

Semantics in Berlin

Alex Lascarides

November 2, 2014

How do we evaluate semantic analyses?

The quality of a semantic analysis can be assessed along several dimensions:

Truth Conditions: The semantic representation of a linguistic unit should capture all its truth conditions and nothing more. But tests for deciding what counts as truth conditional content are controversial (e.g., in my view, the tests proposed in Wilson (1975) involving embeddings are problematic).

Compositional Semantics: The compositional semantic analysis of a linguistic string should capture all and only the meaning that's derivable purely from its form; e.g., it should not entail meaning that holds in many discourse contexts but not all of them. It should also be sufficient, in combination with other resources such as real world knowledge, to support reasoning about the specific pragmatic interpretation of the sentence in context, *without* resorting to direct access to sentential syntax. Again, the boundary between compositional semantics and pragmatics is controversial (e.g., consider the literature on gapping, binding etc).

Derivability: The compositional semantic representation of a string should be derivable via syntax within a wide-coverage grammar; that is, it should not assume idiosyncratic syntactic types or modes of combination that undermine the linguistic analyses of other constructions.

Discourse Effects: The logical form of a linguistic string should not only accurately capture the string's truth conditions (in context) but also capture constraints on the possible interpretations of *subsequent* utterances in the discourse: e.g., the data on modal subordination, anaphoric dependencies across sentence boundaries, and so on.

There are no doubt other criteria we care about. I would like to start our meeting by making the metrics by which we assess a semantic analysis clear and transparent.

So in that vein I will now indulge myself, by declaring my own views about what's needed for modelling discourse interpretation and hence also what demands are placed on compositional semantics.

My own research on discourse has focussed on modelling anaphoric dependencies across clausal boundaries. In all my work, I've observed that two things are primitive and irreducible

components for modelling anaphoric dependencies: the *logical structure* of the individual sentences (e.g., the presence of modals, negation, conditionals etc); and the way the contents of discourse segments combine via *coherence relations* (e.g., *Contrast, Elaboration, Acceptance, Denial*) to create larger segments.

Dynamic semantics (Kamp and Reyle, 1993, Groenendijk and Stokhof, 1991) was designed to deal with the interaction between anaphora and logical structure. It can, for instance, account for the fact that (1ab) is acceptable where (1a'b) is not (Partee, 1984), even though within a classical logic (1a) and (1a') are equivalent:

- (1) a. One of the ten balls is not in the bag.
- a'. Nine of the ten balls are in the bag.
- b. It's under the sofa.

But these orthodox dynamic models both overgenerate and undergenerate the available antecedents to anaphora (Asher and Lascarides, 2003).

In contrast, there's plenty of evidence that inferring coherence relations and resolving anaphora are mutually dependent (Hobbs, 1979). For instance, in (2ab) *he* is John while in (2ab') *he* is Bill; these distinct (but defeasible) interpretations are caused (defeasible) reasoning about what makes the discourse coherent—*Explanation* for (2ab) and *Result* for (2ab').

- (2) a. John can open Bill's safe.
- b. He knows the combination.
- b'. He should change the combination.

Indeed, coherence relations can themselves act as antecedents to anaphoric expressions, even though inferring them is *defeasible*. In dialogue (3), for instance, the (implicated) *Explanation* relation in *A*'s contribution must be the antecedent for the anaphoric expression *no* in *B*'s utterance:

- (3) A: John went to jail. He embezzled the company funds.
- B: No. He embezzled the company funds but he went to jail because he was convicted of tax fraud.

One could not simply use the logical conjunction of the content of the sentences in *A*'s contribution (with the pronoun *he* resolved to John) to resolve the meaning of *no*, because the result would make *B*'s overall contribution inconsistent, given the content of *B*'s second sentence.

While implicated coherence relations seem to be available as antecedents to anaphora, not all implicatures are. For instance, *A*'s contribution in (4) carries the conversational implicature that *A* believes it's raining (Grice, 1975), but even though both *A* and *B* share this knowledge, *B* cannot respond with the elided construction, intending it to mean *I also believe it's raining*.

- (4) A: It's raining.
- B: ???Me too.

So while it's not necessary for conversational implicatures to be a part of the semantic representation of discourse, it is essential that coherence relations are, so that we can adequately specify constraints on how the subsequent discourse can proceed, and how subsequent (anaphoric) contributions can be interpreted. But coherence relations on their own don't account for data like (1). That's why in Segmented Discourse Representation Theory (SDRT, Asher and Lascarides (2003)) we combine dynamic semantics and the discourse structure that's generated by the coherence relations. SDRT is flexible about everything else other than these two ingredients: e.g., the formal language for expressing logical forms can be anything (e.g., first order, typed, DRSS) so long as it includes predicate symbols for representing coherence relations; and the model theory of the formal language can be anything so long as it's dynamic.

An alternative to dynamic semantics for handling the interaction between logical structure and anaphora would be to countenance building syntactic structures for multi-sentence discourse rather than individual sentences in isolation. That way, a constituent in one sentence can c-command another in syntax, and so an existential quantifier, say, that's introduced in one sentence can take semantic scope over the content of a subsequent sentence. This approach can capture the inter-sentential anaphoric bindings within a static semantics; e.g., for *A man walked. He talked.* Dynamic syntax (Cann et al., 2005) and DLTAG (Webber et al., 1999) model syntax this way, although they are neutral about whether the model theory for the logical forms is static or dynamic. But I am sceptical that all the semantic dependencies that a discourse logical form should capture are retrievable from a discourse syntax that is expressed as a tree (the evidence for this is relatively subtle, but I could talk about it at the meeting should people wish). At any rate, I want to maintain a more traditional notion of grammar, in which syntax is confined to analysing sentences in isolation (contra Ginzburg and Sag (2001), for instance). Using dynamic semantics lets me do this.

Conversely, alternatives to using coherence relations for capturing the content of discourse have been proposed, such as using Questions Under Discussion (QUD, Ginzburg (2012)). But I think that in the end, when one completes the QUD model of discourse by articulating a theory of which questions get accommodated (and when), and the principles by which questions persist or get removed from the QUD as the discourse proceeds (something that so far is missing from QUD models), one will find that the QUD model is a notational variant of coherence relations. After all, answers to questions bear all the hallmarks of semantic relations expressed by coherence relations: things like *Explanation* (why?), *Narration* (and then what?), *Denial* (don't you agree?) etc.

I've so far focussed on the *interpretation* of logical forms of discourse (i.e., it should be dynamic), and what those logical forms look like (i.e., they should feature coherence relations). But *constructing* the logical form for discourse should not be modelled in the same (dynamic and undecidable) logic as the one in which one *interprets* logical form—contra the abductive approach in Hobbs et al. (1993), for instance, or the approach taken by Cooper (2014) using Type Theory with Records (TTR). Computing logical form should be decidable: all competent language users by and large agree on what was said including the resolution of anaphora, but they generally disagree on its cognitive effects—e.g., whether what was said should be believed.

All approaches to semantics that use *underspecified logical forms* (ULFs, e.g., Egg et al. (2001), Bos (1995)) assume this separation between the language and logic for constructing logical

form and the language and logic for interpreting it. The language of ULFs is typically propositional, with predicate symbols and variables of various sorts but no quantifiers or modalities. It is a language for *partially describing* the *form* of a fully specific and contextually resolved logical form for discourse. You can conceive of the syntactic form of a fully specific logical form as a *tree*: dominance means that the constructor on the mother node takes the constructors on its children as arguments; the left to right order determines the argument position. So the model theory for ULFs is static and finite: each model M is a finite tree that completely determines a syntactically well-formed fully specific formula in the language for capturing discourse interpretations. ULF formulae denote *tree fragments*: $M \models u$ just in case u is a partial description of the tree M (and so a partial description of the form of the fully specific logical form that M corresponds to). The model theory of all ULFs works this way (e.g., MRS, RMRS, CLLS, hole semantics Copestake et al. (2001), Koller and Lascarides (2009), Egg et al. (2001), Bos (1995)). They differ in formal details and (critically) in the computational complexity for enumerating all models that satisfy a given formula. But I will abstract over those differences here.

Having a static and finite model theory for ULFs is entirely compatible with interpreting discourse logical form dynamically. It's also compatible with the choice of language for representing discourse logical form: although ULFs are expressed in a propositional language with predicate symbols, the discourse logical form can (and should) include quantifiers, and it can be a first order language, higher order, DRSS, or whatever. In other words, the output of grammars like the ERG is compatible with many theories of discourse, such as SDRT, Ginzburg's QUD model etc.

On the other hand, one needs to test whether the ULFs that a grammar assigns to linguistic strings provide sufficiently rich information for supporting inferences about fully specific interpretations in context, given other information that's extraneous to the grammar (e.g., real world knowledge). The fully specific interpretation must include specific coherence relations, antecedents to anaphora and so on.¹ The model theory I've just described for ULFs ensures that ULFs constrain the *form* of a specific logical form but abstracts away almost entirely from its *meaning*. So constructing a discourse interpretation from the ULFs for its sentences won't know about, say, the entailment between "John talked" and "someone talked" unless that is explicitly imported into the mechanisms for constructing logical form. So the linguist must control just how much of the denotation of a word or phrase influences the construction of the logical form of discourse; to allow all aspects of (dynamic) discourse meaning to influence construction would make logical form construction undecidable (because inferring a discourse interpretation is defeasible and so subject to consistency tests). Controlling which aspects of a phrase's meaning influences the construction of logical form is another contentious area.

While I don't expect people at this meeting to hold the same views that I do about constructing and interpreting logical form, I thought I should at least make my prejudices explicit, so that you can interpret what I say below in light of them!

Since I assume this separation between the language and logic for encoding compositional semantics (and for constructing logical form) and the language and logic for encoding the contextually resolved and specific discourse interpretation, in what follows I will attempt to

¹Actually, there's plenty of evidence that people don't resolve all ambiguities during discourse interpretation, but we ignore that here.

minimise confusion by using *italics* for representing the predicate symbols in ULFs, and SMALL CAPS for representing the constructors in the language for representing specific discourse interpretations—i.e., the constructors that predicate symbols in the ULFs describe. Sometimes a predicate symbol at the ULF level maps to a unique constructor (perhaps a predicate symbol or a modal operator) at the discourse level, so that all we get is a change in font! For instance, $l : \textit{girl}(x)$ is satisfied by any and only those discourse logical forms containing the predication $\text{GIRL}(x)$ as a part. But this isn't always so: a ULF predicate symbol may denote one of several constructors in the language of discourse logical form. For instance, the two senses of the noun *mogul* aren't distinguished within the grammar, because the distinct senses don't yield distinct syntactic behaviours for the word. So there is only one predicate symbol in the ULF that the word maps to: $l : \textit{mogul}(x)$. But this ULF formula is satisfied by any logical form that contains $\text{MOGUL_SNOW}(x)$ as a part and it's also satisfied by any logical form that contains $\text{MOGUL_EMP}(x)$ as a part. The discourse logical form must know the difference between snow and Chinese emperors to capture the correct pragmatic inferences, particularly concerning coherence relations. But the syntax/semantics interface can abstract away from these distinct lexical senses.

01 (Closed) Clausal Complements

- [0101] We thought that they barked.
- [0102] They forgot that they had voted.
- [0103] She knew who hadn't arrived yet.
- [0104] He relied on what they said.
- [0105] How they voted surprised me.
- [0106] That they voted early surprised me.
- [0107] Them voting early surprised me.
- [0108] Their voting early surprised me.
- [0109] Their early voting surprised me.

In all the above examples, the verb in the matrix clause should take a *scopal* argument. That way, at the level of discourse interpretation one can account for the *de dicto/de re* readings for “We thought a dog barked” (for instance) while maintaining a relatively straightforward relationship between the scope of constructors in the discourse logical form and the scope constraints in the ULFs that describe it.

It's not obvious how one should derive the ULFs for (5a) and (5b) so that they make similar contributions to discourse interpretation:

- (5) a. John thought Harry is a fool.
- b. John thought of Harry as a fool.

It's tempting to assume that the ULF for (5b) should include a predicate symbol *think* that takes an event and two individuals as arguments, one that denotes John and one that denotes Harry. But this puts $\textit{fool}(e, x)$ outside the scope of the referentially opaque context of *think*. So any straightforward interpretation of the scope constraints expressed in a ULF would entail

that in any discourse interpretation, (5b) publicly commits the *speaker* to Harry being a fool; this clearly isn't the right semantics. So I think that the predicate symbol *think* introduced by (5b) must have a *scopal* variable as an argument, not an individual one. I rule out the argument being the event introduced by the preposition because like (0101) such an analysis doesn't capture the distinct *de dicto/de re* readings. But if the argument to the predicate symbol *think* in the ULF of (5) is scopal, then how does that ULF relate to the ULF for *John thought of Harry*?²

The verb *forget* in (0102), which takes a clausal complement (as opposed to a noun or a base VP), *presupposes* the proposition denoted by the clausal complement. For instance (0102) entails that they had voted and this entailment survives under embeddings: *they didn't forget that they had voted* and *John believes that they forgot that they had voted* also entail that they had voted.³ In other words, this entailment satisfies the tests that it can project out from embeddings.

Some grammatical frameworks separate potentially presupposed content from non-presupposed content. Others, such as the ERG, do not. I actually think that the latter is the right move, because the propositions that can project from embeddings can include content's that's not a part of the compositional semantics of any given phrase in the sentence. For instance, implicative verbs—*manage to X* and *succeed in X*—presuppose that you *tried X* (Karttunen, 1971). The phrase *forget to X* presupposes you intended to *X*.⁴ In these cases, the attempt or the intention is not directly derivable from any syntactic constituent, and to embed it in the syntactic analysis of the word *manage*, *succeed* or *forget* is complex and unnecessary. But these contents cannot be ignored entirely. They must be a part of the pragmatically resolved discourse content, since they can act as an antecedent to a denial move—I find (6ab) fine while (6a'b) is odd (unless in the prior discourse, *A* has already committed to John having tried to open the safe).

- (6) a. A: John didn't manage to open the safe.
 a'. A: John didn't open the safe.
 b. B: You're wrong; he didn't even try.

So these contents influence the ways in which the subsequent discourse can be interpreted.

But if the grammar doesn't identify potential presuppositions (so that subsequent modules for constructing the discourse logical form can reason about their relative scopal position), then something else must do this. That 'something else' must take as input the ULF that's derived within the grammar, and it must output a new ULF that includes the potential presuppositions. It is perfectly possible to write a module that 'parses' the structure of the ULF that's derived by the grammar so as to identify potentially presupposed content, which must then form a part of the discourse logical form (and in particular, in constructing a

²There's also *John thought Harry a fool*, but this is a really weird construction that's not very productive and so is perhaps better handled as a multi-word expression.

³But *They forgot to vote* does not entail or presuppose that they voted.

⁴By "presuppose" here, I mean that there is some content associated with a word or phrase that projects out from the syntactic embedding of that word or phrase. I'm agnostic about whether we actually use the term "presupposition" for such contents or not. But we do need to account for them at the level of discourse logical form so as to capture the right entailments and resolution of subsequent anaphora.

discourse interpretation, one must determine the relative semantic scope of the potentially presupposed content to the asserted content).

Recall that ULFs are partial descriptions of trees. So we add a new formula to the ULF language: $h \approx h'$ is satisfied by any tree M where h and h' denote the root of subtrees T_h and $T_{h'}$ within M that are *strictly isomorphic*: that is, there is an isomorphism f that maps T_h to $T_{h'}$ (and *vice versa*) and for all $l \in T_h$, the constructor at l is identical (syntactically) to the constructor at $f(l)$.⁵ In other words, the elementary predications labelled by h and h' are (syntactically) the same, and (recursively) the arguments to those elementary predications are syntactically the same. Thus $h \approx h'$ basically means that h and h' label (structured) predications that are completely parallel to one another. Note that $h = h'$ entails $h \approx h'$.

We also add a predicate symbol ∂ to the ULF language, corresponding to Beaver’s (2001) symbol ∂ for “potentially presupposed”. Then assuming the grammar outputs something like (R)MRS, the following two rule schemata, which would be extraneous to the grammar, identify the potentially presupposed content of *know* and of *manage* from the ULF that’s output by the grammar, and effectively adds it to the (partial) description of the discourse logical form:⁶

$$\begin{aligned} \mathbf{know:} \quad & (l : \text{know_v_1}(e, x, h_1) \wedge h_1 =_q h_2) \rightarrow \\ & (l' : \partial(h_3) \wedge h_3 \approx h_2) \\ \mathbf{manage:} \quad & (l : \text{manage_v_1}(e, x, h_1) \wedge h_1 =_q h_2) \rightarrow \\ & (l' : \partial(h_3) \wedge h_4 : \text{try_v_1}(e', x, h_5) \wedge h_5 \approx h_2)) \end{aligned}$$

Note that there is a tendency for quantified phrases that were syntactically commanded by *know* to project out from the embedding too. For instance, *a woman* and *every dog* outscope *every man* in (7), but a scope ambiguity remains between *a woman* and *every dog*:

(7) Every man knows that a woman loved every dog.

The above rule schema for *know* captures this because the potentially presupposed content is $\text{love}(e, x, y)$ and if this projects out then it must do so with the relevant quantifiers so that the variables remain bound. On the other hand, if the discourse context is one where the presupposition gets bound or accommodated locally, then the quantifiers may project, or may not. You get a similar interaction between the potentially presupposed content and quantifiers for the verb *manage*.

Rule schemata that are extraneous to the grammar and that relate ULF strings are needed more generally for constructing discourse logical forms; e.g., for computing commonsense entailments from compositional semantics, which in turn is needed for identifying the pragmatically preferred coherence relation, antecedents to anaphora and so on. They’re needed also for reasoning about lexical semantics, and for handling RTE tasks. So I think that using rules extraneous to the grammar to identify potentially presupposed content from the ULF can be captured with such rule schemata too.

⁵This no doubt has horrible consequences in the complexity of inference; in other words, it makes computing all possible specific interpretations that are satisfied by a ULF NP complete at best. But complexity is pretty awful even without this relation being a part of the language. Alex Koller will be more qualified to discuss these issues than I am.

⁶Local binding of the presupposition would correspond to $h \approx h'$ being satisfied by the fact that $h = h'$.

MRS Fingerprint. Given the above discussion, I think that verbs taking clausal complements—whether the complementiser *that* is obligatory or not—should introduce the following within the semantic representation derived in the grammar:

EP: $1:\text{pred}(e)$, $\text{ARGn}(h)$, where $\text{ARGn}=\text{ARG0}$ if the clausal complement is in ‘subject’ position, or its ARG2 if its in ‘object’ position.

HCONS: $h=_q\ 1'$ where $1'$ is the top label of the S complement.

This allows quantifiers introduced within the S complement to take any scope. The projection from embeddings that some verbs exhibit should be handled outside the grammar, even for factive verbs. I don’t think there should be any difference in the ULFs of the sentences *We thought they barked* and *We thought that they barked*, and for sentences like (0106) where the presence of *that* is needed for syntactic well-formedness, there is no trace of it in the semantic representation.

02 Coordination

- [0201] Abrams and Browne danced.
- [0202] The dog, cat, or picture arrived.
- [0203] My friend and colleague arrived.
- [0204] Abrams barked and was old.
- [0205] Abrams wanted and expected to arrive.
- [0206] The dog is old and fierce.
- [0207] The dog barked on Monday and on Wednesday.
- [0208] He hemmed and hawed.

I know virtually nothing about coordination: I find it scary and have always avoided it!

I don’t think (0201) should entail that Abrams and Browne engaged in some kind of joint action. That is, there should be at least one compositional semantic representation that supports a discourse interpretation where Abrams and Browne’s dancing are at different spatio-temporal locations. This is effectively a collective vs. distributive ambiguity.

I’m quite sympathetic to encoding this ambiguity as a sense ambiguity of *and*, in which case it is an ambiguity that is handled *outside* the grammar (just like the sense ambiguity of *mogul*, etc). In Russian, for instance, there are several connectives corresponding to *and*: *i* doesn’t imply the subjects do the action together, while *s* does. So in some languages these senses are lexicalised differently. In English, they’re not. And I think (though I’m not sure) that *and* behaves the same from a syntactic perspective regardless of whether it results in a collective or a distributive reading.

For VP coordination, I think the arguments to *and* should be *handles*. This allows a discourse logical form that features a sentence with VP-coordination to encode *coherence relations* between the content of VPs: there’s lots of evidence that such coherence relations should be a part of logical form, so as to predict the resolution of elided constructions (Asher et al., 1997), word sense disambiguation (Lascarides et al., 1996) and non-sentential fragments (Schlangen,

2003). So this means that *and* in the ULF should take 3 arguments which: (a) are all individual variables (NP coordination); or (b) the first is an event variable and the second and third are handles (for VP and S coordination); or (c) all three arguments are event variables (for ADJ and PREP coordination). For VP and S coordination, the scopal arguments to *and* should outscope the top label of the syntactic head of the conjunction’s first and second complements respectively.

If *and* joins VPs or Ss and so semantically relates two pieces of propositional content, the specific sense of *and* can resolve to any veridical coherence relation (e.g., *Narration*, *Continuation*, *Elaboration*, and so on). Interestingly, if *and* joins an imperative and a indicative then the relation it encodes can be conditional rather than veridical:

(8) Smoke 20 cigarettes a day and you’ll die before you’re 30.

For individuals, $and(x_1, x_2, x_3)$ can resolve to (semantic) equality (i.e., to $x_1 = x_2 = x_3$) or it can resolve to a *member-of* relation (i.e., $member-of(x_2, x_1) \wedge member-of(x_3, x_1)$). However *and* is resolved, the relation between x_1 and x_2 is the same as the relation between x_1 and x_3 ; e.g., you can’t have $x_1 = x_2$ and $member-off(x_3, x_1)$. Resolving *and* to *member-of* relations makes x_1 a complex individual, and so the (dynamic) model theory in which discourse logical form is interpreted should ensure that the context change potential for any formula $P(e, x_1)$ can be satisfied even when e is a complex event, denoting several temporal-spatio locations (in other words, one needn’t assume that the times at which e occurs is a single convex spatio-temporal interval). So this would deal with the fact that (0201) should be satisfied by a model where Abrams and Browne arrived at different times and a different place.

I think that this approach is problematic in analysing (0202) though. My intuition is that the logical form associated with the syntax ((my (friend and colleague)) arrived) should entail that the friend denotes the same person as the colleague (i.e., *and* resolves to equality). But if that ‘equality’ reading is caused by a word sense ambiguity on *and*, then it is handled extraneously to the grammar (even though only one quantifier is introduced by something linguistically explicit, namely *my*). In contrast, I’m fine with the syntax (((my friend) and (colleague)) arrived) entailing either reading (i.e., that there’s one person, or more than one).

Example (0205) involves coordinating two controlling verbs. I would want the discourse logical form for (0205) to be (9)—much simplified, in that I’m ignoring quantifiers:

(9) $\pi_0 : \text{CONTINUATION}(\pi_1, \pi_2)$
 $\pi_1 : \text{WANT}(e_1, a, \text{ARRIVE}(e_3, a))$
 $\pi_2 : \text{EXPECT}(e_2, a, \text{ARRIVE}(e_3, a))$

In other words, *and* resolves to *Continuation*, and the proposition that is wanted and the proposition that is expected are the same: namely, that Abrams arrived, invoking the same *token* arrival event e_3 . So the ULF derived by the grammar should be satisfied by this specific logical form (9). Further, it seems plausible to assume that it is down to the *linguistic form* of (0205) that the proposition that is both wanted and expected is the same. If that is the case, then *any* ULF that is derived by the grammar should entail that there are two formulae $\text{ARRIVE}(e_3, a)$ in its form. We can express this with the *isomorphism* relation $h_1 \approx h_2$ between labels that I introduced earlier for separating asserted content from presupposed content.

Recall that $h_1 \approx h_2$ means that h_1 and h_2 label syntactically identical fully specific formulae. Using this relation, the ULF of (0205) could be something like (10) (again simplified in that I'm ignoring quantifiers):

$$(10) \quad \begin{aligned} h_0 &: \text{and}(e, h_1, h_2) \\ h_3 &: \text{wanted}(e_1, a, h_4) \\ h_5 &: \text{expected}(e_2, a, h_6) \\ h_7 &: \text{arrive}(e_3, a) \\ h_1 =_q h_3, h_2 =_q h_5, h_4 \approx h_6, h_6 =_q h_7 \end{aligned}$$

I have no idea how to construct such a ULF within the grammar, however! The chances are that it's not doable without breaking the derivation of other linguistic constructions; I'm not qualified to judge. If $h_4 \approx h_6$ can't be introduced in a benign way by any linguistic construction, then we would need to add it via postprocessing, much as we add potential presuppositions during post-processing.

MRS Fingerprint. Given the above discussion, can we get away with the following? The MRS for a coordinated construction would include one of the following:

EP: $1:\text{and}(i1, i2, i3)$ where $i1$, $i2$ and $i3$ are resolved within the grammar in one of the following ways:

N and NP Coordination: $1:\text{and}(x1, x2, x3)$, where for N coordination, 1 is also used as the label for the noun complements.

VP and S Coordination: $1:\text{and}(e, h1, h2)$ together with the **HCONS** conditions $h1=_q 11$ and $h2=_q 12$, where 11 and 12 are respectively the top label of the first and second VP or S complements (and so their semantic indices are events).

ADJ and PREP Coordination: $1:\text{and}(e, e1, e2)$, $1:\text{pred1}(e1, x)$, $1:\text{pred2}(e2, y)$, where pred1 and pred2 are the symbols for the relevant prepositions or adjectives.

03 Ellipsis

[0301] Who is?

[0302] Did you?

[0303] There were.

[0304] Abrams tried to.

I supervised David Schlangen's thesis (Schlangen, 2003) on non-sentential fragments, and I still stand by the ULFs he proposed for the compositional semantics of elided constructions, although his syntactic analysis was problematic (both overgenerating and undergenerating well-formed fragments). The compositional semantics for fragments he proposes consist of underspecified predications for the elided bits, which must get resolved outside the grammar (e.g., via inferences concerning discourse coherence (Asher et al., 1997) or higher order unification (Dalrymple et al., 1991)). So, for example: (0304) has a ULF that includes the elementary predications and scoping constraints shown in (11):

- (11) $l : \text{try}(e, a, h)$
 $l' : \text{underspec_pred_rel}(e', a)$
 $h \geq l'$

Note that it's not possible to use $=_q$ to express the scoping constraints, because the elementary predication that l' resolves to via contextual reasoning may include scopal modifiers.

04 'Identity' Copulae

- [0401] Abrams is Browne.
[0402] The theory is that Browne arrived.
[0403] All Browne could do was arrive.
[0404] With Browne the manager, Abrams arrived.
[0405] Browne is a manager.
[0406] The reason: Browne arrived.
[0407] The plan is to sleep more.

I know very little about copulae, but I don't think that the above sentences can generate the same compositional semantics, let alone the specific logical forms in context. (0402), (0403), (0404) and (0407) have focus effects; (0401) and (0405) don't.

Focussing on (0406), I would want the contextually resolved logical form of the discourse (12) to be (12') (simplified, in that quantifiers, tense and presuppositions are ignored):

- (12) a. John left.
b. The reason: Browne arrived.
- (12') $\pi_0 : \text{ELABORATION}(\pi_1, \pi_2)$
 $\pi_1 : \text{LEAVE}(e_1, j)$
 $\pi_2 : \text{EXPLANATION}(\pi_1, \pi_3)$
 $\pi_3 : \text{ARRIVE}(e_2, b)$

How can we infer this logical form from the compositional semantics of (0406)? The noun *reason* is relational: a reason is always a reason why some proposition is true (even the "The reason for the cat" should denote a reason why some proposition involving the cat is true). Further, commonsense reasoning should relate the reason for p being q on the one hand and on the other q being an answer to the question *why p?*. In other words, lexical semantics and commonsense reasoning should validate a default rule for identifying coherence relations during discourse processing that looks as follows in SDRT:

- $(l_1 : \text{reason_for}(x, \alpha) \wedge l_2 : \text{be_id}(x, \beta)) > l_2 : \text{EXPLANATION}(\alpha, \beta)$

In words, if x is the reason why α is true, and x is β , then β explains α . We get this reading supported by the ULF for (0406) if (a) the word *reason* introduces a relational noun whose second argument is scopal and can be anaphorically determined; and (b) we get an 'identity'

relation between the semantic index x of *reason* and the top label of the S (‘identity’ is a loose term here because of the type clash between an individual variable and scopal variable, but we need such identity relations for anaphoric expressions when they resolve to abstract antecedents). In other words, it’s basically the same as the ULF you would get for *The reason is that Browne arrived*. This suggests that the colon should make the same semantic contribution as the phrase *is that* does in the copula.

More generally, I think that there are (at least) two ways in which a string of the form “NP: S” can be well-formed. The first is where the N in the NP is *relational*, and its denotation, together with commonsense reasoning, supports inferences towards certain coherence relations: e.g., *reason* supports *Explanation*; *evidence* supports *Evidence*; *example* and *instance* support *Elaboration*, and *plan* supports *Plan-Elaboration* (I’ll discuss the latter example shortly). In this case, the colon serves to introduce a ‘copula’ predicate symbol between the semantic index of the noun with the top label of the S. In (0406) this would mean that the ULF includes the following:

- (13) $l_1 : reason(x, i_1)$
 $l_2 : arrive(e, b)$
 $l_3 : be_id(x, h_3)$
 $h_3 =_q l_2$

So this sort of construction will produce a well-formed MRS only if the head noun in the NP has a second, *scopal* argument.

The second way in which a string of the form “NP: S” can be well formed is if the NP is an antecedent to some anaphoric expression in the sentence S:

- (14) John’s cat: Mary spoils him.

I can’t think of a discourse context where *him* would denote John (or anyone else) as opposed to his cat. Further, one needs the pronoun and not an elided construction (see (15)) and the construction doesn’t work with all NPs (see (16)):

- (15) *John’s cat: Mary spoils.

- (16) *That: I want it.

Examples like (14) clearly have focus effects (cf. left dislocation, topicalisation). Indeed, I think that sentences like (14) are highly anomalous unless the discourse context supports putting the NP into focus, via a salient alternatives set etc. (Rooth, 1992).

I view the noun *plan* as relational, and at the level of discourse logical form PLAN denotes a relation between an individual variable (that generally denotes a sequence of actions) and a proposition (that denotes the *goal* of the plan). This means that the phrase “the plan for money” denotes in any discourse context a plan whose goal is some proposition where one of its participants is money; given real world knowledge, it’s likely to be a proposition where someone acquires money. Similarly, the phrase “the plan to talk” is a plan whose goal is a proposition that someone (unspecified) talks. So I would analyse (0407) much like (0406),

except that here the verb *is* introduces the identity relation rather than the colon. In other words, its ULF would contain the following:

- (17) $l_1 : plan(x, i_1)$
 $l_2 : sleep(e_1, i_2) \wedge more(e_2, e_1)$
 $l_3 : be_id(e_3, x, h_3)$
 $h_3 =_q l_2$

Similarly to (12), the prior discourse context can provide an antecedent proposition that resolves the missing propositional argument to the noun *plan*:

- (18) I need to have more energy.
The plan is to sleep more.
- (18') $\pi_0 : PLAN-ELABORATION(\pi_1, \pi_2)$
 $\pi_1 : NEED(e_1, i, HAVE(e_2, i, y)) \wedge ENERGY(y) \wedge MORE(e_3, y)$
 $\pi_2 : PLAN(i, \delta(SLEEP(e_5, i) \wedge MORE(e_5, e_6)), HAVE(e_2, i, y))$

(18') assumes that pragmatics resolves the agent of *sleep* to the speaker so as to satisfy the semantics of the coherence relation (Asher and Lascarides, 2003). The copula predicate symbol *be_id* serves to identify the *action* that is the second argument to PLAN with the action that the proposition to sleep more comes about ($\delta\phi$ denotes the action of making ϕ true (Lascarides and Asher, 2004)). The constructor PLAN in the discourse logical form has a quite different structure from the predicate symbol *plan* in the ULF. That's because plans at the domain level must specify the individual(s) who are committed to the plan, the sequence of actions and the specified goal. But the English word *plan* has complements that only specify the goal (*plan to sleep*) or even an NP who is a participant in that goal state (*plan for money*). The participants who are committed to the plan may be specified by the determiner (e.g., *my plan*, *John's plan*), but not necessarily. So identifying the individual(s) that are committed is extraneous to compositional semantics/the grammar.

05 Nominalization

- [0501] To bark bothers Browne.
[0502] Chasing the cat is old.
[0503] What the dog chased bothers Browne.
[0504] The voters support the government's repeal of the law.
[0505] The voters support the law's repeal.
[0506] The student evaluations were negative.

There seems to be some kind of control going on with the infinitival nominalizations while there isn't for gerunds. I.e., (0501) means that Browne barking bothers Browne, while *Barking bothers Browne* doesn't entail that it's Browne doing the barking. I'm not sure that this apparent control is a part of the syntax/semantics interface, or whether it is a pragmatic interpretation borne from defeasible inference. If it is a matter for the grammar, then I'm not

sure how the control can be derived given the standard assumptions made about syntactic and semantic derivations. For cases like (0501) and (0502) I would introduce nominalisation in the discourse language as an NOM that takes a formula as its argument, and semantically its effect is like type raising: i.e., it maps the extension of its argument to its sense (and so $\text{NOM}(\phi)$ does not entail ϕ). This ensures that in the above examples we don't get unwanted entailments, such as the individuals introduced in the nominalisation exist, or that the event introduced by the verb stem occurs etc. Within the grammar, we therefore need a predicate symbol *nom_rel* that takes a scopal argument and denotes the constructor NOM in the (separate) language for discourse logical forms. The predicate symbol *nom_rel* can be introduced in the semantic component of a variety of linguistic constructions, e.g., the construction rule that takes a VP and forms a gerund, or that takes a base-form VP (e.g., *to sleep*) and a VP and forms an S (e.g., (0501)).

“What the dog chased” seems to be the same syntactically and semantically as an embedded interrogative (e.g., *I know what the dog chased*) and so I would analyse “what the dog chased” in (0503) in exactly the same way as it's analysed for embedded interrogatives. Semantically, I agree with Ginzburg (1995) that Groenendijk and Stockhof's (1982) analysis of questions and embedded interrogatives, which entails that an answer must be exhaustive, is problematic. But on the other hand, exhaustiveness of the answers seems to often be a pragmatically preferred interpretation, presumably via Gricean maxims etc. While Ginzburg's semantics for questions doesn't impose exhaustiveness conditions on answerhood, it also doesn't explain this interaction with pragmatics. Further, Ginzburg's analysis abandons the rather nice property that all questions are expressed with a lambda abstract, where the bound variable occupies the trace position left by the *wh*-element. Further, his notion of answerhood is too relaxed, in that *John knows who came to the party* is true if John knows someone came, but not anything about that someone.

Asher and Lascarides (1998, 2003) offer an analysis of questions that follows Groenendijk and Stockhof in defining them in terms of lambda abstracts (e.g., *what the dog chased* is $? \lambda x \iota y (\text{dog}(y), \text{chase}(e, y, x))$), and answers α to $? \lambda x \phi(x)$ are all expressions of the form $\alpha(\lambda x \phi(x))$ such that α entails that some particular object occurs in the extension of $\lambda x \phi(x)$ —thus supporting a *de re* knowledge claim when the question is embedded in attitudinal contexts—or it entails that the extension is empty. Lascarides and Asher (2009b) adapts this semantics to the dynamic semantic setting: following Groenendijk (2003), we treat questions and propositions as the same semantic type (i.e., they both transform an input context into a perhaps different output context). While the context change potential of a question partitions the input context (according to the questions different possible answers), a proposition eliminates elements of the input context. Having questions and propositions be of the same semantic type is very useful for discourse processing because it allows for a uniform analysis of coherence relations when the arguments to the relations are moves of varying surface acts (imperatives, interrogatives or indicatives). For instance the question response to the question in (19) can be captured with the same coherence relation ELABORATION as one would use for (20):

- (19) A: Are Kluwer accepting manuscripts?
 B: What kind of manuscripts?
- (20) A: Kluwer are accepting manuscripts. B: For example, they are accepting monographs.

06 Comparatives

- [0601] The dog is older than the cat.
- [0602] The dog is older than the cat is.
- [0603] The dog is much older.
- [0604] More dogs than cats appeared.
- [0605] A more aggressive dog than mine appeared.
- [0606] It moved more slowly.
- [0607] I fear it more than my brother.

07 Control Relations

- [0701] They forgot to vote.
- [0702] Abrams persuaded the dog to bark.
- [0703] It is easy for the dog to bark.
- [0704] I had a plan to respond.
- [0705] Abrams left without paying.
- [0706] The dog arrived barking.

08 Measure Phrases

- [0801] The dog was ten-and-a-half years old.
- [0802] They finished a ten-mile run.

09 Paranthenicals

- [0901] Abrams (an old friend) arrived.
- [0902] A consultant (maybe it was Abrams) arrived.
- [0903] He is famous -- make that infamous -- in some circles.

Reinhart (1983) highlights the semantic similarities between paranthenicals and *that*-clauses: e.g., *John will be late, I think* and *I think that John will be late*. But this doesn't apply to all paranthenicals. Moreover, I think that the constituent that the paranthenical modifies puts the content of that constituent into focus. I don't think that (0901) is a great answer to the question *What did Abrams do?*, while it would be fine to answer with *Abrams arrived. He's an old friend, by the way*; further, one can use (0901) to answer *Who arrived?*. If paranthenicals interact with focus in this way (and I am very unsure that they do), then *John will be late, I think* has a distinct semantics from *John, I think, will be late*.

Boër and Lycan (1980) observe that paranthenicals are adverbs that can modify the underlying *performative* verb (by which I think they mean that it modifies the performance of the speech act), rather than the main verb of the sentence. This semantic effect of paranthenicals isn't

exhibited in the above examples, but it is in: *Confidentially, John is Catholic* and *John is undiplomatic, frankly*.

I basically have a lot of sympathy for Asher’s (2000) analysis. First, he observes that paratheticals and other discourse particles like *Gee* and *Damn* share a lot of features with presuppositions: they tend to project out from the syntactic context in which they are introduced and they typically contribute a *proposition* to the content of the discourse (whatever the syntactic type of the parathetical). For instance, the most salient reading of (21) is: Mary assures us that the party is over, and if it is, we should leave.

(21) If the party, as Mary assures us, is over, then we should leave.

The examples 0901–0902 also exhibit this kind of behaviour if you embed these sentences in modal or conditional contexts. E.g., *John believes that Abrams (an old friend) arrived* means Abrams is an old friend (of the speaker’s?), and John believes that he arrived. In fact, if you replace *Abrams* with an indefinite, then I think the fact that it comes with a parathetical makes the *de re* reading much more salient than the *de dicto* one:

(22) John believes that a girl (an old friend) arrived

Like our earlier work on presupposition (Lascares and Asher, 2009a), identifying the relative scopal position of the parathetical to the content of the discourse is a byproduct of identifying how it connects to the discourse with a coherence relation, but this is partly constrained by *linguistic form*: the content of the parathetical must be attached with a coherence relation to a discourse unit that outscopes the constituent that the parathetical (syntactically) modifies. Moreover, the proposition that the parathetical expresses is determined by the constituent that it modifies.

In (21), for instance, the parathetical introduces content that attaches with the coherence relation *Background* to the *conditional*; in other words to a piece of content that semantically outscopes the content of the constituent that the parathetical modified in syntax. Nevertheless, the proposition that the parathetical expresses seems to be determined entirely by syntax. Note that the parathetical conveys the proposition *Mary assures us that the party is over*, and not *Mary assures us that if the party is over, then we should leave*. So the proposition expressed must be derived from the parathetical and what it modifies in syntax. But in (21), this proposition outscopes the conditional, even though syntactically it is embedded within it.

Paratheticals don’t always project from embeddings, just like presuppositions don’t always do this. Whether they do or not depends on inferences about the rhetorical role that the parathetical plays in the discourse overall (subject to the constraint that it must rhetorically connect to a unit that outscopes the constituent it modifies). For example, “make that infamous” in (0903) doesn’t project out from an embedding in a belief context:

(23) John believes that he is famous—make that infamous—in some circles.

The parathetical “make that famous” essentially signals a self repair; roughly put, interpret *famous* as *infamous*. A self repair is a kind of coherence relation, and it is one that imposes

its own quite strict constraints on its scopal position in the discourse: it must sit in the discourse structure in the same scopal position as the thing it repairs. So in 0903, inferring that *make that infamous* repairs *famous* and inferring that the content of this paranthetical doesn't project out from the embedding in the belief context are mutually dependent. So the speaker of (23) does not publicly commit to the content that he is infamous, although he does commit to the content that John believes he is infamous.

I am struggling to write down a pragmatically resolved representation of *make that famous* though, because “that” needs to denote the predicate symbol *famous* and not the whole predication *famous(e, x)*. That's because otherwise the elementary predication *in(e, c)* that's introduced by the intersective modifier “in some circles” won't have the right semantic effects (i.e., the event *e* in *in(e, c)* must be the same event that's an argument to the predicate symbol *infamous_rel*).

10 Relative Clauses

- [1001] The dog they adopted barks.
- [1002] The dog that barked disappeared.
- [1003] The dog chased by the cat barked.
- [1004] The dog they thought we admired barks.
- [1005] Money to buy the dog arrived.
- [1006] A dog arrived that barks.
- [1007] A dog arrived whose owner disappeared.
- [1008] A dog and a cat arrived that belong to her.
- [1009] He admired a dog yesterday that barks.
- [1010] Nothing happened, which surprised him.
- [1011] It rained yesterday, which troubled me.

The relatives in (1010) and (1011) are interesting. If you make the relative clause into a scopal sentential modifier then the predicate symbol for trouble takes a scopal argument that will (perhaps barring quantifiers) outscope the proposition expressed by the matrix clause. This seems like the right move for psych verbs like *trouble*, *surprise*, *frighten*, *bore* etc. These predicates should take an individual variable (the experiencer) and a *proposition* as arguments, so that you can capture the different readings of sentences like *Every girl danced, which troubled me* (the reading where each girl dancing troubled you has “every” outscoping “trouble”, and the fact that it was all girls rather than some of them dancing that troubled you has “trouble” outscoping “every”).

Further, I think that psych verbs act like attitude verbs: when reporting them the speaker has the capacity to publicly commit *himself* to the way an object is described, or that description is ascribed to the experiencer of the psychological state. You can't get that distinction if the argument to psych verbs is an event variable as opposed to a proposition.

So, the ULF I would want for (1011) is something like (24), again ignoring quantifiers:

- (24) $l_1 : rain(e_1)$
 $l_1 : temp_loc(e_1, x)$

$l_2 : \text{yesterday}(x)$
 $l_3 : \text{trouble}(e_3, h_1, me)$
 $h_1 =_q l_1$

Making these relatives into scopal S modifiers is no doubt highly problematic for other reasons! But I think the above semantics captures what you would need for getting the right interpretations at the discourse level.

References

- N. Asher. Truth conditional discourse semantics for parentheticals. *Journal of Semantics*, 17(1):31–50, 2000.
- N. Asher and A. Lascarides. Questions in dialogue. *Linguistics and Philosophy*, 23(3):237–309, 1998.
- N. Asher and A. Lascarides. *Logics of Conversation*. Cambridge University Press, 2003.
- N. Asher, D. Hardt, and J. Busquets. Discourse parallelism, scope and ellipsis. In *Proceedings of Semantics and Linguistic Theory (SALT VII)*, Stanford, 1997.
- D. Beaver. *Presupposition and Assertion in Dynamic Semantics*. CSLI, 2001. Revision of 1995 PhD Thesis, University Of Edinburgh.
- S. Boër and W. Lycan. A performadox in truth conditional semantics. *Linguistics and Philosophy*, 4:71–100, 1980.
- J. Bos. Predicate logic unplugged. In *Proceedings of the 10th Amsterdam Colloquium*, pages 133–143, Amsterdam, 1995.
- R. Cann, R. Kempson, and L. Marten. *The Dynamics of Language*. Elsevier, Oxford, 2005.
- R. Cooper. How to do things with types. Technical report, University of Gothenburg, 2014.
- A. Copestake, A. Lascarides, and D. Flickinger. An algebra for semantic construction in constraint-based grammars. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics (ACL/EACL 2001)*, pages 132–139, Toulouse, 2001.
- M. Dalrymple, S. Shieber, and F. Pereira. Ellipsis and higher-order unification. *Linguistics and philosophy*, 14(4):399–452, 1991.
- M. Egg, A. Koller, and J. Niehren. The constraint language for lambda structures. *Journal of Logic, Language, and Information*, 10:457–485, 2001.
- J. Ginzburg. Resolving questions part i. *Linguistics and Philosophy*, 18(5):459–527, 1995.
- J. Ginzburg. *The Interactive Stance: Meaning for Conversation*. Oxford University Press, 2012.
- J. Ginzburg and I. A. Sag. *Interrogative Investigations: The Form, Meaning and Use of English Interrogatives*. CSLI Publications, 2001.

- H. P. Grice. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics Volume 3: Speech Acts*, pages 41–58. Academic Press, 1975.
- J. Groenendijk. Questions and answers: Semantics and logic. In *Proceedings of the 2nd CologNET-ElsET Symposium. Questions and Answers: Theoretical and Applied Perspectives*, pages 16–23, 2003.
- J. Groenendijk and M. Stokhof. Semantic analysis of wh-complements. *Linguistics and Philosophy*, 5(2):175–233, 1982.
- J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14: 39–100, 1991.
- J. R. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142, 1993.
- J.R. Hobbs. Coherence and coreference. *Cognitive Science*, 3(1):67–90, 1979.
- H. Kamp and U. Reyle. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Publishers, 1993.
- L. Karttunen. Implicative verbs. *Language*, 47(2):340–358, 1971.
- A. Koller and A. Lascarides. A logic of semantic representations for shallow parsing. In *Proceedings of the 12th Conference of the European Chapter of Computational Linguistics (EACL)*, pages 451–460, 2009.
- A. Lascarides and N. Asher. Imperatives in dialogue. In P. Kuehnlein, H. Rieser, and H. Zeevat, editors, *The Semantics and Pragmatics of Dialogue for the New Millenium*. Benjamins, 2004.
- A. Lascarides and N. Asher. Agreement, disputes and commitment in dialogue. *Journal of Semantics*, 26(2):109–158, 2009a.
- A. Lascarides and N. Asher. The interpretation of questions in dialogue. In A. Riester and T. Solstad, editors, *Proceedings of Sinn Und Bedeutung 13*, volume 5 of *SinSpeC. Working Papers of the SFB 732*, pages 17–30. University of Stuttgart, 2009b.
- A. Lascarides, A. Copestake, and E. J. Briscoe. Ambiguity and coherence. *Journal of Semantics*, 13(1):41–65, 1996.
- B. Partee. Nominal and temporal anaphora. *Linguistics and Philosophy*, 7:243–286, 1984.
- T. Reinhart. Point of view in language—the use of parentheticals. In G. Rauh, editor, *Essays on deixis*, pages 169–195. Gunter Narr Verlag, 1983.
- M. Rooth. A theory of focus interpretation. *Natural Language Semantics*, 1(1):75–116, 1992.
- D. Schlangen. *A Coherence-based Approach to the Interpretation of Non-Sentential Utterances in Dialogue*. PhD thesis, University of Edinburgh, 2003.

- B. L. Webber, A. Knott, M. Stone, and A. Joshi. Discourse relations: A structural and presuppositional account using lexicalised tag. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL99)*, pages 41–48, Maryland, 1999.
- D. Wilson. *Presuppositions and Non Truth Conditional Semantics*. Academic Press, New York, 1975.