Forthcoming in: Keith Allan and Kasia Jaszczolt, ed. *The Cambridge Handbook of Pragmatics*, Cambridge.

Context and Content: Pragmatics in Two-Dimensional Semantics

Berit Brogaard

April 27, 2010

Abstract:

Context figures in the interpretation of utterances in many different ways. In the tradition of possible-worlds semantics, the seminal account of context-sensitive expressions such as indexicals and demonstratives is that of Kaplan's two-dimensional semantics (the content-character distinction), further pursued in various directions by Stalnaker, Chalmers, and others. This chapter introduces and assesses the notion of context-sensitivity presented in this group of approaches, with a special focus on how it relates to the notion of cognitive significance and whether it includes an intuitively plausible range of expressions within its scope. The chapter concludes with a discussion of the prospects of using two-dimensional semantics to account for context-sensitive expressions in dynamic discourse.

Table of Contents

0. Introduction

Two-dimensional semantics in its most general characterization is a semantics that recognizes two kinds of content: narrow content and wide content.[1] Narrow contents can take the form of linguistic meanings (e.g., functions from context to content) or descriptive Fregean contents. Wide content is a set of possible worlds or a structured Russellian proposition consisting of properties and/or physical objects. On some approaches (for instance, Kaplan's and Stalnaker's), "narrow content" is not semantic content (it is not what is shared in communication) but belongs to what Stalnaker calls 'the metasemantics'. On other approaches (e.g., Chalmers and Jackson), the narrow content is (part of) the semantic content.

Kaplan originally introduced two-dimensional semantics in order to account for the semantics of context-sensitive expressions, primarily indexicals and demonstratives. Names,

[1] I am using the terms 'narrow content' and 'wide content' to refer to kinds of contents thought to play different functional rules. For example, narrow content is likely to play a role in accounting for the lack of substitutability in intensional or hyperintensional contexts. Some will want to deny that narrow content, so defined, is truly narrow (internal) because they adhere to externalism.

indexicals and demonstratives have often been treated as having as their semantic content the individuals to which they refer. But indexicals and demonstratives do not acquire their content in the same way as names. Unlike names, as traditionally construed, their content varies across worlds, times, speakers, and locations. This difference between names and indexicals prompted Kaplan to introduce a further layer of meaning, viz. a function that takes an expression from worlds, times, speakers or locations to its content. This function can then be either constant (names) or variable (indexicals and demonstratives).

Two-dimensional semantics has subsequently been employed by Robert Stalnaker, David Chalmers, Frank Jackson and others as a tool for explaining how sentences that have necessary contents can be cognitively informative and how sentences that are cognitively uninformative can be contingent, and David Chalmers and Frank Jackson have developed the approach further in debates about the role of conceptual analysis and reductive explanation.

Here I will introduce and assess the notion of context-sensitivity presented by the various two-dimensional frameworks, with a special focus on how it relates to the notion of cognitive significance and whether it includes an intuitively plausible range of expressions within its scope. I will argue that the two phenomena (viz. context-sensitivity and cognitive significance) are, to some extent, inseparable. I will conclude with a discussion of the prospects of using epistemic two-dimensional semantics to account for context-sensitive expressions in dynamic discourse.

1. Kaplan

1.1. Kaplan's Two-Stage Theory

In "Demonstratives" David Kaplan introduced a two-stage semantic theory of indexicals and demonstratives. In Kaplan's framework, disambiguated expression types have a linguistic meaning or what Kaplan calls a 'character'. A character is a function from context to content. The context is a sequence of parameters which include (at least) a world, a speaker, a time and a location. In the extended 1989 version of the unpublished text Kaplan included an addressee among the contextual parameters. A demonstration is not directly included in the context as a

contextual parameter. Instead demonstratives associated with different demonstrations are treated as different lexical items.[2] The character of 'I am speaking now' is a function from a world w, a speaker s and a time t of the context to a content that is true just in case s is speaking at t in w. The content of a sentence is the proposition expressed and also what is said by the sentence in context. Context together with a disambiguated sentence type with a determinate character thus yields a proposition.

In Kaplan's original framework, only pure indexicals (e.g., 'I', 'now', and 'here'), true demonstratives (e.g., 'this', 'that') and perhaps a few other types of expression (including complex demonstratives) are context-sensitive expressions. They are the only expressions that have variable character. Noun phrases have constant character.[3] So, the word type 'water' used by Oscar on Earth where the clear potable liquid that comes out of faucets is H_2O and the word type 'water' used by Oscar's physical and phenomenal duplicate Twin-Oscar on Twin-Earth where the clear potable liquid that comes out of faucets is XYZ are different lexical items with different constant characters.

For Kaplan, contexts need not be real speech situations. Strictly speaking, they need not be possible either. A sequence of the actual world, the present author, 3 pm Canberra time, 1535, and The Lounge in Melbourne forms an improper context. However, the sorts of contexts that are relevant for determining contents are possible contexts.

There are two main reasons that contexts, for Kaplan, need not compose real speech situations. First, this notion of context allows that sentence types that are necessarily true or necessarily false relative to real speech situations can be assigned a different truth-value relative to different contexts. These include sentence types such as 'I am not speaking now', 'I do not exist now', and 'I am not here now'. Since there are contexts relative to which these

[2] Note that this move does not guarantee the context-insensitivity of demonstrative-plus-associated-demonstration, as the same demonstration arguably can demonstrate different objects on different occasions. As Kaplan characterizes 'demonstration', a demonstration is 'typically, though not invariably, a (visual) presentation of a local object discriminated by a pointing' (p. 490). Arguably, 'this' and a visual presentation might pick out different (visually indistinguishable) objects at different times. Parameters of the context (e.g. the time-parameter) are thus required to determine the content of the demonstrative-plus-associated-demonstration.

[3] An expression has a constant character if it has the same semantic value across all contexts.

sentence types express true propositions and contexts relative to which these sentence types express false propositions, these sentence types are not logical truths but contingent truths.

Second, if contexts were real speech situations, intuitively valid arguments would come as invalid. Consider the following argument:

If John is hungry now, then he is grumpy now.
John is hungry now.
So, John is grumpy now

Utterances take time. So, if contexts were real speech situations, then there would be contexts in which it is true that if John is hungry now, then he is grumpy now, and true that John is hungry now, but false that John is grumpy now because he eats something before we have a chance to utter the conclusion. Below I will sketch a way to deal with dynamic discourses in a two-dimensional framework.

1.2. Kaplan and Cognitive Significance

Kaplan's theory was intended to give an account of context-sensitive expressions. It was not intended as an account of cognitively informative necessities (Kripke's a posteriori necessities) or cognitively uninformative contingencies (Kripke's a priori contingencies). But Kaplan's framework provides a reasonably good explanation of the phenomenon when indexicals and demonstratives are implicated. For example, 'I am Brit (if I exist)' expresses, relative to my context, a necessary truth of the form 'a = a (if a exists)'. Yet it can be cognitively significant. We can explain its cognitive significance by noting that 'I' and 'Brit' have different characters. The character of 'I' is a function from a context to the speaker of the context, the character of 'Brit' is a function from a context to Brit. The identity claim, it may be said, is cognitively significant, because the expressions flanking the identity sign have different characters. If we discover that the different characters determine the same content in a context, our overall knowledge state has been enriched. Kaplan's original example of cognitive significance runs as follows:

> If I see, reflected in a window, the image of a man whose pants appear to be on fire, my behavior is sensitive to whether I think, 'His pants are on fire' or 'My pants are on fire', though the object of thought may be the same. (p. 533).

Thinking 'his pants are on fire' (referring to myself) and thinking 'my pants are on fire' will elicit different behavioral responses. For Kaplan, what explains the difference and hence what explains cognitive significance is the character of the sentence (the way the proposition expressed is presented), not the proposition expressed.

Kaplan's theory does not by itself offer a general explanation of cognitive significance. 'Hesperus is Phosphorus' is ordinarily cognitively significant, yet if 'Hesperus' and 'Phosphorus' are genuine names rather than disguised descriptions, then the character of 'Hesperus' is identical to the character of 'Phosphorus'. It is a constant one-place function mapping contexts to Venus.

Noting that Kaplan's theory does not offer a general explanation of cognitive significance is not a criticism of his framework thought of as a theory of indexicals and demonstratives. However, there is some reason to think that the correct story of indexicals and demonstratives ought to be generalizable to noun phrases. To see this it may be helpful to introduce the notion of 'twin-earthability'. David Chalmers defines the notion as follows:

> We can say that two possible individuals (at times) are twins if they are physical and phenomenal duplicates; we can say that two possible expression tokens are twins if they are produced by corresponding acts of twin speakers. Then a token is Twin-Earthable if it has a twin with a different 2-intension [Russellian intension]. (2006: section 3.5)

A twin-earthable expression is one which has a twin token with a different 2-intension (or Kaplan content). Note that, on Chalmers' definition, it is not required that the word tokens spoken by the twin-speakers be tokens of the same lexical item.

Twin-earthability plays a crucial role in testing whether an expression is semantically neutral. Semantically neutral expressions, roughly, are the non-twin-earthable expressions (e.g. 'phenomenally conscious', 'cause', and 'friend'). To a first approximation, they are expressions

that do not have wide content (Kaplan content) that does not supervene on the physical-phenomenal make-up of the particular speakers (deferential uses aside). Twin-earthable expressions are required to generate (ideally) cognitively significant sentences.[4] Semantically neutral expressions cannot be used to generate (ideally) cognitively significant sentences.

Note that in order for twin-earthability and semantic neutrality to be relevantly connected, it must be true that narrow content (1-intension) supervenes on the physical-phenomenal make-up of particular speakers.[5] If it does not, then almost any expression (semantically neutral or not) comes out as twin-earthable. Take 'phenomenally conscious'. There is a scenario in which people speak the language Schmenglish. In Schmenlish 'phenomenally conscious' has a somewhat different 1-intension (and hence 2-intension). It means, roughly, *alert*. So, twins on Earth and Schmenglish Earth make corresponding acts and say 'I am conscious', their tokens pick out different properties. So 'phenomenally conscious' comes out as twin-earthable, which was not the result we wanted.

Assuming Kaplan's original framework, indexical and demonstrative expressions are both twin-earthable and context-sensitive. If Oscar and Twin-Oscar both say 'I am human', the twin tokens of 'I' have different 2-intensions (Kaplan contents). One picks out Oscar at every world, the other picks out Twin-Oscar at every world. Likewise, if Oscar and Twin-Oscar demonstrate the clear potable liquid in their drinking glasses and say 'that is water', the twin occurrences of 'that' have different 2-intensions. One picks out H_2O at every world, the other picks out XYZ at every world.

Given Kaplan's original framework, most noun phrases are twin-earthable but not context-sensitive. 'Water' is twin-earthable, because in Kaplan's framework, Oscar and Twin-Oscar's tokens of 'water' pick out different chemical substances. But 'water' is not context-sensitive because its character is a constant function from contexts to contents.

---

[4] I say 'ideally', because we are idealizing away psychological and physiological limitations (e.g. as regards the utterance or comprehension of infinitary sentences) and logical and mathematical non-omniscience.

[5] We need some sort of qualification to rule out Eden worlds, i.e., worlds where color properties are primitive or 'purely qualitative' instantiated properties. In Eden worlds the narrow content of color-terms is extrinsic. So, it doesn't supervene on the intrinsic make-up of the perceivers.

But there is some reason to think that the characters of noun phrases like 'water' ought not to be considered constant functions. Here is one argument. Chalmers' definition of twin-earthability does not require that the twins use tokens of the same lexical item. However, as it stands, it does not rule it out either. But whether we consider Oscar's and Twin-Oscar's uses of the string of letters 'water' tokens of one word type or different lexical items seems somewhat of a methodological choice (I say 'somewhat' because treating Oscar and Twin-Oscar as speaking the same language allows us to reject the requirement that narrow content supervenes on the intrinsic make-up of particular speakers). Suppose, then, that we consider Oscar's and Twin-Oscar's uses tokens of a single lexical item. Oscar and Twin-Oscar then use twin tokens to pick out different chemical substances. So, the lexical item 'water' then must have a variable character yielding different contents relative to different contexts. Relative to Oscar's context the character of 'water' yields H_2O, and relative to Twin-Oscar's context it yields XYZ. So, the character of 'water' is variable. But if the character of 'water' is variable, then it may reasonably be thought to be in the range of broadly context-sensitive expressions. The same sort of argument applies to other noun phrases typically thought to be twin-earthable.

Granting this point, we can account for the cognitive significance of informative necessities in the following way. Being less than omniscient with respect to empirical matters we may not know whether we are on Earth or Twin-Earth (or some other exotic planet). If we are on Earth, then the character of 'water' determines H_2O. If we are at Twin-Earth, then it determines XYZ. Discovering that water is H_2O here thus amounts to discovering that 'water' determines H_2O here, and hence amounts to discovering what the character of 'water' is in our context.

Of course, someone fond of Kaplan's original semantics might reasonably baulk at this move. Strings of letters forming lexical entries with a constant character, it might be said, simply form different lexical entries with a different constant character when used elsewhere. The phenomenon, they will continue, is akin to that of homonymy. The English word 'red' and the Danish word 'red' [rode] are not one lexical item with variable character but two distinct lexical entries spelled in the same way but with distinct characters. Likewise, given Kaplan

semantics, the Earth word ‘water’ and the Twin-Earth word ‘water’ are distinct lexical items spelled the same way with distinct characters.

Personally I find this sort of response somewhat ad hoc. But granting it, we may reasonably ask whether there is a different way of accounting for cognitive significance within the original Kaplan framework. I think there is a way of partially accounting for the phenomenon. As we will see below, one might define a diagonalized character in the following way. Though the characters of names are constant, the strings of letters forming English expressions could have had different characters. For example, ‘water’ could have been a function from context to XYZ rather than a function from context to H_2O. So, if we take sentence types to be associated with different functions from contexts to characters at different worlds, we can take the diagonal character to be a set of functions from contexts to propositions that are true relative to parameters of the context. For example, the diagonal character of ‘water is H_2O’ yields the proposition that H_2O is H_2O at a context that contains the actual world as a parameter and yields the proposition that the clear liquid that fills oceans, rivers and lakes is XYZ at a context relative to which the character of ‘water’ is a function from context to the descriptive content *the clear liquid that fills rivers, oceans and lakes* and relative to which the character of ‘H_2O’ is a function from context to XYZ. So, it might seem that we can account for the intuitive cognitive significance of ‘water is H_2O’ by noting that the diagonal character yields different propositions at different contexts, and so says different things, relative to different contexts. It is plausible that further investigation into our use of the language is needed in order to discover which non-diagonal character sentences actually have and which propositions they actually express.

Despite their virtues, however, the variable character approach and the diagonal approach to cognitive significance have at least two shortcomings. Both approaches trade on the idea that we are less than competent speakers and hence are not sure about what the characters of our expressions are. But this then has the consequence that ‘water is water’ might be cognitively significant. If we are not certain what the character of ‘water’ is, then we are not certain what the character of ‘water is water’ is. So, while we know that the sentence

expresses necessary propositions relative to contexts, we do not know which proposition it expresses relative to our context.

Furthermore, 'Hesperus is the brightest object in the evening sky' is contingent but uninformative if the term 'Hesperus' is introduced as a name of the brightest object in the evening sky. Yet both its variable character and its diagonal character yield different propositions relative to different contexts. So, 'Hesperus is bright at night' comes out as cognitively informative when it should have come out as cognitively uninformative. So, neither the variable character approach nor the diagonalization approach provides a good account of cognitive informativeness.

There is independent reason to think that even when limited to sentences containing context-sensitive expressions neither one of the attempted Kaplan explanations of cognitive significance is within the spirit of Kaplan's framework. In Kaplan's framework, the proposition determined by a sentence in a context is evaluated relative to a circumstance of evaluation. The parameters of the default circumstance of evaluation are parameters of the context. But circumstance-shifting operators can shift the parameters of the circumstance of evaluation. For example, in the case of 'it was the case that John visited Mary', 'it was the case that' functions as a circumstance shifting operator which shifts the time parameter of the default circumstance of evaluation to some time prior to it. 'It was the case that John visited Mary' is true iff John visited Mary at some time prior to the time of speech. However, according to Kaplan, English contains no corresponding operators that operate on character. There are no monsters or context-shifting operators in English (1989: 510). For example, 'John believes that I am hungry' cannot be used to express the proposition that John believes that he is hungry. 'John believes' does not operate on character and so cannot shift the speaker's context to John's context to produce a John-content for 'I'. If, however, Kaplan's framework could explain informative necessities involving indexicals and demonstratives, then we need to allow that some sentential operators operate on character. To see this, consider:

(1) For all I can rule out from the armchair, I am not Brit.

(1) is a plausible specification of the cognitive significance of 'I am Brit'. In the actual context, of course, 'I' and 'Brit' yield the same propositional contents. So, for (1) to be true, the operator 'for all I can rule out from the armchair' must take us to a context that yields different propositional contents for 'I' and 'Brit'. So, the English operator 'for all I can rule out from the armchair' must function as a monster. So, if we extend Kaplan's framework in the ways suggested above, then there are monsters after all, contrary to what Kaplan claimed.

2. Stalnaker

2.1. Assertion

In "Assertion" Robert Stalnaker offers a theory of content that is similar in some respects to Kaplan's. However, where Kaplan opts for structured propositions, Stalnaker takes a proposition to be a set of possible worlds that are the way the world is said to be. Stalnaker thinks of possible worlds as the conditions that must obtain for the proposition to be true, hence they are truth conditions. This idea is motivated by the thought that when one asserts a proposition, one seeks to eliminate the set of possible situations at which what one says is false. Unlike Kaplan, Stalnaker seeks to account for cognitive significance or 'the information that a statement conveys' in those cases in which what one says is informative despite being necessarily true. Necessarily true statements do not exclude any possibilities, yet if one utters an informative necessity, then it seems that one is seeking to rule out possibilities.

In "Assertion" Stalnaker's account of cognitive significance is similar to the diagonalization approach I imaginatively outlined above on behalf of Kaplan. For Stalnaker, utterances are associated with propositional concepts, functions from possible worlds containing the utterance to propositions. The diagonal proposition is the set of worlds at which the utterance's propositional concept yields a true proposition. If, for example, context maps 'water' to the descriptive material *the clear liquid that fills oceans, rivers and lakes* and maps 'H_2O' to XYZ, then the diagonal proposition of 'water is H_2O' is the proposition that the clear liquid that fills oceans, rivers and lakes is XYZ. Cognitive informativeness is explained by our

uncertainty about which context we occupy and hence which proposition is in fact expressed by our assertions.

This approach generates the same problems as diagonalizing on character does. We might be uncertain about which proposition is in fact expressed by the a priori necessity 'water is water' and the a priori contingency 'Hesperus is bright at night', in which case these sentences are cognitively informative despite being a priori.

Of course, our uncertainty might be limited in various ways. We might know 'water' does not refer to tigers but refers to a chemical substance. Likewise, we might know that 'water = water' expresses a necessary proposition even if don't know which. Still, the approach is bound to yield some intuitively mistaken results.

The early Stalnaker approach has other virtues. It can explain communication failure. If A is identical to B, but I lack competence, I may fail to grasp that 'A is F' and 'B is F' express the same proposition. If you are competent and believe I am too, communication failure might ensue. Stalnaker's early account can explain why. However, the early Stalnaker account does not yield a good account of informative necessity and uninformative contingency.

2.2. Assertion Revisited

In "Assertion Revisited" Stalnaker proposes to treat Kaplanian character as a kind of narrow content but understood metasemantically (i.e., not as content proper). He also proposes to treat a wider class of expressions as context-sensitive. Nearly all names and descriptive expressions are treated as having a generalized variable character. For example, 'Socrates' and 'water' are treated on the model of indexicals. So, their character is a variable function mapping context to individuals. The generalized character of 'Socrates lived in Athens' is a variable function mapping context to propositions (or sets of worlds). In this new framework, it is the variability of the character of 'water is H_2O' that explains its cognitive informativeness. Stalnaker furthermore assumes that there may be cases in which speakers know what a sentence says. So, a sentence may be informative relative to one speaker but not relative to another.

Stalnaker's new approach, of course, is similar in a number of ways to the first approach I outlined above on behalf of Kaplan. However, it differs from this approach in treating some expressions which intuitively are not twin-earthable as context-sensitive.

Like the variable character approach we considered above on behalf of Kaplan, Stalnaker's new approach has difficulties distinguishing between the differences in the cognitive value of 'water is water' and 'water is H_2O', and 'Venus is bright at night' and 'Hesperus is bright at night'. If 'water is H_2O' is informative in virtue of the fact that we don't know which content its variable character determines in our context, then so is 'water is water'. Likewise, if 'Venus is bright at night' is informative because we don't know what content its variable character determines in our context, then so is 'Hesperus is bright at night'.

Stalnaker's proposal has further problems. One potential problem is over-generation. If the expression 'The Mayor of Boston' is treated on a par with 'I' and 'now', which the framework seems to allow, then it has a variable character which maps the expression from context to an individual. So, 'the Mayor of Boston is rich' and 'Thomas M. Menino is rich' have the same content (though different characters). This is not a super-controversial proposal. It familiarly goes against Russell's theory of descriptions, according to which descriptions are incomplete symbols which do not denote. But it is in broad agreement with Peter Strawson's proposal that descriptions are referential and Keith Donnellan's proposal that descriptions have attributive and referential uses. Elsewhere both Stalnaker and Kaplan have welcomed this consequence of an extended two-stage approach to semantics.[6]

Furthermore, starting with Strawson many have thought that incomplete descriptions must be treated as context-sensitive. Jason Stanley and Zoltan Szabo, for example introduce nominal restriction to account for the context-sensitivity of quantifiers. If I were to say that every bottle is on the table, I would not ordinarily mean that every bottle in the universe is on the table. What I would mean in the envisaged circumstance is that every <bottle, i> is on the table, where 'i' is contextually completed. So, '<bottle, i>' denotes, say, the set of bottles in the kitchen. The problem of incomplete descriptions can possibly be resolved in the same way (see e.g. Stanley 2002a). 'The book is on the table' can be treated as being of the form 'the <book, i>

[6] See Stalnaker (1970) and Kaplan (1978).

is on the <table, j>', where '<book, i>' and '<table, j>', once the indexical variables are contextually completed, denote, say, the set of books I have in mind and the set of tables in the living room, respectively.

However, granting that incomplete descriptions are context-sensitive does not vindicate a treatment of characters as generalized. Nominal restriction, for example, does not require treating characters as generalized. It does admittedly require treating the restrictor-plus-hidden variable as having a variable character that returns a property or an extension relative to context (e.g., the property of being a bottle in my kitchen or the set of bottles in my kitchen). But this character variability is arguably just of the sort postulated in the original Kaplan text. Nominal restriction does not indicate a need for a more generalized treatment of character.

Moreover, without further constraints on which expressions can have variable character, it would seem that the generalized proposal extends to other quantificational expressions. For example, 'every rich person from Boston' might have a variable character that takes the expression from a context to a set of people. A referential treatment of universally quantified expressions seems less attractive than a referential treatment of definite descriptions.

The generalized character proposal also threatens to turn intuitively non-twin-earthable (or semantically neutral) expressions like 'consciousness', 'friend' and 'cause' into context-sensitive expressions with variable characters. For example, it might be thought that, given some functional description of consciousness, we can define a variable character that yields different properties relative to different contexts. So, relative to the actual world, 'is conscious' and 'is phenomenally conscious' might have the same content, whereas at a zombie world 'is conscious' and 'is able to report part of the content of their computational states' might have the same content. Without further constraints in place, Stalnaker's new proposal seems to threaten to make too many expressions context-sensitive. In response to this sort of worry Stalnaker suggests treating fundamental, natural properties and relations as context-insensitive (and hence as having a constant character).

It might be thought that a limited extension of the variable character approach would be a good way to account for context-sensitivity more generally. For example, it might be

thought that expressions such as 'local', 'nearby', 'tall', 'big', and so on, which beg to be treated as context-sensitive, could be treated as having themselves variables characters rather than as being associated with a hidden constituent with a variable character.[7] But the motivation for this approach is somewhat meager, given that these expressions can be treated, on the model of nominal restriction, as containing a hidden indexical variable (see e.g. Stanley 2002).

Philosophers have in recent years argued that the range of context-sensitive expressions must be widened to include epistemic expressions (e.g. 'know' and 'might'), moral expressions (e.g., 'decent' and 'appropriate'), predicates of personal taste (e.g. 'fun' and 'tasty') and vague expressions (e.g., 'bald' and heap'). If they are right about this, then these expressions may well be best treated on the more generalized variable-character model, though it is possible that the less explored options in these areas of philosophy, viz. the hidden indexical-variable approach or the epistemic two-dimensional approach, ultimately can provide a better account of the contextual nature of these expressions.

3. Chalmers

3.1. Epistemic Two-Dimensionalism

David Chalmers and Frank Jackson have proposed a different two-dimensional framework, according to which epistemic variability and cognitive significance are to be explained at least partially in terms of a basic notion of apriority ('partially' because they recognize that a priori sentences may be cognitively informative to non-ideal reasoners). There are some differences between their approaches. Here I shall focus on Chalmers' framework.

Chalmers offers two versions of his epistemic two-dimensional framework, one which treats the space of possible worlds as plentiful enough to provide a model for deep epistemic

[7] The literature on tense is too big to deal with in this paper but the issues raised in that literature are obviously relevant to the debate about two-dimensional semantics. In traditional tense logic, the tenses were treated as sentential operators, and that is just how Kaplan treats them. In recent years, however, the tenses have more frequently been treated as indexical referential expressions or as indexical quantificational expressions.

possibility (i.e., possibilities not ruled out a priori), and one which constructs epistemic space out of sentences. Let us look at the former proposal first. On the former proposal, there is a space of possible worlds which allows us to define 2-intensions (Kaplan contents or intensions) as either sets of worlds or as functions from worlds to truth-values. 2-intensions are necessarily true in the standard sense when they yield the truth-value true at every world and necessarily false when they yield the truth-value false at every world. On the plenitude proposal, however, the space of possible worlds is large enough to model deep epistemic possibility. A scenario is a possible world in which certain features are marked: in most cases an individual and a time. For every centered world, there is then a maximal hypothesis about the world in question expressed in a canonical language. A canonical language is a semantically neutral language. To a first approximation, an expression is semantically neutral just in case it is not twin-earthable as defined above.[8] Part of the hypothesis about a centered world might be that some heavenly body called 'Venus' is the brightest object in the evening sky and that some heavenly body called 'Venus' is the brightest object in the morning sky. As 'Venus' is not semantically neutral, it is not part of the description that Venus is the brightest object in the evening sky.

The important principle on the plenitude conception of deep epistemic possibility is what Chalmers calls 'Metaphysical Plenitude':

> Metaphysical Plenitude: For all, S, if S is a priori possible, there is a centered metaphysically possible world that verifies S.

Suppose it is not ruled out a priori that Hesperus, the brightest object in the evening sky, is not identical to Phosphorus, the brightest object in the morning sky, despite the fact that it is

[8] Matters are a bit more complicated than this. Semantically neutral concepts are never twin-earthable. But there are some non-twin-earthable concepts that are not semantically neutral. Chalmers offers the following example: Let L be an expression that rigidly designates the speakers height. Then L is non-twin-earthable but it is not semantically neutral.

necessary that Hesperus is identical to Phosphorus. It is then deeply epistemically possible that Hesperus is not identical to Phosphorus. While there is no possible world in which Hesperus is not identical to Phosphorus, there is a canonical description of a centered world which a priori implies that Hesperus is not identical to Phosphorus. This description might, for example, be a description of a centered world in which Jupiter is the brightest object in the evening sky and Mars is the brightest object in the morning sky. As Jupiter is not identical to Mars, the world verifies the sentence 'Hesperus is not identical to Phosphorus'.

On the alternative picture of scenarios, a scenario is an equivalence class of epistemically complete sentences in a canonical language. A specification of a scenario is a sentence in its equivalence class. The rest of the two-dimensional apparatus goes as before. On the constructive treatment of scenarios, one can in principle allow for strong necessities (i.e., deeply epistemically contingent necessities involving semantically neutral terms).

Given either framework, the 1-intension of an expression is defined as a function from scenarios to extensions. A priori (and uninformative) sentences have epistemically necessary 1-intensions, whereas a posteriori sentences have epistemically possible 1-intensions.[9] So, given this framework 'water is H_2O' has a necessary 2-intention that yields the truth-value true at every possible world, but it has a contingent 1-intension that yields the truth-value true at some scenarios and the truth-value false at some scenarios. For example, it yields the truth-value false at scenarios that have a canonical description which a priori implies that the clear liquid that flows in rivers, oceans and lakes is XYZ. All a posteriori necessities, in Kripke's sense, have a 1-intension that comes apart from their two-intension.

[9] Matters are a bit more complicated. An enriched proposition (i.e., a Russellian proposition and a structured primary intension) is a priori at a world w iff the primary intension determined by that proposition is necessary at w, and the proposition is entertainable at w (Chalmers forthcoming b, note 18). An enriched intension is live at a scenario v iff the Russellian component has the same extension as the primary component at v. An enriched proposition is entertainable at a world w iff each of its components is live at a scenario corresponding to w. This additional complication is introduced to avoid that Apriori(P) entails []Apriori(P). For suppose this holds. Then for any a priori truth P, we have []Apriori(P). But by theorem T of system T, Apriori(P). Apriori is factive. So, P. By (restricted) necessitation, necessarily, P. But this result is unfortunate when P is a *contingent* a priori truth.

Chalmers' two-dimensional framework does not inherit the problems of the earlier frameworks. But its explanatory power has some limitations. The framework explains cognitive informativeness for ideal reasoners, but as Chalmers recognizes, it does not explain informativeness for non-ideal reasoners. Despite the fact that all mathematical and logical truths are a priori, many true mathematical conjectures are cognitively informative. Moreover, because the space of scenarios is the space of deep epistemic possibilities, the space of scenarios cannot be used as a way to model more hyper-intensionality more generally, for instance, the belief contexts and strict epistemic possibility contexts. For example, 'I believe that p and not-p' might be true, but there is no scenario that verifies p & not-p. However, as Chalmers (forthcoming) observes, the framework of scenarios can be extended to account for both non-ideal cognitive informativeness and hyperintensionality.

In Chalmers' two-dimensional semantics, scenarios play roughly the role contexts play in Kaplan's framework, narrow contents (1-intensions) play roughly the role that characters play, except that they are part of the semantics proper ("the propositional content") and wide contents (2-intensions) play roughly the role that Kaplan contents play.

Despite the analogy between these approaches, however, there are significant differences. Chalmers explicitly distinguishes epistemic variability from context-sensitivity. An expression is epistemically variable if it has a non-constant 1-intension, whereas an expression is context-sensitive if it has a non-constant character. Epistemic variability is structurally analogous to standard context-sensitivity in certain respects but is conceptually distinct.

Given Chalmers' framework, standard context-sensitive expressions (e.g. pure indexicals and true demonstratives) have 1-intensions that come apart from their 2-intensions. Their 1-intensions are non-constant functions from the marked individuals and times (or features of the individuals) in the centers of scenarios to extensions. Their 2-intensions are constant functions from the speaker's scenario to extensions. For example, the 1-intension of a token of 'I' is a non-constant function from scenarios to the individual in the center. The 2-intension of a token of 'I' is a constant function from the speaker's scenario to the speaker. Likewise, the 1-intension of a token of 'this' is a non-constant function from scenarios to the object

demonstrated by the individual in the center. The 2-intension of 'this' is a constant function from the speaker scenario to the object demonstrated by the speaker.

Context-sensitive expressions are twin-earthable. Oscar and Twin-Oscar's tokens of 'I' have the same 1-intension but different 2-intensions. Other twin-earthable expressions, of course, also have a 1-intension that comes apart from their 2-intension (e.g. 'water', 'Hesperus', etc). But Chalmers treats these other twin-earthable expressions as epistemically variable, not as context-sensitive.

3.2. Twin-Earthability and Context-Sensitivity Broadly Speaking

Despite the differences between epistemic variability and context-sensitivity within the original epistemic two-dimensional framework, I believe one could understand epistemic variability as a kind of context-sensitivity. This proposal is similar to the variable character approach I sketched above on behalf of Kaplan but it avoids some of the most obvious problems with the earlier suggestions. It takes context-sensitivity, broadly construed, to be grounded in twin-earthability. Chalmers, of course, would not endorse this. But here is some motivation for this line of thought. Suppose we have a more localized Twin-Earth phenomenon. Water in and around Australia is XYZ, whereas water in around America is H_2O. It now seems possible to treat the 1-intension of 'water' as a function from locations of individuals in scenarios to chemical substances found on those locations and the 2-intension of 'water' as a function from the speaker's location to the chemical substance found on that location.

Of course, there is still the option of denying that the string of letters 'water' forms is a single lexical item. One could hold that we have here a case of homonymy or polysemy. If this were so, then 'water' as used by speakers in Australia and 'water' as used by speakers in America would be different lexical entries with the same spelling.

But this rejoinder can be empirically falsified. For example, we might discover that competent American and Australian speakers treat the word 'water' in the same way that they treat the word 'I', not as string of letters spelling different words, but as a single lexical entry

and that they are oblivious to the context-sensitivity of the word. To illustrate consider the following dialogue:

(A)

American: I am hungry

Australian: No, you are wrong, I am not hungry.

Just as competent speakers treat (A) as revealing a failure to recognize that 'I' is context-sensitive rather than a failure to recognize that the speakers use two different homonymous words, so we can imagine that competent speakers would treat the following dialogue as revealing a failure to recognize that 'water' is context-sensitive rather than a failure to recognize that the speakers use two different homonymous words:

(B)

American: Water is H_2O

Australian: No, you are wrong, water is XYZ

To say that context-sensitivity, broadly construed, is grounded in twin-earthability is not to deny that there are important differences between, say, indexicals and noun phrases. There is, for example, a difference in how we normally use these two types of expression. As we normally use 'I', its 1-intension is a function from centered worlds to individuals in the center. As we normally use 'water', its 1-intension is, roughly, a function from scenarios to whichever substance satisfies the description 'the clear potable liquid that comes out of our kitchen faucets'.

But it is not clear how much weight this difference carries. First, like the 1-intension of 'water', the 1-intension of 'I' is, roughly, equivalent to the 1-intension of a description, perhaps, 'dthat [the person who utters this token]' (Kaplan 1989: 522). For example, the 1-intension of 'I am not here' is plausibly equivalent to the 1-intension of 'dthat [the person who utters this token] is not here'.

Second, it is not difficult to imagine that 'water' could have the 2-intension it actually does and yet have a 1-intension that is definable in much the same way as the 1-intension of 'I'. Suppose water is person-dependent and happens to have a person-specific chemical structure. For example, chemical structures might be causally connected to the individual essences of people. So, when you pour yourself a glass of clear odorless liquid, the chemical structure is H_2O, and when I pour myself a glass of clear odorless liquid the chemical structure is XYZ. In the envisaged circumstances we might be using 'water' with a 1-intension which, at each scenario, picks out the chemical substance at the center of the scenario. The fact that this 1-intension, even in this scenario, very well could be (roughly) equivalent to the 1-intension of 'the clear, odorless liquid that comes out of the faucet in my kitchen' only emphasizes my point that there is no obvious reason to treat indexicals as context-sensitive and twin-earthable noun-phrases as context-insensitive.

However, as I mentioned above, Chalmers would not endorse this proposal. Moreover, there is an interesting difference between how indexicals and noun phrases are actually used. Presumably Kaplan was right in thinking that there are no operators in English that can change the context to yield a different content for 'I'. For example, 'John believes I am hungry' cannot be used to say that John believes that John is hungry. However, it is possible that there are some monsters in English that can change the context to yield a different content for noun phrases. Here is how:

If twin-earthable expressions can be treated as a kind of context-sensitive expressions, then whenever we have an expression that has a twin token which yields a different 2-intension elsewhere, we have context-sensitivity. If, now, we can have operators relocating speakers from a place where the expression has one 2-intension to a place where it has a different 2-

intension, then we have an operator on context, hence a monster in Kaplan's sense. It might plausibly be thought that there are operators of this kind. Consider a color term like 'redness'. Let's suppose, for simplicity, that its 1-intension picks out a reflectance type in the actual world but some other property in other worlds. For example, in worlds in which colors are purely qualitative properties instantiated by external objects (as color primitivism would have it), the 1-intension of 'red' picks out these purely qualitative properties. Now consider:

(2) In Eden objects instantiate redness in virtue of instantiating a purely qualitative or 'primitive' color property.

If 'redness' as it occurs in (2) picked out a reflectance property, (2) would be false. But (2) looks true. So, it must be that 'In Eden' is capable of changing the scenario (or centered world) from the actual one to an Edenic one. So, to the extent that scenarios are contexts in a broad sense 'In Eden' is a context-shifting operator.

Other cases of context-shifting are more controversial. Consider the following conditional:

(3) Should this turn out to be Twin-Earth, water is not H_2O.

Intuitions seem divided on whether (3) is true or false. I am inclined to think it is true (given the fiction). There are numerous ways to deal with conditionals. But on one plausible account the antecedent 'should this turn out to be Twin-Earth' functions as a context-shifting operator which shifts the speaker's context (i.e., scenario) from Earth to Twin-earth (either by shifting the world or the location within the world). Given a variable character-approach to names, the character of 'water' is something like a description which, relative to Twin-Earth, determines that the content of 'water' is XYZ, not H2O. Given a version of epistemic two-dimensionalism, it is plausible to think that the antecedent 'Should this turn out to be Twin-Earth' shifts us to a

different centered world (or scenario or context) that has a specification that a priori implies that water is XYZ. So, while the 1-intension of 'water' stays the same, the 1-intension now determines a 2-intention that picks out XYZ at every possible world. If Twin-Earth is a planet in our universe, the operator 'On Twin-Earth' shifts the speaker's scenario to a scenario with a different center, viz., a center occupied by Twin-Oscar. The 1-intension of 'water' stays the same, but because we have a new center, the 1-intension determines a different 2-intension that picks out XYZ at every possible world.

So, given a reasonable extension of epistemic two-dimensionalism, some operators in English (in the broad sense of 'operator' that includes antecedents of conditionals) can perhaps be thought of as context-shifting operators which shift context (i.e., the scenario) to determine a new content (2-intension) for broadly context-sensitive expressions (viz. noun phrases).

It may be argued that there are no operators like 'In Eden' or 'Should this turn out to be Twin-Earth' in English. However, I think English is expressively rich, and we can felicitously utter sentences like (2) and (3). Whether they have true readings is an empirical question. If they do, then there are context-shifting operators, in the broad sense, in English.

I also note that while Chalmers does not endorse these considerations and does not propose to treat 1-contingency in terms of context-shift, the foregoing considerations suggest a way for the two-dimensional approach to be understood this way. If (3) is true in my mouth and we evaluate whether (3) is 1-contingent or not, it might reasonably be suggested that we evaluate the sentence within the scope of an envisaged context-shifting operator, for instance, 'Should this turn out to be Twin-Earth' or 'Given hypothesis H'. On this approach then, an evaluation of a sentence's 1-modal features requires shifting the scenario (i.e., the context) either by shifting the world part of the scenario or by shifting the center of the scenario. We then look to see what content is determined given this context-shift. I believe this way of thinking about two-dimensional semantics is broadly within the spirit of Chalmers' approach.

4. Dynamic Two-Dimensional Semantics

As it stands, two-dimensional semantics does not allow for a treatment of dynamic discourse. Consider the following discourse fragment:

(4) John is now entering the room, and he is now taking off his hat and is therefore not now wearing a hat, but he is now putting the hat back on, and is therefore now again wearing a hat.

If the conjuncts in (4) are thought of as parallel-asserted, then (4) is contradiction, not so if (4) is asserted at the right sort of pace in non-parallel fashion. Call a non-parallel assertion a 'dynamical assertion'. So, a parallel-assertion of (4) and a dynamic-assertion of (4) have different intensions.

The fact that discourse fragments like (4) exist in natural language suggest a need for a dynamic two-dimensional semantics. Given length considerations I can only briefly sketch how such an approach might proceed. Consider:

(5) John$_1$ is now$_1$ spotting Susan$_2$.

Following Irene Heim, let us introduce the notion of a filing system, that is, a system that keeps track of variables, names, and descriptive material introduced by the discourse. Here is an illustration:

Filing system F1:

x, y, t_1

Now t_1

John x

Susan y

Spot (x, y)

Additions to the discourse give rise to a new system:

(6) He_1 is now_2 walking over to her_2.

Filing system F2:

x, y, t_1	x, y, t_2
Now t_1	Now t_2
John x	
Susan y	
Spot (x, y) ➔ Walk over to (x, y)	

(7) And is now3 starting a conversation with her2

Filing system F3:

x, y, t_1	x, y, t_2	x, y, t_3
Now t_1	Now t_2	Now t_3
John x		
Susan y		
Spot (x, y) ➔	Walk over to (x, y) ➔	Start a conversation with (x, y)

(8) She_2 is now_4 talking to a man_3.

Filing system F4:

x, y, t1	x, y, t2	x, y, t3	x, y, t4
Now t1	Now t2	Now t3	Now t4
John x			
Susan y			
Man z			

Spot (x, y) ➔ Walk over to (x, y) ➔ Start a conv with (x, y) ➔ Talk to (x, z)

(9) Now_5 he_1 is talking to the man_3 she_2 talked to just a moment ago_4

Filing system F5:

x, y, t_1	x, y, t_2	x, y, t_3	x, y, z, t_4	x, y, z, t_5
Now t_1	Now t_2	Now t_3	Now t_4	Now t_5
John x				
Susan y				

Man z

Spot (x, y) ➔ Walk over to (x, y) ➔ Start a conv with (x, y) ➔ Talk to (y, z) ➔Talk to (x, z)

We can introduce a notion of truth of files as follows. Given an understanding of scenarios as centered possible worlds, we can take 1-models to be pairs of a domain D of actual individuals, some of which serve as a center and an interpretation function I.

M = <D, I>

The 1-interpretation function I maps the non-indexical variables (i.e., discourse referents) of file F to members of D, indexical variables of F to the center of D, and the descriptive material and names (i.e., predicate nominals) of F to properties or relations (or sets of n-tuples) on D. One set of assignments, $i^{@1}$, contains as an element the distinguished interpretations: the actual sets of assignments of extensions to the expressions (or the set of actual centered worlds). Let a 1-assignment A in a model M = <D, I> be a mapping of non-indexical variables onto elements of D, and indexical variables onto the centered elements of D. We can then say that assignment A verifies filing system F in M if there is an extension E of A such that the elements satisfy the descriptive material and names at the scenario. A 1-information state is a space of scenarios (interpretations) with the same filing system. A dynamic 1-intension is a sequence of 1-information states. We can then say that discourse fragments express dynamic 1-intensions which are sequences of spaces of scenarios that share a filing system. If we introduce a designated sequence of scenarios, we can then say that a dynamic 1-intension is true at a designated sequence of scenarios just in case each scenario in the designated sequence of scenarios is in the space of scenarios (or interpretations) with the same filing system as the designated scenario.

A different model is needed to model dynamic 2-intensions. A 2-model is a pair of an uncentered domain of actual individuals and an interpretation function. The 2-interpretation function I maps the indexical and non-indexical variables and the associated names (if any) of

file F to members of D, and the descriptive material of file F to properties or relations (or sets of n-tuples) on D. One assignment, $i^{@2}$ is the distinguished interpretation, the actual assignment of extensions to the expressions (or the actual world). Let an assignment A in a model M = <D, I> be a mapping of variables and associated names (if any) onto elements of D. We can then say that assignment A verifies filing system F in M if there is an extension E of A such that the elements satisfy the descriptive material at the scenario.

A 2-information state is a space of worlds with the same filing system. A dynamic 2-intension is a sequence of 2-information states. A dynamic 2-intension is true at a distinguished sequence of worlds just in case each world in the designated sequence of worlds is in the space of worlds with the same filing system as the designated world.

We can give the following validation conditions for operators:

For universal operators:

$\Box^h$, $\Box^h\Phi$ is true just in case at every h-admissible interpretation i, i assigns true to Φ.

So, 'It is metaphysically necessary that Φ' is true iff at every 2-interpretation i, i assigns **T** to Φ, and 'It is a priori that Φ' is true iff at every 1-interpretation i, i assigns **T** to Φ. To illustrate consider:

(10) The brightest heavenly object is now_1 $Hesperus_1$, and the brightest heavenly object is now_2 $Phosphorus_2$, and $Hesperus_1$ is identical to $Phosphorus_2$.

(11) It is deeply epistemically possible that: the brightest heavenly object is now_1 $Hesperus_1$, and that the brightest heavenly object is now_2 $Phosphorus_2$, and that $Hesperus_1$ is not identical to $Phosphorus_2$.

(10) is true iff the dynamic 2-intension it expresses has the truth-value true at the designated sequence of worlds. So, (10) is true if the brightest heavenly object is Hesperus at the worlds

that assign the utterance time to the variable associated with 'now_1', and the brightest heavenly object is Phosphorus at a designated world that assigns the utterance time to the variable associated with 'now_2', and Hesperus is identical to Phosphorus. (11) is true iff there is a sequence of scenarios at which the dynamic 1-intension of the embedded clause is true. This is so if the brightest heavenly object is the brightest object in the evening sky at the scenarios that assign the time in the center to the variable associated with 'now_1', and the brightest heavenly object is the brightest object in the morning sky at the scenarios that assigns the time in the center to the variable associated with 'now_2', and the brightest object in the evening sky in scenarios in the first set is not identical to the brightest object in the morning sky in the scenarios in the second set.

5. Conclusion

As we have seen, in Kaplan's original framework, context-sensitive expressions are expressions with a variable character—a linguistic meaning that determines different contents relative to different contextual parameters. This approach to context-sensitivity has much to recommend it but it fails to capture a plausible connection between context-sensitivity, broadly understood, and cognitive informativeness. The two phenomena are not unrelated. Cognitively informative necessities arguably contains what David Chalmers calls 'twin-earthable expressions', linguistic strings that have twin tokens with a different content. Twin-earthable expressions are broadly context-sensitive. When I use 'water', the term picks out H2O. When my twin on Twin-Earth uses 'water', it picks out XYZ. On the assumption that we speak the same language, an assumption which is consistent with the definition of Twin-Earthability, 'water' has a variable character and hence is broadly context-sensitive. But if 'water' is context-sensitive, we have a partial explanation of the cognitive significance of sentences such as 'water is H2O'. The sentence, relative to my context, expresses a proposition of the form 'a = a', yet owing to my cognitive limitations I may not know that what content is determined the character of 'water' at my location. Being told that the character of 'water' determines H2O is

informative for me. Thus, if we keep fixed the language spoken, there is a tight connection between twin-earthability and cognitive informativeness.

I have assessed three types of two-dimensional semantic frameworks in terms of how well they account for the connection between cognitive significance and the broader notion of context-sensitivity. One is a diagonalized character approach that takes names and kind terms to have a constant character but allows for the possibility that a diagonal proposition that functions much like a variable character has some explanatory power with respect to the phenomenon of cognitive informativeness. A second approach is to widen the range of expressions that have a variable character to include various noun phrases and descriptions. A third approach is to treat expressions as having epistemic meaning in addition to semantic meaning. I argued that while all three approaches go some way toward explaining the connections between context-sensitivity and cognitive informativeness, the non-epistemic approaches seem to fall short of offering a fully adequate account of cognitively informative necessities and informative contingencies. I concluded by pondering a new problem for the epistemic approaches posed by indexicals in dynamic discourses and sketched a solution.[10]

References

Chalmers, D. 1996. *The Conscious Mind: In Search of a Fundamental Theory*, Oxford: Oxford University Press.

Chalmers, D. 2002: "On Sense and Intension", *Philosophical Perspectives* 16: 135-82.

Chalmers, D. 2006. "Foundations of Two-Dimensional Semantics", in Garcia-Carpintero, M. & Macia, J., eds., *Two-Dimensional Semantics: Foundations and Applications*, Oxford: Oxford University Press: 55-140.

[10] I am grateful to David Chalmers and Melissa Ebbers for helpful comments on an earlier version of this paper.

Chalmers, D. Forthcoming a. “The Nature of Epistemic Space”, in A. Egan and B. Weatherson, ed. *Epistemic Modality*, Oxford: Oxford University Press.

Chalmers, D. Forthcoming b. “Propositions and Attitude Ascriptions”, *Nous*.

Chalmers, D., and Jackson, F. 2001. “Conceptual analysis and reductive explanation”, *Philosophical Review* 110: 315-360.

Donnellan, K. 1966. “Reference and Definite Descriptions”, *Philosophical Review* 77: 281-304.

Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*, doctoral dissertation, Amherst, MA: University of Massachusetts.

Kaplan D. 1978, "DTHAT", *Syntax and Semantics* 9: 221-243.

Kaplan, D. 1989, “Demonstratives”, in J. Almog, H. Wettstein, and J. Perry (eds), *Themes from Kaplan*, New York: Oxford University Press: 481-563.

Lasersohn, P. 2009. “Relative Truth, Speaker Commitment, and Control of Implicit Arguments, in B. Brogaard, ed. *Relative Truth*, special issue of *Synthese* 166: 359-374.

Russell, B. 1905. “On Denoting”, *Mind* 14: 479-493.

Stalnaker R. 1970. “Pragmatics”, *Synthese* 22: 272-289

Stalnaker, R. 1978. “Assertion”, *Syntax and Semantics* 9: 315–32

Stalnaker , R. 1999. *Context* and *Content: Essays on Intentionality in Speech and Thought*, Oxford: Oxford University Press

Stalnaker, R. 2004. “Assertion Revisited, on the interpretation of two-dimensional modal semantics”, *Philosophical Studies 118*: 299-322

Stanley, J. 2002. “Nominal Restriction”, in G. Preyer and G. Peter, eds., *Logical Form and Language*, Oxford: Oxford University Press, 365-88.

Stanley, J. and Z. Szabó. 2000. “On Quantifier Domain Restriction”, *Mind and Language* 15: 219-61.

Strawson, P. F. 1950. “On Referring”, *Mind* 59: 320-344.