourse interpretation is $\underline{\mathcal{C}} \cup \mathcal{R}$. For if $\mathcal{C}|\!\sim\phi$, then even if $\mathcal{R}|\!\sim\neg\phi$, we're still guaranteed that $\underline{\mathcal{C}} \cup \mathcal{R}|\!\sim\phi$ because monotonic information is always preserved. In contrast, $\mathcal{C}|\!\sim\phi$ will *not* guarantee that $\mathcal{C} \cup \mathcal{R}|\!\sim\phi$.

Intuitively, this prioritization of $\mathcal{C}$ over $\mathcal{R}$ during discourse interpretation is justified, on the grounds that people generally choose their speech forms carefully: if changing the speech form but not the compositional semantics would change the implicatures that a competent language user would infer, then the speaker will take this into account during language production.

We now turn to the problem of negated defaults in inference. Though we won't argue for it here, we believe that the pattern of reasoning (b) is intuitively compelling and should be validated even if (a) is valid too:

- `Negated Defaults`
  - (a) $\Gamma, C \;|\!\sim\; B$
  - (b) $\Gamma, C, \neg(C > B) \;|\!\not\sim\; B$

An extension of CE can be made to validate `Negated Defaults` (see Asher and Mao (2000) for details). And this inference pattern is just what we need in order to ensure that `SAB` causes blocking. For suppose that β is a conventionalized ISA, β' is not a conventionalized ISA, and β is similar to β'. These facts verify clauses (a) and (b) of SAB, so clause (c) of `SAB` is inferred (whatever the context, since this is monotonic). Since the antecedent to this default rule (c) is also verified, and all our assumptions are part of $\mathcal{C}$—being, as they are, conventional information—and since `SAB` is also part of $\mathcal{C}$, we are in a situation where $\mathcal{C}|\!\sim\neg\Delta(\beta')$, where $\neg\Delta(\beta')$ is the negated default in `SAB`. So, by the logic CE, $\underline{\mathcal{C}} \vdash \neg\Delta(\beta')$. Expanding out what $\Delta(\beta')$ stands for, this means $\underline{\mathcal{C}} \vdash \neg((\langle\tau,\alpha,\beta'\rangle \wedge \mathit{Info}(\tau,\beta')) > \mathit{dot}(\beta')$.

But given the way we have organized the reasoning which underlies discourse interpretation, recognizing β''s speech act involves reasoning with the following premises: $\underline{\mathcal{C}}$, $\mathcal{R}$ and $\langle\tau,\alpha,\beta'\rangle \wedge \mathit{Info}(\tau,\beta')$. And `Negated Defaults` validates clause (b) of `Negated Defaults in Blocking` in this case, even if clause (a) below is validated too:

- `Negated Defaults in Blocking`
  - (a) $\mathcal{R}, \langle\tau,\alpha,\beta'\rangle, \mathit{Info}(\tau,\beta') \;|\!\sim\; \mathit{dot}(\beta')$
  - (b) $\mathcal{R}, \underline{\mathcal{C}}, \langle\tau,\alpha,\beta'\rangle \wedge \mathit{Info}(\tau,\beta') \;|\!\not\sim\; \mathit{dot}(\beta')$

So interpreting β' does *not* yield an ISA interpretation, even if this interpretation is calculable (i.e., clause (a) above holds). In fact we can say more. The negated default ensures that at least one extension of the >-statements does not imply $dot(\beta')$. Therefore, thanks to the rule `Inferring Dot Types`, in at least one extension we cannot infer for β' SARGs appropriate to unconventionalized ISAs.

6.2 Unblocking

As (32) demonstrates, context can unblock a normally blocked ISA interpration:

(32) [Context: No one except perhaps for B can get to the meeting because they're out of town, and B is sick but in town]

a. A (addressed to B): Are you physically able to get to the meeting?

b. B: OK, I'll go.

Intuitively, unblocking occurs in (32a) because the context sets up a situation where it is precisely B's physical condition that is at issue in getting someone to the meeting, making (32a) an appropriate wording for expressing the request. Clearly, utterances that are equivalent in meaning may give rise to different cognitive effects because of their surface form. And these effects, which are often subtle and difficult to quantify, seem to be sufficient for unblocking. For the sake of simplicity, we'll represent these effects via the relation $\mathit{Cntxt}(\tau, \beta, \beta')$, to signal that the particular linguistic form β' is appropriate in the context τ for performing the speech act that's conventionally associated with β. So in (32), we'll assume that $\mathit{Cntxt}(\tau, \beta, \beta')$ holds, where β is the conventionalized ISA *Can you go the meeting?*, β' is (32a), and τ labels a representation of the context described in (32). $\mathit{Cntxt}(\tau, \beta, \beta')$ reflects the fact that β''s *explicit* linguistic form in this context conveys precisely what's at issue in achieving the goal that's typically associated with the speech act β—in this case, getting someone attend the meeting.

With these ingredients in place, we can stipulate a general axiom of unblocking as a principle in $\mathcal{C}$:

- `Unblocking`
  If
  (a) $\mathit{conv}\text{-ISA}(\beta) \wedge \beta : t_1 \bullet t_2$ and
  (b) $\neg \mathit{conv}\text{-ISA}(\beta')$
  Then
  (c′) $\vdash (\beta \approx \beta' \wedge \mathit{Cntxt}(\tau, \beta, \beta')) > ((\langle \tau, \alpha, \beta' \rangle \wedge \mathit{Info}(\tau, \beta')) > \beta' : t_1 \bullet t_2)$

In words: if β is a conventionalized ISA then the semantically similar β' is also an ISA of the same complex type as β, *as long as* β''s linguistic form makes it a particularly appropriate way of performing the complex speech act within the context τ.

Suppose that the antecedents to `SAB` and to `Unblocking` both hold. Then given that these rules are monotonic, they will generate the >-statements in (c) and (c′) of these rules. But note that the consequent of (c′) is the following default:

$$(\langle \tau, \alpha, \beta' \rangle \wedge \mathit{Info}(\tau, \beta')) > \beta' : t_1 \bullet t_2$$

And since the logic makes $>$ closed on the right under logical consequence, this monotonically implies the following (Asher and Mao, 2000):

$$(\langle \tau, \alpha, \beta' \rangle \wedge \mathit{Info}(\tau, \beta')) > \mathit{dot}(\beta')$$

The consequence of (c) in SAB is exactly this default rule, but negated. So the axioms (c) and (c′) have conflicting consequences. But the antecedent to the $>$-statement (c′) is more specific, because it contains the added conjunct $\mathit{Cntxt}(\tau, \beta, \beta')$. Therefore, CE ensures that if the antecedents to both these defaults apply—i.e., $\beta \approx \beta'$ and $\mathit{Cntxt}(\beta', \gamma)$ are both true—then unblocking wins. That is, $\langle \tau, \alpha, \beta' \rangle \wedge \mathit{Info}(\tau, \beta') > \mathit{dot}(\beta')$ is inferred in $\mathcal{C}$, leading to an interpretation of β as an ISA even though it's normally blocked.

7 Speech Act Blocking vs. Lexical Blocking

Discussions of lexical blocking bear many analogies to this speech act blocking. For both *pig* meaning *meat* and (5), there are default regularities that (potentially) generate the blocked form. For *pig*, it's a default regularity about sense extension of animal nouns. For (5), its interpretation as a request is inferrable from the cognitive axioms $\mathcal{R}$. Furthermore, both *pig* meaning meat and (5) meaning a request are highly marked, because there is an alternative *conventional* way of conveying the same meaning: *pork* in the case of *pig*, and (1) in the case of (5). The lexical entry for *pig* meaning meat would have the same syntactic/semantic representation as *pork*; the difference is in orthography. Similarly, the intended interpretation of (5) would be similar to (1); the difference is largely orthographical. Finally, in both cases, blocking isn't absolute. Given these analogies, it might be illuminating to compare existing approaches to lexical blocking with our model of speech act blocking.

There are at least three ways in which blocking has been modelled at the lexical level: by default inheritance (e.g., Evans and Gazdar 1989, Russell *et al.* 1991); by more general default reasoning (e.g., Briscoe *et al.*, 1996); and by lexical frequencies (Briscoe and Copestake, 1999). Let's examine each of these in turn.

In the default inheritance approach, blocking occurs because there's (a) a specificity relation between two lexical regularities and (b) conflict between these regularities. For example, *slept* blocks *sleeped*, because (a) *sleep* is a subtype of **irr-t-verb**, which in turn is a subtype of **verb**,[17] and (b) the value **+t** that's assigned to the PASTP forms for **irr-t-verb** is incompatible with the value **+ed** that's assigned to PASTP forms for **verb**s. Organizing the information this way allows default inheritance to generate the well-formed past participle *slept* and block the ill-formed one. However, there are problems with this approach. First, blocking is absolute. But it isn't absolute at the semantic level, and so default inheritance can't model blocking as a unified principle of lexical organization. Secondly, a precondition to blocking in this model is that a constraint on a subtype overrides a conflicting constraint on a supertype. While this is plausible for *sleeped* vs. *slept*—because **irr-t-verb** is subsumed by **verb**—it isn't plausible for *pig* vs. *pork*. One cannot plausibly assume that *pork* is subsumed by the count or mass nouns *pig*, or by the sense extension rule.

In reaction to these problems, Briscoe *et al.* (1996) model blocking with more general default

[17] We'll write types/values in bold font, and features in small caps.

reasoning. This account allows for unblocking. But it still suffers from the defect that one cannot block a value without specifying another, alternative value (e.g., to block *pig* meaning meat, an alternative interpretation for *pig* must be stipulated). This may not be possible in all circumstances where blocking seems to be present; e.g., there may be insufficient information for preferring the animal count sense of *pig*, or its metaphorical 'greedy' sense, or the policeman sense, and so on. It's utopian to expect that during interpretation, we can infer an alternative value to replace the blocked one: lexical gaps and varying degrees of conventionalization show that sometimes, the alternative value to the blocked one may be inferrable only on the basis of the *discourse context*, rather than via lexical information alone.

The frequency-approach to blocking treats it as a matter of performance rather than competence (Briscoe and Copestake, 1999). In line with Bauer (1983), lexical rules are treated as fully productive (and so there's an entry in the lexicon for *pig* meaning meat), and one assumes that frequency effects determine assessments of the degree of acceptability of a particular derived word form. In essence, the degree of 'item familiarity' that competent language users have with a particular lexical item determines the user's judgements about its relative novelty/conventionality. In this account, blocking occurs because of an assumption that Grice's maxims predict the following: word forms which frequently mean ϕ are favoured over those that rarely mean ϕ in language production; and conversely during interpretation, one disambiguates the word form w to its more frequent senses rather than its rarer ones. For example, *pig* meaning meat is not blocked at the *symbolic* level by *pork*. However, *pork* is much more frequent than the (estimated) frequency of *pig* meaning meat during parsing. And so the above Gricean assumption predicts that in practice, *pig* meaning meat is usually blocked by *pork*.

The frequency-based approach still shares the problem with the previous approaches that one can block a value only by specifying another value: here, the more frequent value overrides rarer conflicting values. Further, the frequency-based approach cannot apply at the speech act level for practical reasons. Bauer (1993) attests that the range of possible word senses, though large, is not so great that one couldn't use frequency information as a basis for acceptability judgements. But he also points out that we can't capitalize on this at the sentential level, because the sheer size of the combinatoric possibilities prevents us from estimating the probability that a particular sentential form (with corresponding compositional semantics) conveys some particular *intended* content. Finally, it seems to us that modelling blocking via frequencies confuses the process by which something becomes conventionalized with the fact that for a competent language speaker it *is* conventionalized. Frequency plausibly contributes to conventionalization. But we don't believe that the representation of the lexicon of a competent language user need preserve the frequency information that led to this conventionalization. Treating blocking this way also makes it difficult to see how one can predict the contexts where unblocking occurs, and why unblocking typically triggers additional content of the word form than simply its blocked sense.

Let's now consider how these accounts of lexical blocking compare with the account of speech act blocking we've presented. The prerequisites to preemption by synonymy that are utilized in `SAB` are much weaker than in the lexical accounts: Speech act blocking isn't preconditioned by specificity and conflict; nor does it require an alternative (conflicting) interpretation to the blocked one to be stipulated. This is thanks to the logical structure of `SAB`. First, its antecedent doesn't stipulate anything about the relative specificity of β and the form β' that it blocks. Second, when SAB fires, it yields a negated default. The logic CE ensures that this negated default (normally) blocks the ISA interpretation, even though neither SAB nor any other principle need generate an alternative interpretation.

We've suggested that some cases of lexical blocking require weaker preconditions than those demanded by existing accounts. So it might be worth exploring whether the weaker preconditions to blocking we've stipulated in SAB can be applied within the lexicon. That way, we could maintain blocking as a unified principle of lexical organization: the overarching blocking principle will impose relatively weak preconditions, although stronger preconditions could still be stipulated for particular kinds of blocking; e.g., specificity and conflict are necessary for *morphological* blocking.

We offer here one way of reconstructing our approach to blocking within the lexicon. We assume, like the frequency approach above, that the lexicon contains lexical entries that are (usually) blocked interpretations of a given word form. For example, the lexicon contains an entry for pig meaning meat. However, a lexical blocking rule, which is similar in structure to SAB, constrains how words are interpreted during parsing: a similar conventionalized/established form stops one from interpreting a word form from having a derived sense. So, for example, because of *pork*, a word form *pig* that's encountered during parsing isn't usually taken to mean meat (w_o below stands for the word form w):

- `Lexical Blocking for Pig`:
  $(pork_o > meat \wedge pig_o \approx pork) > \neg(pig_o > meat)$

In words, if $pork_o$ is normally interpreted to mean meat and pig_o is 'similar' to *pork* (we'll say what this means shortly), then normally, it's not the case that pig_o is (also) interpreted to mean meat. In fact, this is a particular instance of a more general schema:

- `Lexical Blocking`
  $(w_o > s \wedge w'_o \approx w) > \neg(w'_o > s)$

We assume that $\approx$ is defined so that $pig_o \approx pork$ holds. For example, $w_o \approx w$ means that w_o isn't the same as w'_o (i.e., their orthographies are different), and for at least one sense of w_o, their syntax and semantics are the same. Assuming that rules which encapsulate sense extension interact correctly with `Lexical Blocking`, this will ensure that pig meaning meat is blocked without assuming a subsumption relation between *pork* and *pig*, and without offering an alternative interpretation of the word form pig_o.

The above rule can be used quite straightforwardly, in tandem with other axioms about interpretation, to explain the conversational implicatures that arise during unblocking. For example, suppose the discourse is incoherent unless *pig* means *meat* (e.g., when a vegetarian utters *There's pig on the menu*). Then to avoid a *real* violation of principles of interpretation, one must assume that *pig* denotes meat. But thanks to `Lexical Blocking`, the maxim of Quantity would then be violated for real, if there wasn't a specific reason for using *pork* rather than *pig*. Accommodating a specific reason as part of the interpretation—much as conversational implicatures are in general—leads to the word form *pig* having loaded semantic effects that *pork* would lack.

8 Conclusion

In this paper, we have provided a formal account of speech acts and indirect speech acts. We have argued that logical tools familiar from formal theories of discourse semantics—such as semantic types and type shifting rules—provide a more robust basis for individuating different speech acts than propositional attitudes/intentions can do on their own. We have treated speech act types as relations between utterances rather than as properties of utterances; and we have provided a nonmonotonic framework for inferring these relational speech acts from various clues, including the linguistic forms of the current utterance and its antecedent.

We have followed Gordon and Lakoff (1975) in that some indirect speech acts are calculable from Gricean reasoning. Following Searle (1975) and Morgan (1978), others appear to be conventionalized. We have provided a formally precise account of conventionalized ISAs in terms of complex semantic types. Just as complex types predict at the lexical level that a word like *book* can linguistically behave, simultaneously, as if it's a physical object and an abstract object, they predict at the speech act level that an utterance like (1) can linguistically behave, simultaneously, as if it's both a question and a request.

We also have provided a precise account of the tension between conventionalized and calculable ISAs. Following Horn and Bayer (1984), we advocated a blocking principle on ISAs, which bears analogies to lexical blocking and neg-raising blocking. That is, a conventionalized ISA normally blocks its paraphrases from having an ISA interpretation, even if this interpretation is calculable. Our formal characterization of this principle requires modularization of the knowledge resources: conventional information $\mathcal{C}$ is separated from the cognitive principles $\mathcal{R}$, although there are complex interactions between the two. We also compared our formal characterization of ISA blocking with accounts of lexical blocking.

We close by noting that our account leaves open several problems. We need to examine more closely the criteria for stipulating that two utterances have similar content (as represented in SAB). We also need to provide more rigorous tests for whether a speech act is conventionalized or not. We hope to address these questions in future work.

9 Acknowledgements

We would like to thank Hans Kamp, Colin Matheson, Johanna Moore, Ivan Sag, Mark Steedman, Rich Thomason and anonymous reviewers for helpful comments and feedback. Alex Lascarides is funded by an ESRC (UK) research fellowship.

10 References

Asher, N. [1993] *Reference to Abstract Objects in Discourse*, Kluwer Academic Publishers.

Asher, N. [1999] Discourse Structure and the Logic of Conversation, in K. Turner (ed.) *Current Research in the Semantics Pragmatics Interface.*

Asher, N., J. Busquets and D. Hardt [1997] Discourse Parallelism, Scope, and Ellipsis, in S.

Peters and H. de Swart (eds.) *Proceedings of SALT 7*, Stanford 1997

Asher, N. and A. Lascarides, [1995] Lexical Disambiguation in a Discourse Context, *Journal of Semantics*, **12** number 1, pp69–108, Oxford University Press.

Asher, N. and A. Lascarides [1998a] Questions in Dialogue, *Linguistics and Philosophy*, **23.3**, pp237–309.

Asher, N. and A. Lascarides [1998b] Bridging, *Journal of Semantics*, **15**, pp83–113, OUP.

Asher, N. and A. Lascarides [1998c] The Semantics and Pragmatics of Presupposition, *Journal of Semantics*, **15**, pp239–299, OUP.

Asher, N. and Mao, Y. [2000] Reasoning with Negated Defaults in Commonsense Entailment, *Proceedings of Logic Colloquim*, Paris, also submitted to *Journal of Symbolic Logic*.

Asher, N. and M. Morreau [1991] Commonsense Entailment: A Modal Theory of Nonmonotonic Reasoning, *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI91)*, Sydney, pp387–392.

Asher, N. and J. Pustejovsky [forthcoming] The Metaphysics of Words in Context, submitted to *Journal of Logic, Language and Information*, Kluwer Academic Publishers.

Austin, J. L. [1962] *How to Do Things With Words*, Clarendon Press.

Bach, K. and R. M. Harnish [1979] *Linguistic Communication and Speech Acts*

Bach, K. [1995] Standardization vs. Conventionalization, *Linguistics and Philosophy*, **18.6**, pp677–686, Kluwer Academic Publishers.

Bach, K. and R. M. Harnish [1992] How Performatives Work—A Reply to Searle, *Linguistics and Philosophy*, **15.1**, pp93–110, Kluwer Academic Publishers.

Bauer, L. [1983] *English Word-formation* CUP

Briscoe, E. J. and A. Copestake [1999] Lexical Rules in Constraint-based Grammar, *Computational Linguistics*, **25.4**, pp487–526.

Briscoe, E. J., A. Copestake and A. Lascarides [1996] Blocking, in St. Dizier, P. and Viegas, E. *Computational Lexical Semantics*, pp273–302, Cambridge University Press.

Brown, P. and S. C. Levinson [1987] *Politeness*, Cambridge University Press.

Cole, P. [1975] The Synchronic and Diachronic Status of Conversational Implicature, in Peter Cole and Jerry L. Morgan (eds.), *Syntax and Semantics Volume 3: Speech Acts*, Academic Press, pp257–288.

Copestake, A. and E. J. Briscoe [1995] Semi-productive Polysemy and Sense Extension, *Journal of Semantics*, **12.1**, pp15–67, Oxford University Press.

Davison, A. [1975] Indirect Speech Acts and What to Do with Them, in Cole, P. and J. Morgan (eds.) *Syntax and Semantics 3: Speech Acts*, pp143–185, Academic Press, New York.

Evans, R. and G. Gazdar [1989] Inference in DATR *4th Conference of the European Chapter of the Association for Computational Linguistics (EACL-1989)*, Manchester, England, pp66–71.

Fernando, T. [1997] Ambiguity under Changing Contexts, *Linguistics and Philsophy*, **20.6**, pp575–606, Kluwer Academic Publishers.

Ginzburg, J. and I. A. Sag [2000] *English Interrogative Constructions*, CSLI publications.

Gordon, D. and Lakoff, G. [1975] Conversational Postulates, in Peter Cole and Jerry L. Morgan (eds.), *Syntax and Semantics Volume 3: Speech Acts*, Academic Press, pp83–106.

Grice, H. P [1975] Logic and Conversation, in Cole, Peter and Morgan, Jerry L. (eds.) *Syntax and Semantics Volume 3: Speech Acts*, pp41-58. Academic Press.

Grosz, B. and Sidner, C. [1986] Attention, Intentions and the Structure of Discourse. *Computational Linguistics*, **12**, 175–204.

Grosz, B. and Sidner, C. [1990] Plans for Discourse, in P. R. Cohen, J. Morgan and M. Pollack (eds.) *Intentions in Communication*, pp.417–444, MIT Press.

Hintikka, J. [1983] *The Game of Language*, Reidel.

Hobbs, J. R. [1985] On the Coherence and Structure of Discourse. Report No. CSLI-85-37, Center for the Study of Language and Information, October, 1985.

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

Horn, L. R. [1989] *A Natural History of Negation*, University of Chicago Press.

Horn, L. R. and S. Bayer [1984] Short-circuited Implicature: A Negative Contribution, *Linguistics and Philosophy*, **7.4**, pp397–414, Kluwer Academic Publishers.

Kamp, H. and Reyle, U. [1993] *From Discourse to Logic*, Kluwer Academic Publishers.

Lakoff, G. [1973] Some Thoughts on Transderivational Constraints, in B. B. Kachru *et al.* (eds.), *Issues in Linguistics*, pp442–452, University of Illinois Press.

Lascarides, A. and N. Asher [1993] Temporal Interpretation, Discourse Relations and Commonsense Entailment, in *Linguistics and Philosophy*, **16**, pp.437–493.

Lascarides, A. and N. Asher [1999] Cognitive States, Discourse Structure and the Content of Dialogue, in *Proceedings from Amstelogue 1999*, pp1–12, Amsterdam, May 4th–7th 1999.

Levinson, S. C. [2000] *Presumptive Meanings: The Theory of Generalized Conversational Implicature*, MIT Press.

Lewis, D. [1969] *Convention: A Philosophical Study*, Blackwells.

Litman, D. and J. Allen [1990] Discourse Processing and Commonsense Plans, in P. R. Cohen, J. Morgan and M. Pollack (eds.) *Intentions in Communication*, pp.365–388, MIT Press.

Montague, R. [1974] *Formal Philosophy: Selected Papers of Richard Montague*, Reidel, Dordrecht.

Morgan, J. L. [1975] Some Interactions of Syntax and Pragmatics, in Peter Cole and Jerry L. Morgan (eds.), *Syntax and Semantics Volume 3: Speech Acts*, Academic Press, pp289–303.

Morgan, J. L. [1978] Two Types of Convention in Indirect Speech Acts, in Peter Cole (ed.), *Syntax and Semantics Volume 9: Pragmatics*, Academic Press, pp261–280.

Nunberg, G., I. A. Sag and T. Wasow [1994] Idioms, *Language*, **70**, pp491–538.

Perrault, R. [1990] An Application of Default Logic to Speech Act Theory, in P. R. Cohen, J. Morgan and M. Pollack (eds.) *Intentions in Communication*, pp.161-186, MIT Press.

Pollard, C. and I. Sag [1994] *Head Driven Phrase Structure Grammar*, CSLI Lecture Notes, University of Chicago Press, Stanford CA.

Pustejovsky, J. [1995] *The Generative Lexicon*, MIT Press.

Russell, G., J. Carroll and S. Warwick [1991] Multiple default inheritance in a unification-based lexicon, *29th Annual Meeting of the Association for Computational Linguistics (ACL-91)* Berkeley, California, pp215–221.

Sag, I. A. and T. Wasow [1999] *Syntactic Theory: A Formal Introduction*, CSLI Publications.

Sadock, J. [1974] *Towards a Linguistic Theory of Speech Acts*, Academic Press.

Searle, J. [1969] *Speech Acts*, Cambridge University Press.

Searle, J. [1975] *Indirect Speech Acts*, in Peter Cole and Jerry L. Morgan (eds.), *Syntax and Semantics Volume 3: Speech Acts*, Academic Press, pp59–82.

Searle, J. [1976] The Classification of Illocutionary Acts, *Language and Society*, **5**, pp1–24.

Segerberg, K. [1990] Validity and Satisfaction in Imperative Logic, *Notre Dame Journal of Formal Logic*, **31.2**, pp203–221.

Mann, W. C. and S. A. Thompson [1987] Rhetorical Structure Theory: A Framework for the Analysis of Texts. *International Pragmatics Association Papers in Pragmatics*, 1, pp79-105.

Vanderveken, D. [1990] *Meaning and Speech Acts*, Cambridge University Press.