

Natural Language & How we Use it

Critical Notice of Soames, *Philosophical Essays*, vol 1.

Ofra Magidor

[Penultimate Draft – forthcoming in Analysis]

1. Introduction

Soames (2009) is a superb collection of essays by one of the world's leading philosophers of language. The collection contains fifteen essays published over the last thirty years, with four previously unpublished essays, and a helpful introduction.

The essays are divided into five parts: Presupposition, Language and Linguistic Competence, Semantics and Pragmatics, Descriptions, and Lessons for Legal Interpretation. This division may seem a little puzzling at first with different parts covering topics of very different scope, but the essays in each part cohere very well together making the division entirely justified (perhaps with the one exception of Part 5, which contains only one essay and could have incorporated into Part 3). In fact, one advantage of this collection is in bringing together essays that relate to each other, offering the reader an excellent overview of Soames's views. The essays within each section seem to be ordered mostly by publication date, and in some places a different ordering would have been preferable. For example, Essay 9 is quite hard to follow without first reading Essay 10, so I would have preferred the reverse order.

Two features are particularly noteworthy throughout the collection. First, even the essays that were written a few decades ago are still entirely relevant, and researchers in the field would no doubt be rewarded from revisiting them. Second, Soames manages to skilfully weave together detailed and careful discussion of specific phenomena with insights on broad foundational issues. Both features are particularly well demonstrated by Part 1 of the collection, on which I have chosen to focus my discussion. I will begin, though, by making some very brief comments on Parts 2 and 3.

2. Linguistics and Psychology

The central thesis defended in Part 2 is that linguistics is a fundamentally independent branch from psychology, and that linguistic theories of syntax and semantics should not be understood as the psychologically real mechanisms employed by competent speakers.¹

It is worth noting that the importance of this thesis has become particularly vivid in recent years. This is due to new empirical work which suggests that many mainstream linguistic theories may not be adequate if construed as theories of psychological processing. For example, advancements in the in field of machine learning in computer science have yielded impressively successful algorithms for various natural language processing tasks. What is interesting about these algorithms is that they apparently receive no more linguistic input than a child learning a language, and apply only very general learning techniques, rather than any specific linguistic knowledge. This seems to undermine one of Chomsky's chief arguments for the framework of Universal Grammar: the Poverty of Stimulus Argument. If linguistic competence can be acquired by a computer program without building into it an innate grammar, then there is no reason to think that humans cannot acquire language in a similar fashion. But this in turn undermines one of the main reasons for believing that Chomskian linguistics represents a psychologically real theory.² All this is highly controversial, but the general point stands: if linguistics is construed as a branch of psychology, then its theories can principle be refuted by empirical evidence concerning processing. If Soames is right, on the other hand, linguistics need not be troubled by such empirical results.

I am generally sympathetic to Soames's thesis, but I don't find his arguments entirely convincing. One argument, for example, relies on the claim that semantic facts are dependant on one's environment (semantic externalism), but psychological states are internal to the subject (psychological internalism). But both claims can be contested: one might (as many linguists do) construe semantics as an entirely internalist theory, strictly speaking applicable only to idiolects of single subjects. On the other hand, those sympathetic to semantic externalism will naturally construe psychological states as having wide contents, and hence themselves externalist.³ Another problem is that since many of Soames's arguments focus on semantic notions such as truth-conditions and entailment, it is hard to see why his conclusion should apply to syntax as well.

I would have also liked to see Soames offer a more detailed positive proposal concerning what linguistics *is* ultimately about. Specifically, in what way is linguistics constrained by

¹ See especially Essay 3 and 4.

² For excellent discussions of this issue, see Lapin (2005) and Lapin & Shieber (2007).

³ Soames acknowledges the latter point in the introduction (footnote 6 on p. 5).

psychology? (It is hard to imagine that psychology poses *no* constraints on linguistics, or else there would be no reason to link semantics to any linguistic behaviour of speakers).⁴ Also, given Soames's thesis, how do we motivate some of the fundamental constraints that are often placed on adequate linguistic theories, such as the requirement of compositionality? After all, such constraints are standardly motivated on psychological grounds, which Soames will obviously want to reject.⁵

3. Meaning, Assertion, and Pragmatic Enrichment

In Part 3, Soames defends a novel theory of the relationship between meaning, semantic content, and assertion. On the traditional picture, when a speaker utters a sentence *S* (literally and assertively) the speaker thereby asserts the semantic content of *S*. Soames offers a substantial revision of this picture: first, speakers can assert more than one proposition by an utterance of a sentence *S*. Second, it may be that the semantic content of *S* is not among those propositions that the speaker asserts. Third, the semantic content of *S* may not be a complete proposition. The semantic content of *S* does not determine, but merely constrains the propositions asserted by its utterances: the primary proposition asserted is always a *pragmatic enrichment* of the semantic content, and other propositions asserted are a priori relevant entailments of the primary proposition. The meaning of a sentence is taken to be the 'least common denominator': what is common to all contents that can be asserted by (literal) utterances of the sentence.

Soames justifies this revision by what is essentially an 'inference to the best explanation' argument: he demonstrates how the revised conception can help address an impressive range of well-known puzzles in semantics, including Frege's puzzle, Kripke's puzzle about belief, the problem of incomplete definite descriptions, and the problem of interpreting bare numerical quantifiers ('*n* Fs'). Consider for example Frege's puzzle. According to Soames, an assertion of 'Hesperus is bright' has a content which predicates brightness of the *D* who is Hesperus (where *D* is the descriptive material the speaker associates with 'Hesperus'). If the speaker associates different descriptive material than *D* with 'Phosphorus', they would assert a different content by uttering the sentence 'Phosphorus is bright'. Nevertheless the semantic content of both sentences is the same: since there is no single common description that different speakers associate with 'Hesperus' (/ 'Phosphorus'), for each of the sentences, the only common content asserted by its utterances is the singular proposition that predicates brightness of Hesperus.

⁴ Soames does maintain that purely logical connections between sentences supervene on psychological facts (p. 164), but this is a very weak constraint, and at any rate we are not told why this constraint should hold.

⁵ Interestingly, Soames views some such constraints as ones that linguistics places on psychology (see pp. 175-77), but this leaves open the question of their justification.

I cannot do justice to Soames's intriguing proposal in this short space, but let me raise three brief points. First, many of the applications Soames suggests are not entirely problem free. For example, in discussing Kripke's puzzle, Soames maintains that when Pierre utters the sentence 'I don't believe that Paderewski is musical', the semantic content of this sentence is not among the propositions Pierre asserts. But now imagine that Paula (who associates a different description with 'Paderewski') utters the sentence 'Pierre doesn't believe that Paderewski is musical'. Intuitively, it would be correct to claim that Pierre and Paula both said the same thing. But on Soames's proposal, there seems to be no shared asserted content one could correctly attribute to both speakers.⁶

Second, the notion of meaning as 'least common denominator' requires further clarification. Note first that since Soames accepts that in some cases the semantic content is not among the propositions asserted by uses of a sentence, meaning is not after all a common dominator (let alone the least common denominator) of such assertions. In Essay 9, Soames recognises this problem, and proposes instead that meaning be a propositional matrix *pm*, such that all (relevant) assertions of *S* are assertions of "an acceptable completion of *pm*".⁷ But this raises the worry that the semantic content of *S* might turn out a lot more austere than even Soames wants to allow. I take it that all assertions of 'Jill is tall' involve acceptable completions of the propositional matrix involving nothing but the property of tallness, but this propositional matrix is clearly not the meaning of the sentence. Relatedly, since on Soames's framework when one asserts *P*, one also asserts a priori relevant entailments of *P*, then assertions of 'Jill is tall' always involve assertions of the proposition that someone is tall. Why then, shouldn't we count the proposition that someone is tall as the semantic content of the sentence?

Finally, it is worth noting that for any case where Soames appeals to pragmatic enrichment, one can formulate a parallel indexical theory: one that replaces the pragmatic enrichment with a context sensitive variable, and remains entirely within the framework of traditional semantics. Soames acknowledges this possibility in places, but objects to it on the grounds that "with such additions, context become very abstract, and to find out what context...is relevant to understanding a speaker's remark, a hearer would often *first* have to find out what the speaker is saying. Because of this, one who adopts this extension of contextual semantics can't *explain* how speakers-hearers interpret what is said by...applying their knowledge of semantics".⁸ What I find puzzling about this complaint is that it seems at odds with Soames's strong insistence on separating semantics from psychology. It would appear that given such a strict separation, one need not be worried by a semantic theory that does not correctly characterise the order of processing.

⁶ It's true that there are some very weak propositions that both will have asserted (e.g. that there is something that Pierre doesn't believe), but these do not seem to be the shared content one is attributing here.

⁷ p. 260.

⁸ P. 290. See also footnote 3 in p. 280.

4. Presupposition

The two essays on presupposition, which occupy about a quarter of the book, are perhaps the greatest gems in this collection. As we will see, despite being written 30 and 20 years ago respectively, the issues they raise are still very much pressing.

4.1 Presupposition Projection

In Essay 1 ('A projection problem for speaker presupposition'), Soames presents a forceful argument against compositional theories of presuppositions such as those relying on trivalent logics, and the prominent theory of Karttunen and Peters (henceforth KP).⁹

According to KP, each expression A has two dimensions to its conventional meaning: A 's contribution to the standard truth-conditions of sentences of which it is a constituent (denoted by A^e), and A 's contribution to the presuppositions (or in KP's terminology 'conventional implicatures') generated by those sentences (denoted by A^i). Within this framework, KP offers a powerful theory of presuppositions which is both foundationally solid (it is based on a view of presuppositions as a second conventional dimension of meaning, grounded in Grice's notion of conventional implicature), as well as producing an impressive range of correct descriptive predictions.

Soames's argument consists of two related claims: first, despite its apparent success, KP fails to be fully descriptively adequate; second, in order to accommodate the cases that KP fails to handle, one would have to substantially revise KP's understanding of presuppositions, thus undermining the foundational basis of the theory.

The issue Soames focuses on is that of presuppositions of disjunctions, a notoriously difficult problem in the theory of presupposition. According to KP, the presupposition of a disjunction is computed as follows: $(A \vee B)^i = (A^e \vee B^i) \wedge (B^e \vee A^i)$

Consider, for example,

- (1) Either there is no king of France or the king of France is in hiding.

Assuming that the first disjunct presupposes that there is a king of France, and the second generates no presuppositions (a sentence A can be represented as generating no presuppositions by letting A^i be some tautology T), then according to KP the presuppositional component of (1) will be:

⁹ Karttunen & Peters (1979).

(2) (There is no king of France \vee There is a king of France) \wedge (The king of France is not in hiding \vee T)

Since (2) is a tautology, it follows that (1) generates no presuppositions, which is precisely the right prediction.

Similarly,

(3) Either the king of France is in hiding or there is no king of France

is predicted to generate a tautological presupposition. It is worth noting that this latter prediction is more controversial (it is not clear that current utterances of (3) are felicitous), and indeed many theorists have opted instead for the following asymmetric projection rule for disjunction: $(A \vee B)^i = (A^e \vee B^i) \wedge A^i$, which predicts that (3) generates the presupposition that there is a king of France.¹⁰

Soames raises two problems for the KP disjunction projection rule. The first involves cases where the disjuncts have a common presupposition. Consider for example

(4) Either the chairman of the department isn't a sceptic or the chairman of the department isn't a very good philosopher.

Both disjuncts share the presupposition that the department has a chairman, and intuitively (4) retains this presupposition. However, according to the KP projection rule, this presupposition is cancelled.

While correct, this problem is not very serious. First, as Soames himself notes, the problem can easily be solved by amending the KP symmetric projection rule to contain a third conjunct as follows: $(A \vee B)^i = (A^e \vee B^i) \wedge (B^e \vee A^i) \wedge (A^i \vee B^i)$. Second, it is worth noting that the asymmetric projection rule mentioned above does not suffer from this problem.

The second problem Soames raises for the KP rule runs a lot deeper. This time, the problem concerns disjuncts with conflicting presuppositions. Consider the disjunction

(5) Either Bill met the king of Slobovia or Bill met the president of Slobovia

relative to contexts where it is common ground that no country has both a king and a president. The first disjunct generates the presupposition that Slobovia has a king, and the

¹⁰ For a good discussion of symmetric vs asymmetric projection rules see Rothschild (2008).

second that Slobovia has a president, which conflict. The KP projection rule predicts that the presupposition of the disjunction will be

- (6) (Bill met the king of Slobovia \vee Slobovia has a president) \wedge (Bill met the president of Slobovia \vee Slobovia has a king)

The problem is that (given the contextual assumptions) this is equivalent to ‘Either the Bill met the king of Slobovia or Bill me the president of Slobovia’. This means that the sentence presupposes exactly what it asserts, which seems wrong: a speaker can felicitously utter the disjunction without already taking its content for granted, and Soames argues convincingly that this intuition cannot be easily explained away as a case of obtaining information via accommodation, because that would fail to explain why embeddings of the disjunction as in

- (7) If Bill either met the king of Slobovia or the president of Slobovia then he’ll have great stories to tell.

do not encourage us to accommodate the context with the material in the antecedent.

It is important to recognise that amending the symmetric KP rule as suggested above will not help here (that will ensure that (5) doesn’t presuppose *only* what it asserts, but it will continue to presuppose what it asserts). Also, the aforementioned asymmetric projection rule does no better (it predicts that the sentence presupposes that Bill met the King). Soames also rejects another possibility: disjoining the KP projection rule with the condition $\neg(A^i \wedge B^i)$. While this will block (5) from presupposing what it asserts, it will also block any *other* presuppositions in these kinds of disjunctions, even presuppositions that intuitively ought to filter through. For example, in

- (8) Either Bill met the king of Slobovia and insulted him *too* or Bill met the president of Slobovia

the presupposition that Bill insulted someone else will be incorrectly blocked.¹¹ This is indeed a serious problem, and it is striking that despite being raised by Soames thirty years ago, most contemporary theories of presuppositions offer projection rules for disjunction that still suffer from this problem.

What then can be done? One possibility that Soames does not consider is to amend the KP apparatus so that instead of generating a single presupposition, each sentence will generate a *set* of presuppositions, and where the projection rules are amended accordingly. The simplest way of doing this is to systematically replace any rule of the form $(A*B)^i = f(A^i, B^i, A^e, B^e)$,

¹¹ Note that Soames assumes here that the first disjunct presupposes that Bill insulted someone else, rather than that *if* Bill met the king, the insulted someone else (as KP predicts). I ignore this complication in my discussion.

where $*$ is some two-place propositional connective and f the function that computes the relevant presuppositional component, with: $(A*B)^i = \{p \in A^i, q \in B^i: f(p, q, A^e, B^e)\}$. This will ensure that if the original projection rule was $(A \vee B)^i = \neg(A^i \wedge B^i) \vee ((A^e \vee B^i) \wedge (B^e \vee A^i) \wedge (A^i \vee B^i))$, then the relevant presupposition of (8) will not be entirely blocked. One would of course need to be convinced that this proposal works across the board, but if it does, then it is a fairly small departure from the original KP system, and one that does not undermine its foundational underpinnings.

Soames proposes a different solution to the problem. He argues that the heart of the issue lies in the interaction between presuppositions and conversational implicatures. The idea is that when an utterance of a sentence generates a presupposition that conflicts with a conversational implicature, the requirement posed by the conversational implicature trumps the one posed by the presupposition, and cancels it. Thus for example, since in uttering (5) a speaker conversationally implies that (5) is not already taken for granted, the problematic presupposition is cancelled. Soames also argues that merely keeping the original KP projection rules while allowing conversational implicatures to trump presuppositions is insufficient to address the problem. Instead, he offers an entirely different conception of presupposition projection: one according to which *all* presuppositions of the constituents of a sentence are projected (thus eliminating the need for the compositional projection rules), and where all putative cases of cancellation are explained via trumping by conversational implicatures.¹² Consider for example a classic case of presupposition cancellation:

(9) If John has a son, then his son is bald.

As opposed to KP, Soames maintains that (9) retains the presupposition that John has son. However, since by uttering the conditional one also generates the conversational implicature that the antecedent is not taken for granted (otherwise, it would have been pointless to utter the conditional, rather than its consequent), the presupposition need not (and in fact must not) be taken for granted in this context.

Soames's proposal is certainly very elegant, but unfortunately it too suffers from serious problems. One general concern with the proposal is its foundational underpinnings. Soames rightly argues against a hybrid view that combines the KP projection rules with trumping, that it does not sit well with the KP conception of presuppositions as conventional implicatures.¹³ However, a similar worry applies to Soames's own proposal: if presuppositions are neither conventional implicatures nor part of asserted content, one would like to know more about what exactly they are, and how they fit with our general linguistic picture.

¹² It is worth noting that a very similar proposal was developed independently in Gazdar (1979). (See Soames (2009), p. 58, footnote 62).

¹³ See pp. 55-56.

The theory also faces some serious descriptive problems. Let us return to (5), the very example that was used to motivate Soames's proposal. Soames is indeed able to explain why the obvious presuppositions of the disjuncts (that Slobovia has a king, and that it has a president) are ultimately cancelled. The problem, however, is that Soames also wants to accommodate the (correct) prediction that (5) generates the presupposition:

(10) Either Slobovia has a king or Slobovia has a president.¹⁴

He argues that each of the sentence's (trumped) presuppositions entails (10), but that since (10) is no conflict with any conversational implicatures of (5), it is not trumped. The problem with this reasoning is that Soames is implicitly assuming here that any entailment of a presupposition is itself a presupposition. While harmless in the KP system, this is an assumption Soames cannot make. Consider for example the conditional (9). The sentence presupposes, according to Soames, that John has a son. If any entailment of a presupposition was itself a presupposition, then the sentence would also presuppose

(11) Either John has a son or spiders have eight legs.

But clearly, someone uttering (9) need not take (11) for granted, and this cannot be explained via trumping (there is no reason to think that someone uttering (9) must not take (11) for granted).¹⁵

Soames himself raises other cases that are correctly handled by KP, but not by his system. For example:

(12) If someone at the conference solved the problem, it was Julius who solved it.

The difficulty is that even if the speaker must not take for granted that someone at the conference solved the problem, there is no reason why they cannot take for granted the weaker presupposition of the consequent (namely, that someone solved the problem).

Throughout the paper Soames offers an insightful analysis of the principled difference between his proposal and systems such as KP. We can divide the cases where a complex sentence S contains a constituent with a presupposition P, to ones where the speaker

(i) must take P for granted (as in 'If the king is here, we will bow')

¹⁴ Generating this presupposition is also a problem for the revised KP system I proposed above, but can be addressed by amending the projection rule to say: $(A \vee B)^i = (\neg(A^i \wedge B^i) \vee ((A^e \vee B^i) \wedge (B^e \vee A^i))) \wedge (A^i \vee B^i)$.

¹⁵ Soames's other justification for why (5) presupposes (10) is even more suspect: he notes that (5) entails (10) (footnote 65, p. 59). But surely, not every entailment is a presupposition.

(ii) must not take P for granted (as in (5))

(iii) is permitted to take P for granted, and is permitted not to (as in (12))

Soames explains that while both KP and his system handle cases of type (i) fairly smoothly (that is, they both have correct predictions when the presuppositions *do* filter through), his system has an advantage in handling cases of type (ii) (because when a presupposition is trumped, the speaker must not to take it for granted), and KP has an advantage in cases of type (iii) (because these are cases where the presupposition is cancelled, without being trumped).

While I find this analysis extremely helpful, I'm not sure it fully captures the differences between the two systems. There are some other cases of type (i) that are handled by Soames's theory more successfully than the KP system. One such case (that has not, to my knowledge, been previously discussed in the literature) is the case of disjunctions with one trivially true disjunct that generates no presuppositions. Consider for example

(13) Either London is the capital of England or the king of France is bald.

According to the KP rule, the sentence generates the presupposition

(14) (London is the capital of England \vee There is a king of France) \wedge (The king of France is bald \vee T)

which is trivially satisfied in any context in which it is taken for granted that London is the capital of England. On the other hand, Soames's theory, which takes a more liberal approach to filtering, predicts that (13) presupposes that there is a king of France, which strikes me as the right prediction. On the other hand, Soames has trouble with some cases of type (i) where the complex sentence generates a presupposition which is a manipulation of the presuppositions of its constituents, rather than a direct inheritance. This is demonstrated by the problem I raised above in accounting for why (5) presupposes (10).

4.2 Overview

Essay 2 ('Presupposition'), contains an extremely rich overview of the main approaches to presupposition and some of the challenges each of them faces both on the foundational and the descriptive side. Soames distinguishes between three fundamentally different approaches to presupposition: the logical approach, the expressive approach, and the pragmatic approach. According to the *logical approach*, a sentence S presupposes a proposition P if relative to any world in which S is true or false, P is true. According to the *expressive approach*, S presupposes P if S only expresses a proposition in contexts in which P is true. It is initially tempting to think of the expressive approach as a subspecies of the logical approach: one reason a sentence S might fail to be true or false when P is false, is that it fails to express a

proposition. However, I think Soames is right to distinguish the two approaches. Assume, for example, that any sentence containing the description ‘The queen of England’ presupposes that there is a queen of England, on the expressive conception. If one utters the sentence ‘The queen of England doesn’t exist’ in current contexts, the presupposition is satisfied and the sentence expresses a proposition Q . Now consider a possible world w relative to which there is no Queen of England. It is at least an open possibility for the expressive approach to maintain that relative to w , Q is simply false, and correspondingly so is S . On the other hand, according to the logical conception, S cannot be false in w , because P is false in w .¹⁶ According to the *pragmatic approach*, presuppositions are constraints that sentences (or utterances thereof) place on the common ground assumptions of conversational participants. Thus characterised, the pragmatic approach is very general, and encompasses many different views of presupposition. In particular, on this characterisation, a proposition being a pragmatic presupposition does not preclude it from also being a logical presupposition.

One interesting issue that arises in this connection (and which Soames discusses at length) is Stalnaker’s proposal that the pragmatic requirements posed by (at least some) pragmatic presuppositions are explained precisely via their being logical presuppositions. Stalnaker’s idea is that when a sentence S is asserted, one ought to keep in the context-set those worlds in which S is true and remove from the context-set those worlds in which S is false. Thus if relative to a world w in the context-set a logical presupposition P of S is false, S will be neither true or false, leaving us with no instruction as to whether to keep w in the context-set or remove it. This explains, according to Stalnaker, why a context in which a logical presupposition P is not common ground (namely, P is not true in all worlds of the context-set) is also pragmatically defective.

Soames argues against Stalnaker’s proposal as follows. First, he draws a key distinction between three ways in which a sentence might fail to be true or false. The first way is the case of expressive failure, where the sentence does not express a proposition. The second is where a sentence expresses a proposition but this proposition is neither true nor false. The third is where the sentence expresses a proposition, but this proposition is somehow indeterminate, in the sense that it cannot be said to be true nor said to be not true (following certain analyses of sentences containing vague or partially-defined predicates). It is important to realise that underlying this third option is the retention of the Tarskian truth-schema, which entails that although the proposition cannot be characterised as being true or false, neither can it be characterised as being neither true nor false.

¹⁶ See Soames (2009), p. 82.

Soames argues that on either of the first two ways for bivalence failure, Stalnaker's argument does not go through. This is because one can naturally reformulate the rules for context update so that one keeps a world w in the context-set if S is true relative to w , and one removes from the context-set all worlds relative to which S is either false or not true. I think this argument is correct, and can be reinforced by adding two points against each of the first conceptions. The first conception will be insufficient for Stalnaker's purposes, because there are many cases of apparent presupposition failure where the presupposition is satisfied relative to the actual world, but is nevertheless not common ground. For example, if in the actual world there is a queen of England but this is not common ground, then even given the expressive approach, 'The queen of England is tall' will succeed in expressing a proposition despite being pragmatically inappropriate. Regarding the second conception, it is worth noting that it is far from clear that this is a legitimate way for bivalence to fail: after all, on the plausible assumption that truth is governed by the Tarskian truth-schema, it is simply a blatant contradiction to say that a proposition is neither true nor false.¹⁷

The third conception does not suffer from the Soames's objection to the first two: if relative to a world w , we cannot determine that a proposition Q is true but cannot determine that it is not true, then indeed we would not know how to handle w . Soames argues, however, that this conception is not naturally suited to deal with the problem of presupposition. For one thing, it does not seem to exhibit the right kind of projection behaviour. For example, if 'bald' is a partially-defined predicate, then even in cases where John is a borderline case of baldness, 'It is likely that John is bald' or 'If John is bald, then Bill will like him' might still be perfectly true.

I doubt this is a knockdown argument: cases of presupposition failure might be only analogous to cases of partially-defined predicates in the kind of bivalence failure involved, but not fully mimic their projection behaviour. But the argument certainly demonstrates that thinking of presuppositions in terms of the third conception is not entirely natural. Moreover, it is worth keeping in mind just how hard it is to accept a framework where some propositions 'cannot be said to be true or false' but where this is not due to their *being* neither true nor false, nor is it due to any epistemic ignorance on our part. The difficulties of accepting such a framework might well lead us to reject the third conception altogether, and adopt bivalence across the board.

I think there is another fundamental problem with linking a pragmatic account with a logical-semantic one (a problem which Soames does not discuss). The problem is that the conception of context that underlies logical-semantic accounts is crucially different than that which

¹⁷ See Williamson (1994), p. 188. Soames also expresses some hesitation as to whether there are cases of this second kind (see p. 113).

underlies the pragmatic account. On the traditional understanding of logical-semantic accounts, the crucial notion of context is one that is linked to the actual world (at the time and location of utterance). On the pragmatic accounts, on the other hand, it is linked to the beliefs of the participants in the conversation. In particular, note that on the pragmatic approach the actual world need not be a member of the context-set.¹⁸ Now consider a conversation where as a matter of fact there is a queen of England, but where all conversational participants erroneously take it for granted that there isn't one. According to the natural interpretation of the logical account, an utterance of 'The queen of England is tall' does not suffer from presupposition failure: the relevant presupposition is in fact true, 'The queen of England' successfully denotes a particular woman (say Elizabeth), and the sentence is true or false depending on whether or not Elizabeth satisfies the predicate 'tall'. According to the pragmatic conception, however, the sentence *does* suffer from presupposition failure, because the presupposition is not true across the context-set (moreover, since it is false across the context-set, it cannot be easily accommodated). One could try to insist on a non-standard interpretation of the logical account, where the relevant presupposition is stipulated to be interpreted relative to the (pragmatic) context-set, rather than relative to the semantic context of utterance, but this seems to undermine the idea that the logical conception is independently motivated on semantic grounds, and that pragmatic presuppositions are explained via logical ones. Alternatively, one could insist on a non-standard semantics where 'The queen of England' does not semantically denote the actual unique queen of England, but whoever participants in the conversation *believe* is the queen of England. But this takes us into a radical revision of semantics, one that goes far beyond the scope of the current discussion.¹⁹ The upshot is that while sentences might in principle have both logical and pragmatic presuppositions (and moreover, the same presupposition might act both as a logical and a pragmatic one), the two conceptions are fundamentally different and play very different roles in our conversational practices.

Another theme in the essay, which links it to the previous one, is the issue of presupposition projection. Soames takes here a more pluralistic approach, which allows for three different ways in which a presupposition may in effect be cancelled. The first involves the standard compositional rules for presupposition projection, as in the KP system. The second involves cases where the negation of a presupposition P is explicitly added to the common ground, prior to an utterance of S, as in

(15) There is no king of France.

Followed by

(16) Therefore, the king of France isn't in hiding.

¹⁸ See Stalnaker (1999), p. 100.

¹⁹ It is worth noting in other work other work (e.g. Stalnaker (1978)), Stalnaker suggests a revision along these lines. For criticism of this semantic framework see Hawthorne & Magidor (2009).

The apparent presupposition of (17) (that there is a king of France) seems to be cancelled by (16), and an utterance of (17) is thus felicitous even though its presupposition is not taken for granted or accommodated.

Soames is certainly right that this kind of example cannot be easily explained via standard compositional accounts such as KP,²⁰ but I don't think his own explanation, which involves the presupposition being cancelled simply because its negation is added to the common ground, succeeds in accounting for this case.²¹ After all, if whenever a presupposition explicitly contradicted the common ground it would have been cancelled, we would never have observed the anomalies associated with presupposition failure. But it is precisely these anomalies which motivate the need for a theory of presuppositions in the first place. Moreover, the kind of cancellation exhibited by (16)-(17) does not seem to be replicated in contexts that do not involve negation. Consider for example

(17) Only one person came to dinner last night.

Followed by

(18) Bill came to dinner too.

The presupposition of (18) does not seem to be cancelled by (17). (Note that it is not easy to explain the anomaly of (18) on other conversational grounds). It seems then, that the phenomenon exhibited by (16)-(17) has something specific to do with negation, and not with a general mechanism for cancellation.

The third way in which cancellation can occur, according to Soames, is the one introduced in the previous essay, namely trumping of a presupposition by a conflicting conversational implicature. This time, though, Soames does not argue for trumping as the primary mechanism of cancellation, but focuses on special cases. Consider the sentence

(19) If I realise later I haven't told the truth, I will confess it to everyone.

In general, we expect the presuppositions of the antecedent of a conditional to filter through the conditional (indeed, this is what KP predicts). Interestingly, though, in the case of (19), the presupposition of the antecedent (namely that I haven't told the truth) seems to be cancelled. Soames explains this cancellation via an appeal to trumping by a conversational implicature. His idea is that by uttering (19), one generates the conversational implicature that one does not know its antecedent. But in most contexts, if you were taking it for granted that you haven't told the truth, then you would also be taking the antecedent of (19) for

²⁰ Unless one posits a separate negation operator that blocks presuppositions, as suggested by Karttunen & Peters (1979), pp. 47-48. For criticism of this proposal see Horn (1985).

²¹ See Soames (2009), p. 98 for the suggestion that this is a general cancellation mechanism.

granted, and hence the conversational implicature conflicts with the presupposition, ultimately cancelling it.

As with the previous example, I agree with Soames that this is a difficult case that cannot be easily accommodated by standard projection rules, but again I find that his own analysis falls short. To see why, consider the following example. Suppose I'm about to undergo an operation that will make me lose all my pre-operation memories. I know that after my operation my family might teach me some facts about my pre-operation life (some of which I used to know before and others which I didn't). Now I utter the following sentence:

(20) If after the operation I realise that before the operation my partner was cheating on me, I will be very upset.

An utterance of (20) in this context seems to me perfectly felicitous, even if I currently do not take it for granted (or even believe) that my partner is cheating on me. So the presupposition of the antecedent seems to be cancelled here, just as with (19). However, this time the conversational explanation offered by Soames will not work: even if it were taken for granted that my partner is cheating on me, that does not entail that I will realise this after the operation, because this may not be amongst the facts that I will relearn. Another example which seems to me to exhibit the same phenomenon as (19) and (20) is the following. Suppose we are at a lung cancer fundraiser. The host announces that anyone that recently stopped smoking will win a prize. None of us have any idea about the current or past smoking habits of John (who is not part of the conversation), and I say:

(21) If John recently stopped smoking, he will win a prize.

This utterance seems to me perfectly felicitous, even though it is not taken for granted in this context that John used to smoke. But again, this cannot be explained via a clash with conversational implicatures: we might have failed to know whether or not John stopped smoking even if we had known that John used to smoke. In all these examples, we are faced with a special context in which one focuses on the truth-conditional content of the antecedent, and largely ignores the presuppositions it normally generates. The phenomenon is certainly a puzzling one, but a mere appeal to conversational implicatures cannot account for it.

A final theme running throughout the essay that I would like to comment on is the 'explanation problem': the need to explain why the different connectives exhibit the projection behaviour that they do, by relying entirely on their standard semantic content. It is hard to overestimate the influence this problem had on thinkers in the field: it was one of the main motivations for the radical turn to dynamic semantics, and solving it is taken by many to be the benchmark of a successful theory of presupposition.²² Soames played a central role

²² See e.g. Heim (1983) and Schlenker (forthcoming).

bringing this problem to the forefront, and in this essay he makes an important contribution to the issue. He points out (pp. 119-120) that although Heim's dynamic system was motivated in a big part by a need to address this problem, it ultimately fails to answer it. Heim offers a semantic treatment of connectives such as 'and' which explains their projection behaviour, but her explanation does not rule out the possibility of alternative connectives which have the same truth-conditional effect as 'and', but with a different presuppositional projection behaviour.

The explanation problem has certainly played a central role in this field, but I am inclined to think it may have been given *too* central a role.²³ Underlying the problem is an assumption that there is a fundamental difference between the status of the basic presuppositional triggers, and the way in which the different connectives project presuppositions. As Soames puts it (p. 102): "[On KP], pragmatic presuppositional requirements are reduced to arbitrary linguistic conventions associated with lexical items and constructions. This is plausible for examples like 'even' or clefts; but it is implausible for other cases, particularly the connectives. Surely, there is some connection between the truth conditional content of the connectives and the pragmatic presuppositions of sentences containing them".

But this chasm between the connectives and the basic triggers is often exaggerated. Sure, the projection rules for connectives such as 'and' are not completely arbitrary, and have *something* to do with these connectives' truth-conditional contributions. But the connection between the truth-conditional and presuppositional contribution of many basic triggers does not seem to be completely arbitrary either.²⁴ After all, the fact that 'X stopped phi-ing' generates the presupposition that X used to phi (rather than, for example, that X plays baseball) is surely somehow related its truth-conditional meaning. Of course, there is still a difficult puzzle in explaining why 'stop' yields precisely this presupposition and not other equally plausible alternatives (e.g. that X no longer smokes). But as is seen from Soames's criticism of Heim, a similar problem arises in the case of the connectives. On the other hand, in both cases there is a need to explain why the same presuppositional behaviour seems stable across a wide range of languages. I suggest then, that the explanation problems for the connectives and the basic triggers are not fundamentally different. These problems are certainly interesting, but perhaps it is time they are repositioned to the field lexical meta-semantics, rather than playing a central role in our theories of semantics and pragmatics.²⁵

²³ This is not in particular a criticism against Soames, who explicitly downplays the importance of the problem (p. 120).

²⁴ Triggers like 'even' or 'too' are limiting cases because they have no truth-conditional contribution.

²⁵ Thanks to Cian Dorr for helpful discussion of this paper.

References

- Gazdar, G., *Pragmatics: Implicature, Presupposition, and Logical Form*, Academic Press, 1979
- Hawthorne, J. and Magidor, O., 'Assertion, context, and epistemic accessibility', *Mind* 118 (2009), pp. 377-397.
- Heim, I., 'On the projection problem for presuppositions', in Barlow, M., and Flickinger, D., (eds), *Proceedings of the Second West Coast Conference on Formal Linguistics* (1983), pp. 114–125.
- Horn, L., 'Metalinguistic negation and pragmatic ambiguity', *Language* 61 (1965), pp.121-174.
- Karttunen, L., and Peters, S., 'Conventional implicatures', *Syntax and Semantics 11* (1979), pp. 1-56.
- Lappin, S., 'Machine learning and the cognitive basis of natural language', *Proceedings of Computational Linguistics in the Netherlands 2004* (2005), pp. 1-11.
- Lappin, S., and Shieber, S., 'Machine learning theory and practice as a source of insight into universal grammar', *Journal of Linguistics* 43 (2007), pp. 393-427.
- Rothschild, D., 'Presupposition projection and logical equivalence', *Philosophical Perspectives* 22 (2008), pp. 473-497.
- Schlenker, P., 'Be articulate: A pragmatic theory of presupposition projection', *Theoretical Linguistics* 34 (forthcoming), pp. 157-212.
- Soames, S., *Philosophical Essays, vol. 1: Natural Language, What it Means & How We Use It*, Princeton University Press (2009).
- Stalnaker, R., 'Assertion', *Syntax and Semantic* 9 (1978), pp. 315-332.
- Stalnaker, R., *Context and Content*, Oxford University Press, 1999.